Volume Title: Quantifying Systemic Risk

Volume Author/Editor: Joseph G. Haubrich and Andrew W. Lo, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-31928-8; ISBN-13: 978-0-226-31928-5

Volume URL: http://www.nber.org/books/haub10-1

Conference Date: November 6, 2009

Publication Date: January 2013

Chapter Title: Introduction to "Quantifying Systemic Risk"

Chapter Author(s): Joseph G. Haubrich, Andrew W. Lo

Chapter URL: http://www.nber.org/chapters/c12066

Chapter pages in book: (p. 1 - 10)

Introduction

Joseph G. Haubrich and Andrew W. Lo

Introduction

In the wake of the financial crisis of 2007 through 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

Joseph G. Haubrich is a vice president of and an economist at the Federal Reserve Bank of Cleveland. Andrew W. Lo is the Charles E. and Susan T. Harris Professor, professor of finance, and director of the Laboratory for Financial Engineering at the Massachusetts Institute of Technology. He is a research associate of the National Bureau of Economic Research.

The views expressed here are those of the authors only, and do not represent the views of the Federal Reserve Bank of Cleveland, the Board of Governors of the Federal Reserve System, AlphaSimplex Group, MIT, the National Bureau of Economic Research, or their affiliates and employees. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see http://www.nber.org/chapters/c12066.ack.

shareholders 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.

However, the challenge is not just measuring systemic risk, but also implementing it within the existing regulatory infrastructure; at issue is institutional design as well as statistical inference. After all, the ultimate goal is not just prediction, but also prevention—failing that, intervention to mitigate the severity of an impending crisis. Achieving this lofty goal requires detailed knowledge of the dynamics of the financial sector.

The technical side of risk measurement has received the most attention, particularly from academics, but risk management involves three distinct elements according to Lo (1999)—prices, preferences, and probabilities—and we can frame the discussion of systemic risk in a similar fashion. Centuries of work by scientists and mathematicians have advanced the understanding of probability, but the practical difficulties in estimating the distribution of financial market data remain formidable. Estimating extreme events from everyday behavior can be seriously misleading. The space shuttle booster O-rings performed acceptably in cool temperatures, but failed disastrously in the freezing conditions of the *Challenger* launch on January 28, 1986 (Tufte, 1990). Though not as rare as we once thought, financial crises remain extreme events. Making the problem even harder, figuring the odds means aiming at a moving target. The past twenty-five years of finance have stressed how changes in variance affect stock prices, interest rates, and spreads (Engle, 2001). Recent financial crises, from the sovereign defaults of the late 1990s to the Panic of 2007 to 2009 to the current problems in Europe highlight how quickly the correlations between different investments can change, encapsulated in the folk wisdom that "in a crisis, all correlations go to one."

Decisions also require a way to rank different risks—investors, even regulators, must confront their own (or society's) preferences, which are inevitably subjective. Exactly how does a particular investor value different payoffs and probabilities? Still, as an aid to decision making, a variety of objective measures have been proposed, mainly variations of statistical concepts used to describe the "risk." These include traditional measures such as mean and variance, along with the various flavors of the popular value at risk (VaR) measure such as expected shortfall. Other notions such as stochastic dominance and its extensions, for example, the economic risk measure of Aumann and Serrano (2008) or the operational measure of Foster and Hart (2009), provide a way to think about the trade-off between risk and return but must ultimately involve preferences. A more mathematical approach postulates a set of axioms that any "good" risk measure must obey. As might be expected, however, different axioms can produce very different risk measures, producing such varying measures as expected utility, coherent measures of risk, or uncertainty aversion. Regulators face the problem on a higher level, seeking to implement the trade-offs that society prefers.

Finally, prices play a dual role in thinking about risk, as both the input and the output of the process. Price movements—the profits and the losses—drive the need for hedging and risk management. At the same time, prices are the output, the outcome of supply and demand expressed through preferences and probabilities. But it is when prices do not properly capture the economic value of the corresponding commodity that problems arise. As with air pollution, systemic risk arises when market prices do not reflect the full impact of a firm's decisions on the rest of the economy. This creates the need for something beyond business as usual.

Perhaps in an ideal world, market discipline alone would induce firms to measure and manage risk properly. But systemic risk, like other negative externalities, means that individual firms do not consider how their actions affect the system as a whole. There is a sense that we are not starting from the Garden of Eden. Safety nets such as deposit insurance or implicit too-big-to-fail policies reduce the incentive to manage risk. How supervision best responds is another matter. Recently, the conversation has shifted from the safety and soundness of individual banks to the appropriate level of "macroprudential" supervision; that is, the total risk in the system. An early proponent of macroprudential supervision, Claudio Borio of the Bank for International Settlements explains it as having both a cross-sectional (distribution of risk) and a time series (change over time) dimension (Borio 2003). Other regulators have also voiced the intent to make regulation more macroprudential (Tarullo 2010). One output of this philosophy is the Basel III proposal to require globally systemically important banks (G-SIBs) to hold additional capital buffers (BCBS 2011).

Clearly, macroprudential regulators need valid measures of systemic risk. Operationally, they need these measures to set capital requirements and to consider other aspects of supervision such as merger decisions. Measures that are not available on a timely basis, difficult to interpret, or easily manipulated are of little use. But if risk measurement needs change, so to do regulators. This "changing face of supervision" (FRBC 2010) is, in fact, becoming apparent in both the skill sets and organization of supervisors and regulators. One example is the horizontal reviews conducted by the Federal Reserve System and the Financial Stability Board, which has introduced cross-functional, horizontal reviews for capital (SCAP, CCAR) and executive compensation. These involve diverse groups of supervisors, economists, and lawyers, who make a point of comparing results across similar firms.

But reacting properly takes more than technical expertise. Will the regulators have the commitment to react as they should, forcing firms into resolution or requiring higher capital? How the public reacts to a crisis depends on how they expect the regulators to behave, and thus credibility and reputation become paramount. The best systemic risk measures should support this, and enable the public to hold regulators accountable. Ed Kane, among others, has called for the creation of a military-style academy for supervisors, as much to provide the *esprit de corps* needed to resist lobbying pressure as to provide advanced risk training (Kane 2011). But this also suggests that there are limits to what supervision and regulation can accomplish, even based on advanced measures of systemic risk. If so, the financial system should be designed to be robust to mistakes. But if finding the correct statistical measure of systemic risk is hard, redesigning the financial system is orders of magnitude more difficult (Haubrich 2001).

Of course, any successful attempt to measure and supervise systemic risk must be based on understanding the financial markets, on how actual institutions behave and interact. This is a tall order, as any list of the major players would include banks, brokers/dealers, hedge funds, exchanges, mutual funds, pensions plans, insurance companies, and government-sponsored enterprises. The financial crisis provided many examples of the byzantine connections between these players: consider the failure of AIG. They were writing credit protection in the form of credit default swaps (CDS) on tranches of collateralized debt obligations (CDOs) based on subprime mortgages (Stultz 2010). Risk was transferred from home lenders via the derivatives market to an insurance company. Uncovering such connections is difficult even for highly regulated entities such as banks. For example, a new, fairly priced swap arrangement has no effect on a company's balance sheet, as the two legs are priced to offset each other. Future price changes can shift the relative valuation of the legs, however, and so the swap does constitute risk to the balance sheet. The crisis has renewed discussions of more extensive data collection, but such collection is expensive and, inevitably, a trade-off must be made. In an ironic twist of fate, in 2006 the Federal Reserve stopped reporting the M3 monetary aggregate that contained a (limited) measure of repurchase aggreements (Repos, RPs), which, barely more than a year later, emerged at the center of the financial crisis of 2007 to 2009 (Gorton 2010).

The limits of accounting information have led some to look for connections via price information: Adrian and Brunnermeier (2010) propose conditional value at risk (CoVaR), Acharya et al. (2010) use the marginal expected shortfall, and Billio et al. (2011) use principal component analysis and Granger-causality networks. This brings the discussion back full circle, in that advancements on the technical side of measuring risk can uncover structural connections. Even here, the analysis does not eliminate the need for wisdom. The connections are dynamic and changing—there was no correlation between monoline insurers and mortgage-backed securities until the monolines started writing insurance on mortgage-backeds. The whole process might be compared to a card-counting blackjack player in Las Vegas trying to find patterns in a multideck sort. A few hands do not reveal much about the remaining cards, but now start swapping in decks with extra face cards, and on top of that, every once in a while let a Tarot card from the Major Arcana turn up.

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 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 these insightful discussants' remarks after each contribution.

In "Liquidity Risk, Cash Flow Constraints, and Systemic Feedbacks," Sujit Kapadia, Matthias Drehmann, John Elliott, and Gabriel Sterne introduce a theme that reappears 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 find it harder to obtain liquidity, the problem can become systemic. The process illustrates, as do several other chapters in this volume, how the fallacy of composition can hold in the financial markets: individual defenses against risk lead to greater risk overall.

The chapter emphasizes 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) under way at the Bank of England, using detailed balance sheet information from UK banks encompassing macrocredit risk, interest and noninterest 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 accommodate both aggregate distributions and scenario analysis: large losses at some banks can be exacerbated by liability-side feedbacks, leading to system-wide instability.

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 also use the resulting price volatility to gauge risk. 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 increased risk, do the same, leading to an even larger drop in price, starting a downward spiral toward even more risk. A crisis can arise quickly, because the process is highly nonlinear, with larger movements appearing suddenly. The critical threshold depends on the specifics of each market: risk management strategies, leverage, and capital plans. The chapter applies this insight to a variety of markets, explaining the implied volatility skew for options, the procyclical impact of Basel II bank capital requirements, and the optimal design for derivatives clearing and lenders of last resort. Spelling out the precise mechanism, though a challenge, takes a vital first step in the design of more robust institutions and policies.

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 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 a large set of quarterly indicators of financial and real activity for the G7 economies over the 1980Q1 to 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 G7 economies from 2008Q4 to 2009Q1.

In "Hedge Fund Tail Risk" Tobias Adrian, Markus K. Brunnermeier, and Hoai-Luu Q. 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 chapter 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 chapter thus provides a built-in solution to the problem it uncovers. Implementing this solution may not be easy, however. In fact, the chapter 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 "How to Calculate Systemic Risk Surcharges," 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 explore the implications of taxing each firm based on its contribution to systemic risk. Specifically, 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, they also examine letting the market set the price. In such a scenario, 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 firm's 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 percent to 95 percent) 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 setting, the payoff goes not to the firm, but to the regulator. This adds a measure of precommitment to the government rescue policy.

Applying this measure of systemic risk to the recent crisis provides some encouraging results. The chapter 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 type of contingent support compares with other related proposals such as forced debt-for-equity conversions.

In "The Quantification of Systemic Risk and Stability: New Methods and Measures," Romney B. Duffey approaches the problem of predicting financial systemic risk from the standpoint of a general theory of technical systems with human involvement. Discussions about the financial crisis often borrow 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 is emphasized in the keynote address by Henry Hu on "Systemic Risk and Financial Innovation: Toward a 'Unified' Approach." Hu argues that a proper understanding of systemic risk requires understanding financial innovation as a process, focusing 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 prophecy: 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 and 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 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 assigned 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.

References

Acharya, Viral, Lasse H. Pedersen, Thomas Philippon, and Matthew Richardson. 2010. "Measuring Systemic Risk." Working Paper. New York University.

Adrian, Tobias, and Markus Brunnermeier. 2010. "CoVaR." Working Paper. Princeton University.

Aumann, Robert J., and Roberto Serrano. 2008. "An Economic Index of Riskiness." *Journal of Political Economy* 116:810–36.

Basel Committee on Banking Supervision (BCBS). 2010. *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems.* Basel: Bank for International Settlements, December. (Revised June 2011.)

Billio, Monica, Mila Getmansky, Andrew W. Lo, and Loriana Pelizzon. 2011. "Econometric Measures of Systemic Risk in the Finance and Insurance Sectors." Working Paper. MIT Sloan School.

Borio, Claudio. 2003. "Towards a Macroprudential Framework for Financial Supervision and Regulation?" *CESifo Economic Studies* 49 (2): 181–215.

Engle, Robert. 2001. "GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics." *Journal of Economic Perspectives* 15 (4): 157–68.

Federal Reserve Bank of Cleveland (FRBC). 2010. *Putting Systemic Risk on the Radar Screen.* Annual Report. Cleveland: Federal Reserve Bank of Cleveland.

Foster, Dean P., and Sergiu Hart. 2009. "An Operational Measure of Riskiness." *Journal of Political Economy* 117:785–814.

Gorton, Gary B. 2010. *Slapped by the Invisible Hand: The Panic of 2007.* New York: Oxford University Press.

Haubrich, Joseph G. 2001. "Risk Management and Financial Crises." *Economic Commentary,* February. Cleveland: Federal Reserve Bank of Cleveland.

Kane, Edward J. 2011. "Statement on Distortions in Systemic-Risk Measurement and Risk-Taking in the Financial Sector." Financial Institutions and Consumer Protection Subcommittee, US Senate Committee on Banking, Housing and Urban Affairs, Hearing on Debt Financing in the Domestic Financial Sector, August 3.

Lo, Andrew W. 1999. "The Three P's of Total Risk Management." *Financial Analysts Journal* 1999:13–26.

Stulz, Rene M. 2010. "Credit Default Swaps and the Credit Crisis." *Journal of Economic Perspectives* 24 (1): 73–92.

Tarullo, Daniel K. 2010. "Financial Regulatory Reform." Speech at the US Monetary Policy Forum, New York, February 26.

Tufte, Edward R. 1990. *Envisioning Information.* Cheshire, CT: Graphics Press.