

Hedge Funds, Systemic Risk, and the Financial Crisis of 2007–2008

Written Testimony of Andrew W. Lo*

Prepared for the U.S. House of Representatives
Committee on Oversight and Government Reform
November 13, 2008 Hearing on Hedge Funds

1. Introduction

Chairman Waxman, Ranking Minority Member Davis, and other members of the House Oversight Committee, I would like to start by thanking you for giving me an opportunity to testify at this hearing on the role of hedge funds in our financial system and their regulatory and tax status. In the interest of full disclosure, I wish to inform the committee that I am a principal investigator in a project funded by the National Science Foundation, and in addition to my academic position at MIT, I am affiliated with an asset management company that manages several hedge funds and mutual funds.

Before turning to the substance of my testimony—hedge funds, systemic risk, and regulatory oversight—in this introductory section I would like to summarize the most important themes:

1. Financial crises may be an unavoidable aspect of modern capitalism, a consequence of the interactions between hardwired human behavior and the unfettered ability to innovate, compete, and evolve. But even if crises cannot be avoided, their disruptive effects can be reduced significantly by ensuring that the appropriate parties are bearing the appropriate risks, and this is best achieved through greater transparency, particularly in the so-called “shadow banking system”. Government can play a central role in providing such transparency.
2. Before we can hope to manage the risks of financial crises effectively, we must be able to define and measure those risks explicitly. Therefore, the first order of business for designing new regulations is to develop a formal definition of systemic risk and to construct specific

* Harris & Harris Group Professor, MIT Sloan School of Management, and Chief Scientific Officer, AlphaSimplex Group, LLC. Please direct all correspondence to: Andrew W. Lo, MIT Sloan School of Management, 50 Memorial Drive, E52-454, Cambridge, MA 02142. The views and opinions expressed in this article are those of the author only, and do not necessarily represent the views and opinions of AlphaSimplex Group, MIT, any of their affiliates and employees, or any of the individuals acknowledged below. The author makes no representations or warranty, either expressed or implied, as to the accuracy or completeness of the information contained in this article, nor is he recommending that this article serve as the basis for any investment decision—this article is for information purposes only. I am grateful to Jerry Chafkin, John Cox, Arnout Eikeboom, Jacob Goldfield, Gilles Guerin, Joe Haubrich, Bob Jarrow, Amir Khandani, Bob Merton, Maureen O’Hara, and Phil Vasan for helpful comments and discussion, and I thank AlphaSimplex and the MIT Laboratory for Financial Engineering for research support.

measures that are sufficiently practical and encompassing to be used by policymakers and the public. Such measures may require hedge funds and other parts of the shadow banking system to provide more transparency on a confidential basis to regulators, e.g., information regarding their assets under management, leverage, liquidity, counterparties, and holdings.

3. The most pressing regulatory change with respect to the financial system is to provide the public with information regarding those institutions that have “blown up”, i.e., failed in one sense or another. This could be accomplished by establishing an independent investigatory agency or department patterned after the National Transportation Safety Board, e.g., a “Capital Markets Safety Board”, in which a dedicated and experienced team of forensic accountants, lawyers, and financial engineers sift through the wreckage of every failed financial institution and produces a publicly available report documenting the details of each failure and providing recommendations for avoiding such fates in the future.
4. To the average American, the current financial crisis is a mystery, and concepts like subprime mortgages, CDO’s, CDS’s, and the “seizing up” of credit markets only creates more confusion and fear. A critical part of any crisis management protocol is to establish clear and regular lines of communication with the public, and a dedicated inter-agency team of public relations professionals should be formed for this express purpose, possibly within the Capital Markets Safety Board.
5. Current GAAP accounting methods are backward-looking by definition and not ideally suited for providing risk transparency, yet accounting measures are the primary inputs to corporate decisions and regulatory requirements. A new branch of accounting—“risk accounting”—must be developed and widely implemented before we can truly measure and manage systemic risk on a global scale.
6. All technology-focused industries run the risk of technological innovations temporarily exceeding our ability to use those technologies wisely. In the same way that government grants currently support the majority of Ph.D. programs in science and engineering, new funding should be allocated to major universities to greatly expand degree programs in financial technology.
7. The complexity of financial markets is straining the capacity of regulators to keep up with its innovations, many of which were not contemplated when the existing regulatory bodies were first formed. New regulations should be adaptive and focused on financial functions rather than institutions, making them more flexible and dynamic. An example of an adaptive regulation is a requirement to standardize an OTC contract and create an organized exchange for it whenever its size—as measured by open interest, trading volume, or notional exposure—exceeds a certain threshold.

I would like to add three caveats to the discussion that follows below. The first is that while the need for regulatory reform may seem clear in light of the current financial crisis, the underlying causes are complex, multi-faceted, and not yet completely understood. Therefore, I would urge the Committee and other parts of government to refrain from reacting too hastily to market events, but to deliberate thoughtfully and broadly to craft new regulations for the financial

system of the 21st century. Financial markets do not need more regulation; they need more effective regulation.

Second, although much of the material in this testimony is based on and informed by my academic research, a significant portion of the inferences surrounding systemic risk and hedge funds is indirect and circumstantial because of the hedge-fund industry's lack of transparency. Without more comprehensive data on hedge-fund characteristics such as assets under management, leverage, counterparty relationships, and portfolio holdings, it is virtually impossible to draw conclusive inferences about the systemic risks posed by hedge funds. I will attempt to point out the most fragile of my claims below, but I ask the Committee members to please bear in mind the tentative and potentially controversial nature of some of my conclusions and recommendations.

Third, since my testimony will become part of the public record, I wish to emphasize that this document is not a formal academic research paper, but is intended for a broader audience of policymakers and regulators. In particular, academic readers may be alarmed by the lack of comprehensive citations and literature review, the imprecise and qualitative nature of certain arguments, and the abundance of illustrative examples, analogies, and metaphors. Accordingly, such readers are hereby forewarned—this paper is not research, but is instead a summary of the policy implications that I have drawn from my interpretation of that research.

I begin in Section 2 with a proposal to measure systemic risk, and argue that this is the natural starting point for regulatory reform since it is impossible to manage something that cannot be measured. In Section 3, I review the relation between systemic risk and hedge funds, and show that early warning signs of the current crisis did exist in the hedge-fund industry as far back as 2004. However, I argue in Section 4 that financial crises may be an unavoidable aspect of human behavior, and the best we can do is to acknowledge this tendency and be properly prepared. This behavioral pattern, as well as traditional economic motives for regulation—public goods, externalities, and incomplete markets—are relevant for systemic risk or its converse, “systemic safety”, and in Section 5 I suggest applying these concepts to the functions of the financial system to yield a rational process for regulatory reform. In Section 6, I propose the formation of a new investigative office patterned after the National Transportation Safety Board to provide the kind of information aggregation and transparency that is called for in the previous sections, and in Section 7, I discuss fair-value accounting, which involves another critical aspect of transparency and systemic risk. The role of financial technology and education in the current crisis is considered in Section 8, where I argue that more finance training is needed, not less. I conclude in Section 9.

2. Measures of Systemic Risk

The well-known adage that “one cannot manage what one cannot measure” is particularly timely with respect to the notion of *systemic risk*, a term that has come into common usage but which has so far resisted formal definition and quantification. Like Justice Potter Stewart's definition of the obscene, systemic risk has historically been defined in a similar fashion, mainly by central bankers who know it when they see it. Systemic risk is usually taken to mean the risk of a broad-based breakdown in the financial system, often realized as a series of correlated defaults

among financial institutions, typically banks, that occurs over a short period of time and typically caused by a single major event. The classic example is a banking panic in which large groups of depositors decide to withdraw their funds simultaneously, creating a “run” on bank assets that can ultimately lead to multiple bank failures. Banking panics were not uncommon in the United States during the nineteenth and early twentieth centuries, culminating in the 1930–1933 period with an average of 2,000 bank failures per year during these years (Mishkin, 1997), and which prompted the Glass-Steagall Act of 1933 and the establishment of the Federal Deposit Insurance Corporation (FDIC) in 1934.

Although today banking panics are virtually non-existent thanks to the FDIC and related central banking policies, systemic risk exposures have taken shape in other forms. In particular, many financial institutions now provide some of the same services that banks have traditionally provided, but are outside of the banking system. For example, securitization has opened up new sources of capital to finance various types of borrowing that used to be the exclusive province of banks, including credit-card debt, trade credit, auto and student loans, mortgages, small-business loans, and revolving credit agreements. This so-called “shadow banking system”—consisting of investment banks, hedge funds, mutual funds, insurance companies, pension funds, endowments and foundations, and various broker/dealers and related intermediaries—provided a significant fraction of the liquidity needs of the global economy over the past two decades, supporting the growth and prosperity that we have enjoyed until recently. And after the repeal of the Glass-Steagall Act in 1999, the shadow banking system grew even more rapidly in size and importance. However, as its moniker suggests, the shadow banking system is neither observable nor controlled by the regulatory bodies that were created to manage the risks of potential liquidity disruptions. Therefore, it is not surprising that we were unprepared for the current financial crisis, and that we lack the proper tools to manage it effectively.

The starting point for regulatory reform is to develop a formal definition of systemic risk, one that captures the linkages and vulnerabilities of the entire financial system, not just those of the banking system. From such a definition, several quantitative measures of systemic risk should follow, with which we can monitor and manage the overall level of risk to the financial system. Even the most conservative central bank would agree that attempting to eliminate all systemic risk is neither feasible nor desirable—risk is an unavoidable by-product of financial innovation. But unless we are able to measure this type of risk objectively and quantitatively, it is impossible to determine the appropriate trade-off between such risk and its rewards.

Given the complexity of the global financial system, it is unrealistic to expect that a single measure of systemic risk will suffice. A more plausible alternative is a collection of measures, each designed to capture a specific risk exposure. For example, any comprehensive collection of risk measures should capture the following characteristics of the entire financial system:

- Leverage
- Liquidity
- Correlation
- Concentration
- Sensitivities
- Connectedness

Leverage refers to the aggregate amount of credit that has been extended in the financial system, and liquidity refers to the ease with which investments may be liquidated to raise cash. The precise mechanism by which these two characteristics combine to produce systemic risk is now well understood. Because many investors make use of leverage, their positions are often considerably larger than the amount of collateral posted to support those positions. Leverage has the effect of a magnifying glass, expanding small profit opportunities into larger ones, but also expanding small losses into larger losses. And when adverse changes in market prices reduce the market value of collateral, credit is withdrawn quickly and the subsequent forced liquidation of large positions over short periods of time can lead to widespread financial panic, as we have witnessed over the past several months. The more illiquid the portfolio, the larger the price impact of a forced liquidation, which erodes the investor's risk capital that much more quickly. Now if many investors face the same "death spiral" at the same time, i.e., if they become more highly correlated during times of distress, and if those investors are obligors of a small number of major financial institutions, then small market movements can cascade quickly into a global financial crisis. This is systemic risk. However, the likelihood of a major dislocation also depends on the degree of correlation among the holdings of financial institutions, how sensitive they are to changes in market prices and economic conditions, how concentrated the risks are among those financial institutions, and how closely connected those institutions are with each other and with the rest of the economy.

Although these six characteristics are simple to state, developing quantitative measures that can be applied to the global financial system may be more challenging. By looking at the financial system as a single portfolio, several useful measures of systemic risk can be derived by applying the standard tools of modern portfolio analysis. For example, Bodie, Gray, and Merton (2007) and Gray and Malone (2008) apply the well-known framework of contingent claims analysis to the macroeconomy, which yields several potentially valuable early warning indicators of systemic risk including aggregate asset-liability mismatches, nonlinearities in the risk/return profile of the financial sector, and default probabilities for sovereign debt. Getmansky, Lo, and Makarov (2004) propose simple measures of illiquidity risk exposures that can also be applied to the financial system. Chan et al. (2006, 2007) and Lo (2008) contain other risk analytics that are designed to measure sensitivities, correlations, and concentration in traditional and alternative investments, and these measures did provide early warning signs of potential dislocation in the hedge-fund industry in 2004 and 2005 (see, for example, Gimein, 2005).

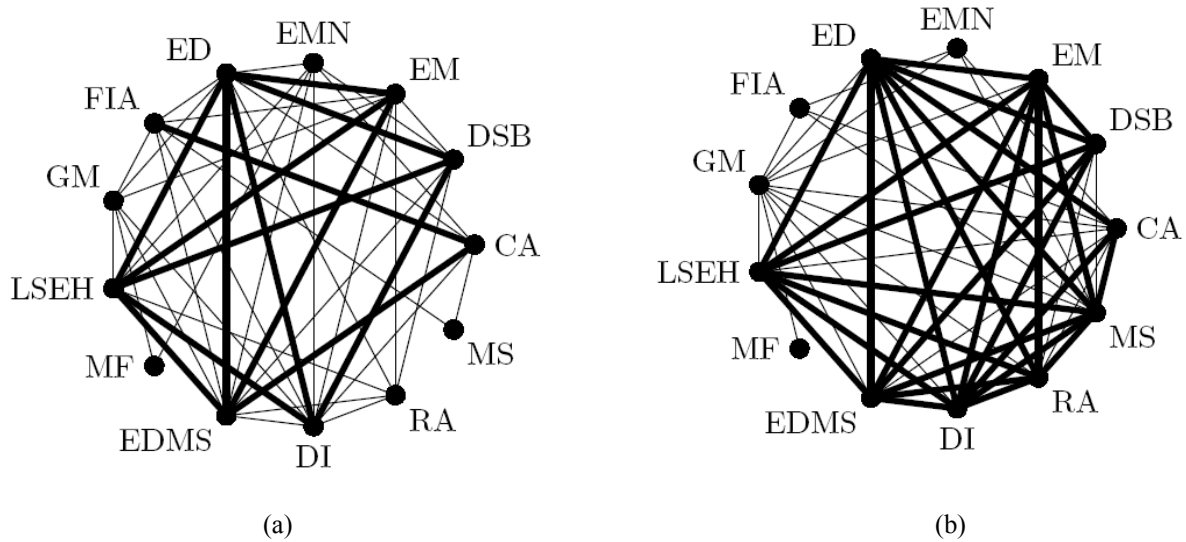


Figure 1. Network diagrams of correlations among 13 CS/Tremont hedge-fund indexes over two sub-periods: (a) April 1994 to December 2000 (excluding the month of August 1998), and (b) January 2001 to June 2007. Thicker lines represent absolute correlations greater than 50%, thinner lines represent absolute correlations between 25% and 50%, and no connecting lines correspond to correlations less than 25%. CA: Convertible Arbitrage, DSB: Dedicated Short Bias, EM: Emerging Markets, EMN: Equity Market Neutral, ED: Event Driven, FIA: Fixed Income Arbitrage, GM: Global Macro, LSEH: Long/Short Equity Hedge, MF: Managed Futures, EDMS: Event Driven Multi-Strategy, DI: Distressed Index, RA: Risk Arbitrage, and MS: Multi-Strategy. (source: Khandani and Lo, 2007)

Finally, a number of recent advances in the theory of networks (for example, Watts and Strogatz, 1998, and Watts, 1999), may be applicable to analyzing vulnerabilities in the financial network. A simple example of this new perspective is contained in Figure 1, which displays the absolute values of correlations among hedge-fund indexes over two periods, April 1994 to December 2000 and January 2001 to June 2007, where thick lines represent absolute correlations greater than 50%, thinner lines represent absolute correlations between 25% and 50%, and no lines represent absolute correlations below 25%. A comparison of the two sub-periods shows a significant increase in the absolute correlations in the more recent sample—the hedge-fund industry has clearly become more closely connected. More recently, Soramäki et al. (2007) have adopted this network perspective by mapping the topology of the Fedwire inter-bank payment system, which has generated a number of new insights about the risk exposures of this important network, including where the most significant vulnerabilities are (see Figure 2).

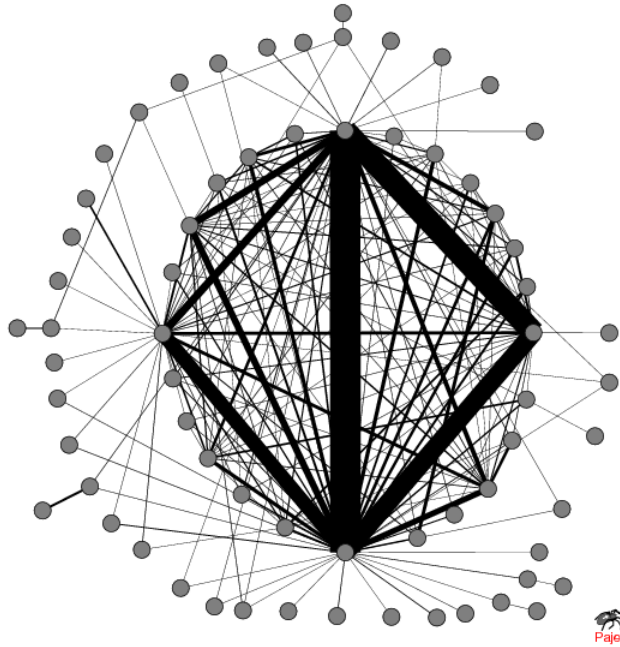


Figure 2. Core of the Fedwire Interbank Payment Network, from Soramäki et al. (2007, Figure 2).

Although a number of indirect measures of systemic risk can be computed from existing data, the biggest obstacle is the lack of sufficient transparency with which to implement these measures directly. While banks and other regulated financial institutions already provide such information, the shadow banking system does not. Without access to primary sources of data—data from hedge funds, their brokers, and other counterparties—it is simply not possible to derive truly actionable measures of systemic risk. Therefore, the need for additional data from all parts of the shadow banking system is a pre-requisite for regulatory reform in the hedge-fund industry. In particular, I propose that hedge funds with more than \$1 billion in gross notional exposures be required to provide regulatory authorities such as the Federal Reserve or the SEC with the following information on a regular, timely, and confidential basis:

- Assets under management
- Leverage
- Portfolio holdings
- List of credit counterparties
- List of investors

Given the large number of hedge funds versus the much smaller number of prime brokers (these are brokers that have hedge funds as clients), it may be more efficient for regulatory authorities to obtain these data directly from the prime brokers, or even to ask prime brokers to compute certain risk analytics specified by regulators and provide them electronically on an automated basis to regulators so as to preserve confidentiality and streamline the reporting process.

However, it is important to balance the desire for transparency against the necessity of preserving the intellectual property that hedge funds possess. Unlike other technology-based

industries, the vast majority of financial innovations are protected through trade secrecy, not patents.¹ Hedge funds are among the most secretive of financial institutions because their franchise value is almost entirely based on the performance of their investment strategies, and this type of intellectual property is perhaps the most difficult to patent. Therefore, hedge funds have an affirmative obligation to their investors to protect the confidentiality of their investment products and processes. If hedge funds are forced to reveal their strategies, the most intellectually innovative ones will simply cease to exist or move to other less intrusive regulatory jurisdictions. This would be a major loss to U.S. capital markets and the U.S. economy, hence it is imperative that regulators tread lightly with respect to this issue. One compromise is for regulators to obtain aggregated, redacted, and coded hedge-fund information—possibly pre-computed risk analytics described above—from the prime brokers that service hedge funds. This approach is operationally more efficient (there are only a few prime brokers, and they service the majority of hedge funds), and by assigning anonymous codes to every fund (so that the identities of the hedge funds are not divulged, but their information is stored in a consistent fashion across multiple prime brokers) or by transmitting pre-computed risk analytics, the proprietary aspects of the hedge funds' portfolios and strategies are protected.

3. Hedge Funds and Systemic Risk

One of the most vibrant parts of the financial sector over the last decade has been the hedge-fund industry. Relatively unconstrained by regulatory oversight, motivated by profit-sharing incentive fees, and drawn to far-flung corners of the investment universe, hedge funds have taken on a broad array of risks that would have otherwise been borne by less willing market participants. The increased risk-sharing capacity and liquidity provided by hedge funds over the last decade has contributed significantly to the growth and prosperity that the global economy has enjoyed. For example, hedge funds have raised tens of billions of dollars over the past three years for infrastructure investments, i.e., highways, bridges, power plants, and waste treatment and water purification facilities in India, Africa, and the Middle East. In their quest for greater profitability, hedge funds now provide liquidity in every major market, taking on the role of banks in fixed-income and money markets, and marketmakers and broker/dealers in equities and derivatives markets.

However, as part of the shadow banking system, hedge funds lie outside the purview of the Federal Reserve, the Office of the Comptroller of the Currency, the SEC, and CFTC, and the Treasury. Therefore, it is impossible to determine definitively what their contribution to systemic risk is. As early as 2004, Chan et al. (2004) presented indirect evidence that the level of systemic risk in the hedge-fund industry had increased; in particular, they conclude with the following summary:

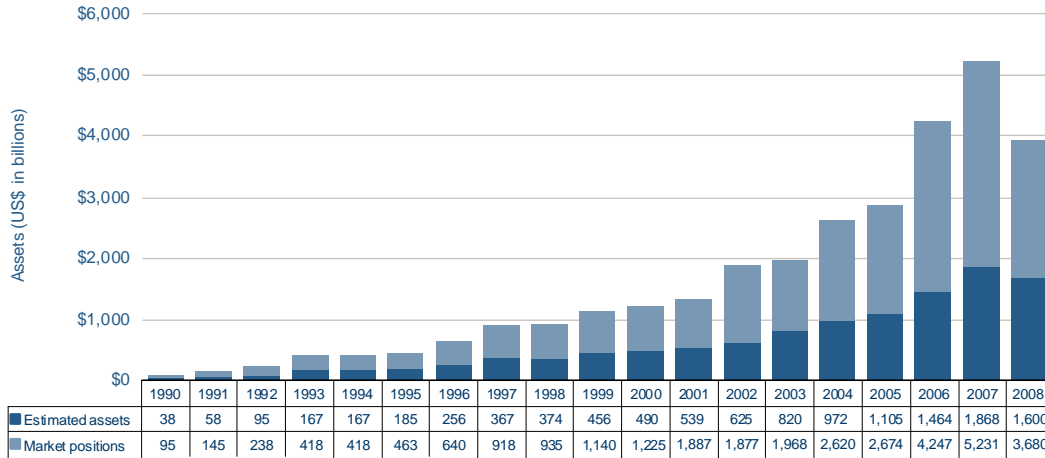
1. The hedge-fund industry has grown tremendously over the last few years, fueled by the demand for higher returns in the face of stock-market declines and mounting pension-fund liabilities. These massive fund inflows have had a material impact on hedge-fund returns and risks in recent years, as evidenced by changes in correlations, reduced performance, increased illiquidity as

¹ See Lerner (2002) for a review of financial patents.

measured by the weighted autocorrelation ρ^*_t , and the large number of hedge funds launched and closed.

2. The banking sector is exposed to hedge-fund risks, especially smaller institutions, but the largest banks are also exposed through proprietary trading activities, credit arrangements and structured products, and prime brokerage services.
3. The risks facing hedge funds are nonlinear and more complex than those facing traditional asset classes. Because of the dynamic nature of hedge-fund investment strategies, and the impact of fund flows on leverage and performance, hedge-fund risk models require more sophisticated analytics, and more sophisticated users.
4. The sum of regime-switching models' high-volatility or low-mean state probabilities can measure the aggregate level of distress in the hedge-fund sector. Recent measurements suggest that we may be entering a challenging period. This, coupled with the recent uptrend in the weighted autocorrelation ρ^*_t , implies that systemic risk is increasing.

Although based on indirect technical research findings, these conclusions were not hard to justify from casual empirical observation of general economic conditions over the past decade. The low interest-rate and low credit-spread environment of the 1990's created greater competition for yield among investors, hence large sums of money from retail and institutional investors flowed into virtually every type of higher-yielding investment opportunity available, including hedge funds, mutual funds, residential real estate, mortgages, and, of course, CDO's, CDS's, and other "exotic" securities. This push for yield also manifested itself in significant legislative pressure to relax certain constraints, resulting in the repeal of the Glass-Steagall Act and the growth of government-sponsored enterprises such as Fannie Mae and Freddie Mac. The overall impact of these conditions was to create an over-extended financial system—part of which was invisible to regulators and outside their direct control—that could not be sustained indefinitely. Moreover, the financial system became so "crowded" in terms of the extraordinary amounts of capital deployed in every corner of every investable market, that the overall liquidity of those markets declined significantly. The implication of this crowdedness is simple: the first sign of trouble in one part of the financial system will cause nervous investors to rush for the exits, but—as the analogy suggests—it is impossible for everyone to get out at once, and this panic can quickly spread to other parts of the financial system. To develop a sense for the potential scale of such a panic, consider the growth of hedge-fund assets from 1990 to the present plotted in Figure 3, and note the sharp decline in assets and leverage in 2008 (with 2008Q4 estimated by Credit Suisse). The responsiveness of hedge-fund investors to underperformance is well-known, and these relatively rapid changes in risk capital can lead to wild market gyrations as we have experienced recently (see also Khandani and Lo's, 2007 and 2008, analysis of the August 2007 "quant meltdown").



Sources: through Q3 2008 – HFR industry report. Q42008 projections – based on CS analysis

Figure 3. Growth of assets and leverage in the hedge-fund industry from 1990 to 2008 (source: HFR industry report and Credit Suisse for 2008Q4 projections)

In 2005 and 2006, Chan et al. (2006, 2007) extended these tentative conclusions with additional data and analytics, and with each iteration, they uncovered more indirect evidence for increasing levels of systemic risk. The recurring themes from their analysis were increasing assets flowing into all parts of the hedge-fund industry, correspondingly lower returns presumably as a result of these increased asset levels, greater illiquidity risk and leverage as hedge funds undertook more exotic investments using greater leverage to boost their returns, and finally, greater correlation among different hedge-fund strategies, particularly with respect to losses. These themes built to a crescendo in the first half of 2007 with the demise of several prominent multi-billion-dollar hedge funds involved in mortgage-backed securities and credit-related strategies, and apparently caused significant dislocation in August 2007 in a completely unrelated part of the hedge-fund industry—long/short equity market-neutral funds—because of desperate attempts by investors to reduce risk and raise cash to meet margin calls (see Khandani and Lo, 2007, 2008).

But why should we be concerned about the fortunes of private partnerships or wealthy investors? The reason is that over the past decade, these investors and funds have become central to the global financial system, providing loans, liquidity, insurance, risk-sharing, and other importance services that used to be the exclusive domain of banks. But unlike banks—which are highly regulated entities (but less so, since the repeal of the Glass Steagall Act in 1999), with specific capital adequacy requirements and leverage and risk constraints—hedge funds and their investors are relatively unconstrained. This freedom is important. By giving managers a broad investment mandate, hedge-fund investors are able to garner higher returns on their investments in various economic environments, including market downturns and recessions. The dynamic and highly competitive nature of hedge funds also implies that such investors will shift their assets tactically and quickly, moving into markets when profit opportunities arise, and moving out when those opportunities have been depleted. Although such tactics benefit hedge-fund investors, they can also cause market dislocation in crowded markets with participants that are not fully aware of or prepared for the crowdedness of their investments.

Beyond the proposal of Section 2 requiring hedge funds to provide additional data to regulators, it may be necessary to expand the scope of the Federal Reserve system to include direct oversight for the very largest hedge funds, their prime brokers, and other related financial institutions such as certain insurance companies engaged in bank-like activities, e.g., highly leveraged loans, credit guarantees, and retail liquidity provision. If the Fed is expected to serve as lender of last resort to non-bank financial institutions during times of distress, such institutions should be part of the Fed's permanent regulatory mandate during times of calm, which includes capital adequacy requirements, leverage restrictions, and periodic bank examinations.

4. Behavioral Foundations of Systemic Risk

Apart from the obvious and indisputable need to develop measures of systemic risk and to require financial institutions to provide additional data to regulators, proposing more specific regulatory reforms requires a deeper understanding of the underlying causes of the current crisis. There is, of course, no shortage of culprits on which the crisis can be pinned; the following is a partial list of participants who were complicit in the rise and fall of the real-estate market and the financial side-bets that went along with the bubble:

- Homeowners
- Commercial banks and savings and loan associations
- Investment banks and other issuers of MBSs, CDO's, and CDS's
- Mortgage lenders, brokers, servicers, trustees
- Credit rating agencies
- Insurance companies
- Investors (hedge funds, pension funds, sovereign wealth funds, mutual funds, endowments, and other investment institutions)
- Regulators (SEC, OCC, CFTC, Fed, etc.)
- Government sponsored enterprises
- Politicians and their constituents

Not surprisingly, with a crisis of this magnitude, all of us have played a part in its care and feeding and there is plenty of blame to go around. But while the finger-pointing may continue over the coming months and years, a more productive line of inquiry is to identify causal factors that can *only* be addressed through regulatory oversight. To that end, there are two observations that may be useful in identifying such factors.

The first observation is that the current crisis is not unique. Despite the number of seemingly unprecedented events that have transpired in 2007 and 2008, from a longer and global historical perspective, credit crises occur with some regularity. Consider, for example, the following set of concerns regarding the strength of the banking system expressed by a U.S. central banker:

...first, the attenuation of the banking systems' base of equity capital; second, greater reliance on funds of a potentially volatile character; third, heavy loan commitments in relation to resources; fourth, some deterioration in the quality of assets; and, fifth, increased exposure to the larger banks to risks entailed in foreign exchange transactions and other foreign operations.

This seems as relevant today as it was in 1974 when Fed chairman Arthur Burns spoke before the American Bankers Association about the soundness of the banking system (see Minsky, 2008, p. 57). In his conclusion, Burns observed that “our regulatory system failed to keep pace with the need,” and “a substantial reorganization [of the regulatory machinery] will be required to overcome the problems inherent in the existing structural arrangement”.

Reinhart and Rogoff (2008a) provide a more systematic analysis of the uniqueness of the subprime mortgage meltdown of 2007–2008 by identifying eighteen bank-centered financial crises that have occurred around the world since 1974,² and coming to the following conclusion after comparing them to the current crisis:

Our examination of the longer historical record, which is part of a larger effort on currency and debt crises, finds stunning qualitative and quantitative parallels across a number of standard financial crisis indicators. To name a few, the run-up in U.S. equity and housing prices that Graciela L. Kaminsky and Carmen M. Reinhart (1999) find to be the best leading indicators of crisis in countries experiencing large capital inflows closely tracks the average of the previous eighteen post World War II banking crises in industrial countries. So, too, does the inverted v-shape of real growth in the years prior to the crisis. Despite widespread concern about the effects on national debt of the early 2000s tax cuts, the run-up in U.S. public debt is actually somewhat below the average of other crisis episodes.

Figure 4 displays four graphs from Reinhart and Rogoff (2008a) that highlight the remarkable parallels in real housing prices, real equity prices, real GDP growth per capita, and public debt as a fraction of GDP between the current crisis and the eighteen others that have occurred in various countries since 1974.

The second observation is that the common theme in the majority of these crises is a period of great financial liberalization and prosperity that preceded the crisis (see Kaminsky and Reinhart, 1999, and Reinhart and Rogoff, 2008b). While this boom/bust pattern is familiar to macroeconomists, who have developed complex models for generating business cycles, there may be a simpler explanation based on human behavior. During extended periods of prosperity, market participants become complacent about the risk of loss—either through systematic underestimation of those risks because of recent history, or a decline in their risk aversion due to increasing wealth, or both. In fact, there is mounting evidence from cognitive neuroscientists that financial gain affects the same “pleasure centers” of the brain that are activated by certain narcotics.³ This suggests that prolonged periods of economic growth and prosperity can induce a collective sense of euphoria and complacency among investors that is not unlike the drug-induced stupor of a cocaine addict. Moreover, the financial liberalization that typically

² In particular, Reinhart and Rogoff (2008) identify five “Big” crises (with the year in which the crisis started in parentheses)—Spain (1977), Norway (1987), Finland (1991), Sweden (1991), and Japan (1992)—and thirteen other banking and financial crises—Australia (1989), Canada (1983), Denmark (1987), France (1994), Germany (1977), Greece (1991), Iceland (1985), Italy (1990), New Zealand (1987), United Kingdom (1974, 1991, 1995), and United States (1984).

³ In particular, the same neural circuitry that responds to cocaine, food, and sex—the mesolimbic dopamine reward system that releases dopamine in the nucleus accumbens—has been shown to be activated by monetary gain as well. See, for example, Delgado et al. (2000), Breiter et al. (2001), Montague and Berns (2002), Schultz (2002), and Knutson and Peterson (2005).

accompanies this prosperity implies greater availability of risk capital, greater competition for new sources of excess expected returns, more highly correlated risk-taking behavior because of the “crowded trade” phenomenon, and a false sense of security derived from peers who engage in the same behavior and with apparent success.

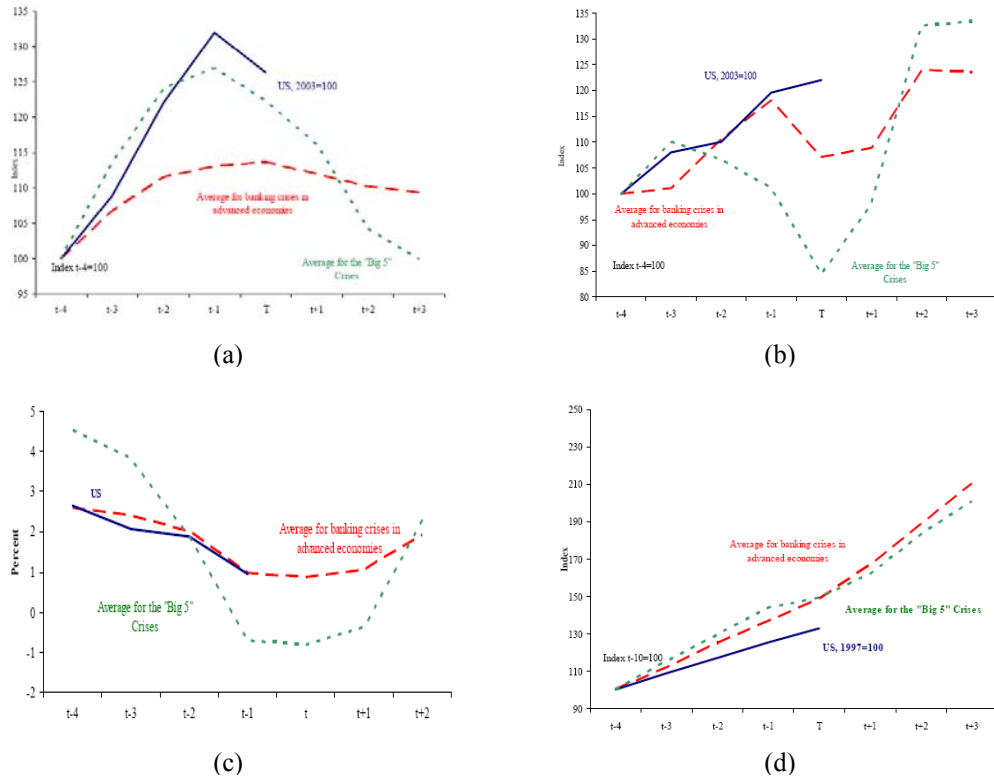


Figure 4. Reinhart and Rogoff’s (2008a) comparison of current (a) real housing prices, (b) real equity prices; (c) real GDP growth per capita; and (d) public debt as a fraction of GDP to those of other countries before, during, and after eighteen financial crises since 1974.

Consider, for example, the case of a Chief Risk Officer (CRO) of a major investment bank XYZ, a firm actively engaged in issuing and trading collateralized debt obligations (CDO’s) in 2004. Suppose this CRO was convinced that U.S. residential real estate was a bubble that was about to burst, and based on a simple scenario analysis, realized there would be devastating consequences for his firm. What possible actions could he have taken to protect his shareholders? He might ask the firm to exit the CDO business, to which his superiors would respond that the CDO business was one of the most profitable over the past decade with considerable growth potential, other competitors are getting into the business, not leaving, and the historical data suggest that real-estate values are unlikely to fall by more than 1 or 2 percent per year, so why should XYZ consider exiting and giving up its precious market share? Unable to convince senior management of the likelihood of a real-estate downturn, the CRO suggests a compromise—reduce the firm’s CDO exposure by half. Senior management’s likely response would be that such a reduction in XYZ’s CDO business will decrease the group’s profits by half, causing the

most talented members of the group to leave the firm, either to join XYZ's competitors or to start their own hedge fund. Given the cost of assembling and training these professionals, and the fact that they have generated sizable profits over the recent past, scaling down their business is also difficult to justify. Finally, suppose the CRO takes matters into his own hands and implements a hedging strategy using OTC derivatives to bet against the CDO market.⁴ From 2004 to 2006, such a hedging strategy would likely have yielded significant losses, and the reduction in XYZ's earnings due to this hedge, coupled with the strong performance of the CDO business for XYZ and its competitors, would be sufficient grounds for dismissing the CRO.

In this simple thought experiment, all parties are acting in good faith and, from their individual perspectives, acting in the best interests of the shareholders. Yet the most likely outcome is the current financial crisis. This suggests that the ultimate origin of the crisis may be human behavior—the profit motive, the intoxicating and anesthetic effects of success, and the panic sell-off that inevitably brings that success to an end.

Economists do not naturally gravitate toward behavioral explanations of economic phenomena, preferring, instead, the framework of rational deliberation by optimizing agents in a free-market context. And the ineluctable logic of neoclassical economics is difficult to challenge. However, recent research in the cognitive neurosciences has provided equally compelling experimental evidence that human decisionmaking consists of a complex blend of logical calculation and emotional response (see, for example, Damaso, 1994, Lo and Repin, 2002, and Lo, Repin, and Steenbarger, 2005). Under normal circumstances, that blend typically leads to decisions that work well in free markets. However, under extreme conditions, the balance between logic and emotion can shift, leading to extreme behavior such as the recent gyrations in stock markets around the world in September and October 2008.

This new perspective implies that preferences may not be stable through time or over circumstances, but are likely to be shaped by a number of factors, both internal and external to the individual, i.e., factors related to the individual's personality, and factors related to specific environmental conditions in which the individual is currently situated. When environmental conditions shift, we should expect behavior to change in response, both through learning and, over time, through changes in preferences via the forces of natural selection. These evolutionary underpinnings are more than simple speculation in the context of financial market participants. The extraordinary degree of competitiveness of global financial markets and the outsize rewards that accrue to the "fittest" traders suggest that Darwinian selection is at work in determining the typical profile of the successful investor. After all, unsuccessful market participants are eventually eliminated from the population after suffering a certain level of losses. For this reason, the hedge-fund industry is the Galapagos Islands of the financial system in that the forces of competition, innovation, natural selection are so clearly discernible in that industry.

This new perspective also yields a broader interpretation of free-market economics (see, for example, Lo, 2004, 2005), and presents a new rationale for regulatory oversight. Left to their own devices, market forces generally yield economically efficient outcomes under normal

⁴ In fact, most CRO's do not have unilateral authority to engage in such hedging strategies, but let us endow him with such powers just to illustrate how difficult it would be for any single individual to reign in the risk budgets of firms profiting from subprime-related activities in the most recent years leading up to the current crisis.

market conditions, and regulatory intervention is not only unnecessary but often counter-productive. However, under atypical market conditions—prolonged periods of prosperity, or episodes of great uncertainty—market forces cannot be trusted to yield the most desirable outcomes, which motivates the need for regulation. Of course, the traditional motivation for regulation—market failures due to externalities, natural monopolies, and public-goods characteristics—is no less compelling, and the desire to prevent sub-optimal behavior under these conditions provides yet another role for government intervention.

A simple example of this dynamic is the existence of fire codes enacted by federal, state, and local governments requiring all public buildings to have a minimum number of exits, well-lit exit signs, a maximum occupancy, and certain types of sprinklers, smoke detectors, and fire alarms. Why are fire codes necessary? In particular, given the costs associated with compliance, why not let markets determine the appropriate level of fire protection demanded by the public? Those seeking safer buildings should be willing to pay more to occupy them, and those willing to take the risk need not pay for what they deem to be unnecessary fire protection. A perfectly satisfactory outcome of this free-market approach should be a world with two types of buildings, one with fire protection and another without, leaving the public free to choose between the two according to their risk preferences.

But this is not the outcome that society has chosen. Instead, we require all new buildings to have extensive fire protection, and the simplest explanation for this state of affairs is the recognition—after years of experience and many lost lives—that we systematically under-estimate the likelihood of a fire.⁵ In fact, assuming that improbable events are impossible is a universal human trait (see, for example, Plous, 1993, and Slovic, 2000), hence the typical builder will not voluntarily spend significant sums to prepare for an event that most individuals will not value because they judge the likelihood of such an event to be nil. Of course, experience has shown that fires do occur, and when they do, it is too late to add fire protection. What free-market economists interpret as interference with Adam Smith’s invisible hand may, instead, be a mechanism for protecting ourselves from our own behavioral blind spots. Just as Odysseus asked his shipmates to tie him to the mast and plug his ears with wax as they sailed past the three Sirens of Circe’s islands, we use regulation as a tool to protect ourselves from our most self-destructive tendencies.

Finally, beyond the natural predilection of human behavior to excess, there is another reason to suspect that financial crises are inevitable. In studying accidents across many industries and professions—including nuclear meltdowns, chemical plant explosions, power grid failures, and airplane crashes—Perrow (1984) identifies two common elements that routinely lead to disaster: complexity and tight coupling. The former concept is clear. The latter is defined by Perrow (1984: 89–90) as “a mechanical term meaning there is no slack or buffer or give between two items [in a system]. What happens to one directly affects what happens in the other...”. Perrow concludes that accidents are normal and to be expected in complex systems that are tightly coupled. The current financial system certainly satisfies the complexity criterion (see Figures 6 and 7, and the discussion in Section 8), and the credit relationships between various

⁵ This phenomenon is a special case of the more general behavioral bias of under-estimating the likelihood of negative outcomes, and the heuristic of assigning zero probability to low-probability events (see, for example, Plous, 1993, Chapter 12, and Slovic, 2000).

counterparties—and the legal, accounting, and regulatory constraints on collateral and liquidity—have created tight coupling among many parts of the financial system. Financial crises are normal accidents.

5. A Process for Regulatory Design and Reform

With respect to regulatory reform, I acknowledge that an academic may not be in the best position to make recommendations—legislation is perhaps best left to professional legislators. However, I do think that academic research can inform the process of regulatory reform, and provide useful input in considering priorities, structure, and even implementation. This is the spirit in which I propose a somewhat different perspective to financial regulation in this section.

The behavioral and traditional rationales for regulation lead naturally to a broader approach for formulating policy and regulatory reform, an approach first advocated by Merton and Bodie (1993) for deposit insurance reform which focuses on financial functions, not financial institutions (see, also, Crane et al., 1995, and Hogan and Sharpe, 1997). The functional approach to studying financial institutions and regulation begins with the observation that there are six functions of the financial system—a payments system, a pooling mechanism for undertaking large-scale investments, resource transfer across time and space, risk management, information provision for coordinating decisions, and a means of contracting and managing agency problems. Because functions tend to be more stable than institutions, regulations designed around functional specifications are less likely to generate unintended consequences.

From a functional perspective, the standard economic approach to determining the need for regulatory oversight—identifying “market failures”—may be applied to various functions of the financial system. Among the possible sources of market failures are:

1. Public Goods (commodities like national defense that benefit everyone, but where no one has an incentive to pay for it because it is not possible to exclude anyone from its benefits once it is produced)
2. Externalities (unintended costs or benefits of an activity that are not incorporated into the market price of that activity, e.g., pollution from a factory, live music from a neighborhood bar)
3. Incomplete Markets (the absence of certain markets because there are insufficient suppliers or demanders of that product or service, e.g., unemployment insurance)
4. Behavioral Biases (certain patterns of human behavior that are recognized as undesirable and counterproductive but which are likely to occur under particular circumstances, e.g., over-eating, driving while intoxicated, panic selling of investments)

Systemic risk is a public good (or, more accurately, a public “bad”), hence government can play a positive role in addressing this market failure. To see this more clearly, consider the converse of systemic risk, i.e., systemic safety. Everyone in the global economy benefits from systemic

safety—the assurance that financial markets are stable, liquid, and reliable—but no single individual is willing to pay for this public good (in fact, it is unclear whether most individuals were even aware of the importance of this public good in their own lives until recently). The public goods aspect of liquidity in the banking system was clearly recognized by the public and private sectors at the turn of the 20th century in the United States, which led to the development of the Federal Reserve System. However, the recent growth in importance of shadow banking system has significantly reduced the ability of the Federal Reserve to maintain the same level of systemic safety as before. Several new regulations for addressing this issue are proposed in Sections 6 and 7 below.

Once a particular market failure is identified, the appropriate regulatory tools needed to address the failure will follow naturally, e.g., subsidies or taxes, proper disclosure of private information, government provision, or new securities regulation. In the case of systemic risk in the financial system, all four characteristics apply to some degree. The government is the natural provider of systemic safety because of the public goods nature of liquidity, stability and reliability, the positive externalities of a well-functioning financial system, the inability of the private sector to credibly provide systemic safety, and because individual behavior is not reliably rational during just those times when systemic safety is in jeopardy.

The functional perspective can also be applied to the organization of regulatory agencies. In the aftermath of the terrorist attacks of September 11, 2001, the need for better coordination among the FBI, CIA, NSA, and other government agencies became painfully obvious. A similar case can be made for financial regulation, which is currently distributed among several agencies and offices including the SEC, CFTC, OCC, OTS, Treasury, and the Federal Reserve. Rather than creating a new super-agency to coordinate among existing regulators, it may be more cost-effective to re-organize regulatory responsibilities according to functional lines. For example, the SEC has traditionally focused on protecting investors and ensuring fair and orderly markets, whereas the focus of the Federal Reserve has been to provide liquidity to the banking system as lender of last resort. This suggests a natural division of new regulatory responsibilities in which the management of systemic risk falls within the Federal Reserve's mandate and the creation of new exchanges continues to be part of the SEC's mandate. By focusing on functions rather than institutions, a more efficient regulatory infrastructure may be achieved.

The multi-faceted nature of systemic risk implies that several approaches to regulatory reform will be necessary. Moreover, because of the competitive and adaptive nature of financial markets, the most effective regulations are those that can adapt to changing market conditions. From an archaeological perspective, the body of securities laws may be viewed as the fossil record of the unbounded creativity of unscrupulous financiers in devising new ways to separate individuals from their money, and the multitude of strata of securities regulations trace out the co-evolution of financial misdeeds and the corresponding static regulatory responses. The most durable regulations are those that recognize the adaptive nature of markets and their participants, and are allowed to adapt accordingly.

To illustrate how adaptive regulations may be formulated, consider the credit default swaps (CDS) market. The magnitude and importance of this OTC market have led to the very sensible proposal to establish a CDS exchange with standardized contracts, daily mark-to-market and

settlement, and a clearing corporation. An adaptive version of this proposal would be to require the establishment of a similar organized exchange and clearing corporation for any set of OTC contracts that exceeds certain thresholds in volume, open interest, and notional exposure, where such thresholds should be defined in terms of percentages of those quantities in existing markets, e.g., 5% of the combined dollar volume of all organized futures markets. Such an adaptive regulation would promote the orderly and organic creation of new exchanges as the need arises, and reduce the likelihood of systemic shocks emanating from the failure of a small number of too-big-to-fail institutions. Of course, exemptive relief can always be provided under certain conditions, but the benefit of an adaptive regulation is that an orderly transition from small heterogeneous OTC trading to a market that has become vital to global financial system is permanently institutionalized.

6. The Capital Markets Safety Board

With any form of technological innovation, there is always the risk that the technology outpaces our ability to use it properly, bringing unintended consequences. The current threat of global warming is perhaps the most dramatic example of this common pattern of human progress. But in the face of space shuttle explosions, nuclear meltdowns, bridge collapses, and airplane crashes, we rarely blame the technology itself, but, instead, seek to understand how our possibly inappropriate use of the technology may have caused the accident. The outcome of that evaluation process may yield improvements in both the technology and how it is used, and this is how progress is made. Technological innovation of any form entails risk, but as long as we learn from our mistakes, we reduce the risk of future disasters.

In this respect, the financial industry can take a lesson from other technology-based professions. In the medical, chemical engineering, and semiconductor industries, for example, failures are routinely documented, catalogued, analyzed, internalized, and used to develop new and improved processes and controls. Each failure is viewed as a valuable lesson, to be studied and reviewed until all the wisdom has been gleaned from it, which is understandable given the typical cost of each lesson.

One successful model for conducting such reviews is the National Transportation Safety Board (NTSB), an independent government agency whose primary mission is to investigate accidents, provide careful and conclusive forensic analysis, and make recommendations for avoiding such accidents in the future. In the event of an airplane crash, the NTSB assembles a team of engineers and flight-safety experts who are immediately dispatched to the crash site to conduct a thorough investigation, including interviewing witnesses, poring over historical flight logs and maintenance records, and sifting through the wreckage to recover the flight recorder or “black box” and, if necessary, reassembling the aircraft from its parts so as to determine the ultimate cause of the crash. Once its work is completed, the NTSB publishes a report summarizing the team's investigation, concluding with specific recommendations for avoiding future occurrences of this type of accident. The report is entered into a searchable database that is available to the general public (see <http://www.nts.gov/ntsb/query.asp>) and this has been one of the major factors underlying the remarkable safety record of commercial air travel.

For example, it is now current practice to spray airplanes with de-icing fluid just prior to take-off when the temperature is near freezing and it is raining or snowing. This procedure was instituted in the aftermath of USAir Flight 405's crash on March 22, 1992. Flight 405 stalled just after becoming airborne because of ice on its wings, despite the fact that de-icing fluid was applied before it left its gate. Apparently, Flight 405's take-off was delayed because of air traffic, and ice re-accumulated on its wings while it waited for a departure slot on the runway in the freezing rain. The NTSB Aircraft Accident Report AAR-93/02—published February 17, 1993 and available through several internet sites—contains a sobering summary of the NTSB's findings (Report AAR-93/02, page vi):

The National Transportation Safety Board determines that the probable cause of this accident were the failure of the airline industry and the Federal Aviation Administration to provide flightcrews with procedures, requirements, and criteria compatible with departure delays in conditions conducive to airframe icing and the decision by the flightcrew to take off without positive assurance that the airplane's wings were free of ice accumulation after 35 minutes of exposure to precipitation following de-icing. The ice contamination on the wings resulted in an aerodynamic stall and loss of control after liftoff. Contributing to the cause of the accident were the inappropriate procedures used by, and inadequate coordination between, the flightcrew that led to a takeoff rotation at a lower than prescribed air speed.

Rather than placing blame on the technology, or on human error, the NTSB conducted a thorough forensic examination and concluded that an incorrect application of the technology—waiting too long after de-icing, and not checking for ice build-up just before take-off—caused the crash. Current de-icing procedures have no doubt saved many lives thanks to NTSB Report AAR-93/02, but this particular innovation did not come cheaply; it was paid for by the lives of the 27 individuals who did not survive the crash of Flight 405. Imagine the waste if the NTSB did not investigate this tragedy and produce concrete recommendations to prevent this from happening again.

Financial crashes are, of course, considerably less dire, generally involving no loss of life. However, the current financial crisis, and the eventual cost of the Paulson Plan, should be sufficient motivation to create a “Capital Markets Safety Board” (CMSB) dedicated to investigating, reporting, and archiving the “accidents” of the financial industry. By maintaining teams of experienced professionals—forensic accountants, financial engineers from industry and academia, and securities and tax attorneys—that work together on a regular basis over the course of many cases to investigate every single financial disaster, a number of new insights, common threads, and key issues would emerge from their analysis. The publicly available reports from the CMSB would yield invaluable insights to investors seeking to protect their future investments from similar fates, and in the hands of investors, this information would eventually drive financial institutions to improving their “safety records”.

In addition to collecting, analyzing, and archiving data from financial blow-ups, the CMSB should also be tasked with the responsibility of obtaining and maintaining information from the shadow banking system—hedge funds, private partnerships, sovereign wealth funds, etc.—and integrating this information with that of other regulatory agencies (see Section 2 for further discussion). By having one single agency responsible for managing data related to systemic risk, and creating high-level risk analytics such as a network map of the financial system, estimates of

illiquidity exposure, leverage, and asset flows, the repository of data will be far easier to access and analyze.

The NTSB provides an additional valuable service that the CMSB should also take on: at the very start of its investigative process, the NTSB establishes itself as the clearinghouse for all information related to the accident, and communicates frequently and regularly with the press to provide as much transparency as possible to an undoubtedly anxious public. Specifically, the following is an excerpt from the NTSB's standard operating procedure for major accident investigations (see <http://www.nts.gov>):

At a major accident, the NTSB will send several public affairs officers (PAOs) to accompany the Go-Team and facilitate information dissemination. Often, one of the five Presidentially-appointed Board Members will accompany the team and serve as principal spokesperson. The Go-Team is led by a senior career investigator designated as Investigator-in-Charge (IIC).

While the Board's investigative team includes representatives from other agencies and organizations, only the Safety Board may release factual information on the investigation. Representatives of other organizations participating in our investigation risk removal and exclusion from the process if they release investigative information without NTSB permission.

The NTSB will establish a command post near the crash site, usually in a hotel. On-site public affairs operations will be organized from the Command Post. Local phone numbers for public affairs will be announced when they have been established.

Although not possible in every circumstance, the Safety Board strives to conduct two press conferences a day when on scene, one at mid- to late-afternoon and the other in the evening following the progress meeting held by the investigative team. The Board's spokespersons discuss factual, documented information. They do not analyze that information, nor speculate as to the significance of any particular piece of information.

If conditions permit, Safety Board PAOs will attempt to gain admittance for the news media, either in total or in a pool arrangement, to the accident scene itself, keeping in mind limitations posed by physical and biomedical hazards.

The Board will maintain a public affairs presence on scene for as long as circumstances warrant, usually 3 to 7 days. After that, information will be released from the Public Affairs Office in Washington, D.C., (202) 314-6100.

By taking such an active role in providing information to the public immediately and continuously throughout its investigations, the NTSB reduces the likelihood of panic and overreaction, and over the years, this policy has earned the NTSB the public's trust and confidence. Compare this approach to the sporadic and inconsistent messages that were communicated to the public regarding the current financial crisis, which may well have magnified the dislocation that ensued in the stock market and money market funds during September and October 2008. Of course, the Treasury and Federal Reserve can hardly be faulted for not providing polished presentations of every aspect of their deliberations—public relations has never been a significant component of their mandate. But in times of crisis, when emotions run high, it is particularly important to communicate directly, truthfully, and continuously with all stakeholders no matter how bad the news, as any experienced emergency-room doctor will acknowledge.

Of course, formal government investigations of major financial events do occur from time to time, as in the April 1999 *Report of the President's Working Group in Financial Markets on Hedge Funds, Leverage, and the Lessons of Long-Term Capital Management*. However, this inter-agency report was put together on an *ad hoc* basis with committee members that had not worked together previously and regularly on forensic investigations of this kind. With multiple agencies involved, and none in charge of the investigation, the administrative overhead becomes more significant. Although any thorough investigation of the financial services sector is likely to involve the SEC, the OCC, the CFTC, the US Treasury, and the Federal Reserve—and inter-agency cooperation should be promoted—there are important operational advantages in tasking a single office with the responsibility for coordinating all such investigations and serving as a repository for the expertise in conducting forensic examinations of financial incidents.

The establishment of a CMSB will not be inexpensive. The lure of the private sector poses a formidable challenge to government agencies to attract and retain individuals with expertise in these highly employable fields. Individuals trained in forensic accounting, financial engineering, and securities law now command substantial premiums on Wall Street over government pay scales. Although the typical government employee is likely to be motivated more by civic duty than financial gain, it would be unrealistic to build an organization on altruism alone. However, the cost of a CMSB is trivial in comparison to the losses that it may prevent. For example, if regulators had fully appreciated the impact of the demise of Lehman Brothers—which a fully operational CMSB with the proper network map would have been able to forecast—the savings from this one incident would be sufficient to fund the CMSB for half a century. Moreover, the benefits provided by the CMSB would accrue not only to the wealthy, but would also flow to pension funds, mutual funds, and retail investors in the form of more stable financial markets, greater liquidity, reduced borrowing and lending costs as a result of decreased systemic risk exposures, and a wider variety of investment choices available to a larger segment of the population because of increased transparency, oversight, and ultimately, financial security.

It is unrealistic to expect that market crashes, manias, panics, collapses, and fraud will ever be completely eliminated from our capital markets, but we should avoid compounding our mistakes by failing to learn from them.

7. Transparency and Fair-Value Accounting

A common theme among the issues and proposals throughout my testimony has been the importance of transparency for managing systemic risk. Financial markets are highly competitive and adaptive, and—apart from occasional dislocations due to overwhelming behavioral reactions—are extremely effective in aggregating, parsing, internalizing, and disseminating information, the quintessential illustration of the “wisdom of crowds”. As a general principle, the more transparency is provided to the market, the more efficient are the prices it produces, and the more effective will the market allocate capital and other limited resources. When the market is denied critical information, its participants will infer what they can from existing information, in which case rumors, fears, and wishful thinking will play a much bigger role in how the market determines prices and quantities. Therefore, from a

systemic risk perspective, as well as a social welfare perspective, it is difficult to justify any regulatory change that interferes with or otherwise reduces transparency.

One example of such a change is the recent proposal to suspend “Fair-Value Accounting” (FASB Statement No. 157). Fair-value or mark-to-market accounting requires firms to value their assets and liabilities at fair market prices, not on a historical-cost basis, and this practice has been blamed for the current financial crisis because it has forced many firms to write down their assets, thereby triggering defaults and insolvencies. At first blush, this proposal seems ill-conceived because it calls for less transparency. After all, in the current credit crisis, banks are refusing to lend to each other because they have no idea what the other banks’ assets are worth, and suspending fair-value accounting will not improve this state of affairs. Imagine a doctor advising the parents of a feverish child to discontinue hourly temperature readings because the frequent readings only serve to alarm them. Instead, the doctor suggests that the parents either wait until the child is feeling better before taking the next reading, or that they construct an estimate of the child’s temperature based on readings taken last week when the child was feeling better.

Nevertheless, the proposal is worth more serious consideration because it involves several subtle issues surrounding the economic nature of markets, prices, and the importance of transparency. There is no doubt that a suspension of fair-value accounting will reduce current pressures on a number of potential insolvent financial institutions. However, this reduction in current pressures comes at a cost, which depends on whether the suspension of fair-value accounting is temporary or permanent. If it is permanent, market participants lose a significant degree of transparency regarding corporate assets and liabilities, and will price securities accordingly. Borrowing costs will likely increase across the board, and because firms with higher-quality assets may not have any mechanism to convince the market of this fact, such firms may refrain from participating in capital markets, thereby reducing market liquidity and also creating adverse selection (where only firms with lower-quality assets remain in the market), which raises borrowing costs even more. Moreover, on an ongoing basis, firms will have to maintain larger reserves to achieve the same credit quality because of the increased risk of their less-transparent portfolios, further reducing liquidity and increasing borrowing costs.

If the suspension of fair-value accounting is temporary, then there must be a day of reckoning when firms will have to mark their assets and liabilities to market, and the suspension is merely a postponement of that eventuality. A postponement is reasonable under two conditions: (1) the existence of extraordinary circumstances that cause market prices of the firm’s assets and liabilities to deviate significantly from economic value; and (2) the extraordinary circumstances are temporary and unrelated to the economic value of the firm’s assets and liabilities. For example, suppose a terrorist attack on U.S. soil creates a massive but temporary flight-to-quality, during which time the value of an insurance company’s assets, which are largely invested in AAA-rated corporate debt, falls precipitously. In this scenario, the flight-to-quality is temporary, and the decline in the insurance company’s assets is largely (although not completely) unrelated to its economic value, hence a temporary suspension of fair-value accounting may be defensible since the insurance company is likely to be solvent once the flight-to-quality passes.

However, this argument implicitly assumes that the flight-to-quality is a form of temporary insanity that should be dismissed or, at the very least, discounted. But if, for example, the flight-to-quality is not a temporary phenomenon, but rather a change in regime that is likely to last for years, e.g., because the terrorist event portends ongoing threats that cannot easily be eliminated, then the suspension of fair-value accounting is delaying the inevitable and interfering with the appropriate re-pricing of business enterprises under a new economic regime.

Also, the temporariness of impairments to economic value is insufficient justification for suspending fair-value accounting—the condition that the impairment be unrelated to the economic value of the impaired asset is also critical. The reason is simple: even if an asset's market value is only temporarily impaired, if the impairment is directly related to the nature of that asset then it should be taken into account and marked to market. For example, suppose a bank holds part of its assets in a hurricane insurance company. During an unusually active hurricane season, the value of these holdings may be temporarily impaired, but this impairment is directly related to the economic value of the insurance company, and suspension of fair-value accounting only interferes with the price discovery process.

A more subtle argument for suspending fair-value accounting has been put forward by Plantin, Sapra, and Shin (2008), who observe that during periods where liquidity is very low, a forced liquidation of an asset by firm A can depress the market price of that asset, which affects the value of firm B if it also holds that same asset and is required to mark that asset to market. In such situations, fair-value accounting inadvertently creates correlation among the assets of many firms, even those that are not attempting liquidations, and this increased correlation can lead to the kind of “death spiral” discussed above, where liquidations cause deterioration of collateral that leads to more liquidations, further deterioration of collateral, and so on. They consider this spillover effect a negative externality—a negative consequence of the liquidation that is not part of the economic value of the asset being liquidated—and argue that in some cases (long-lived and illiquid assets), it is socially optimal to use historical-cost accounting instead of fair-value accounting.

However, their conclusion rests heavily on the interpretation of the spillover effect as negative externality. Another interpretation is that such spillover effects are, in fact, part of the economic value of an asset. In particular, if the asset in question were short-term U.S. Treasury Bills, then presumably the spillover effects of a liquidation would be minimal. But then the price of T-Bills should reflect this property, and the price of less liquid assets should reflect the potential spillover effects as well. Therefore, the potential for spillover effects is a characteristic that can be known in advance and priced accordingly, in which case the welfare implications of switching from fair-value to historical-cost accounting is unclear. In any case, Plantin, Sapra, and Shin (2008) do not consider the case where fair-value accounting is temporarily suspended, and the impact of moving from one regime to another in their framework is an open question. The answer will likely depend on whether the two conditions described above hold for the fair-value postponement.

Other recent studies have argued that during liquidity crises, market prices are not as meaningful as they are during normal times, hence marking securities to market may not always yield the same information content (see, for example, Allen and Carletti, 2008, Easley and O'Hara, 2008,

and Sapra, 2008). While no consensus has yet emerged regarding which alternative to mark-to-market pricing is best, the fact that the distinction between “liquidity pricing” and “normal pricing” is being drawn more frequently by accountants and economists suggests that neither historical-cost accounting nor fair-value accounting can be appropriate for all circumstances. This implies that neither approach is the correct one, but that a more flexible mechanism for pricing assets and liabilities is needed—one that can accommodate both normal and distressed market conditions.

A more general observation about accounting methods is that they are designed to yield information about value, not about risk, a point made by Merton and Bodie (1995) as part of their call for a separate branch of accounting focusing exclusively on risk. They use the example of a simple fixed/floating interest-rate swap contract which has zero value at the start, hence is considered neither an asset nor a liability, but is an “off-balance-sheet” item. We have learned from experience that off-balance-sheet items can have enormous impact on a firm’s bottom line, hence it is remarkable that our accounting practices have yet to incorporate them more directly in valuation. In fairness to the accounting profession, accounting methods are designed to be *backward-looking*, involving the allocation of revenues and costs that have already occurred to various categories. Accountants tell us what *has* happened, leaving the future to corporate strategists and fortunetellers. But this exclusive focus on realized results implies that risk is not part of the accountant’s lexicon, i.e., there is no natural way to capture risk from the current GAAP accounting perspective. Yet accounting concepts like capital ratios and asset/liability gaps are used to formulate and implement regulatory requirements and constraints.

A modest beginning for developing risk accounting methods is to define the concept of a “risk balance sheet”, which is simply the risk decomposition of a firm’s mark-to-market balance sheet where both assets and liabilities are considered to be random variables, i.e., unknown quantities with certain statistical properties. Since assets must always equal liabilities, the variance of assets must always equal the variance of liabilities, hence the risk balance sheet is just the variance decomposition of both sides (see Figure 5). Note that the variance of both total assets and total liabilities is given by the sum of the variances of the individual assets and liabilities, plus their pairwise covariances. These are the terms that have created so much controversy with respect to subprime mortgage-backed securities—they swelled to unprecedented levels in 2007 as subprime mortgage defaults became highly correlated throughout the country. Risk accounting standards—which have yet to be developed—must address both the proper methods for estimating the variances and covariances of assets and liabilities, and the potential instabilities in these estimates across different economic environments.

This challenge is not just a regulatory one, but requires regulators to collaborate with accountants and financial experts to develop a completely new set of accounting principles focused exclusively on risk budgeting. This new branch of risk accounting may be one of the most critical pieces of the financial and regulatory infrastructures of the 21st century.

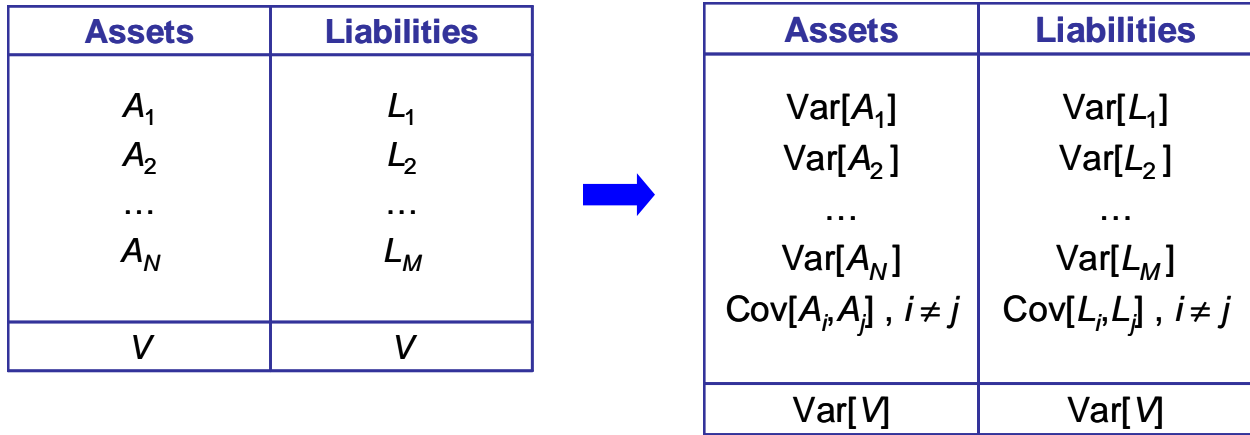


Figure 5. The “risk balance sheet”, defined as the risk decomposition of a firm’s market-value balance sheet.

8. The Role of Technology and Education

Given the complexity of the financial structures involved in the current crisis—mortgage-backed securities, collateralized debt obligations, credit default swaps, SPV’s, SIV’s, and other exotic entities and securities—a natural question is whether the crisis was due to financial technology, e.g., derivative securities, structured products, and the mathematical methods involved in pricing and hedging them? If, as Warren Buffett claims, derivatives are financial weapons of mass destruction, should we consider outlawing derivatives altogether?

There is no doubt that financial technology has had an indelible impact on the financial system over the past two decades. However, that technology has been used as often to reduce risk as it has to augment it. For example, as of July 2008, the notional amount outstanding of interest-rate derivatives—one of the most popular instruments among non-financial corporations for *hedging* interest-rate risk—was \$465 trillion according to the International Swaps and Derivatives Association (see <http://www.isda.org>). In contrast to that market, the comparable notional amount for credit default swaps over the same period was \$55 trillion, and only \$12 trillion for equity derivatives. These figures suggest that the most common use of derivatives today is not as financial weapons of mass destruction, but as hedging vehicles for transferring risk from part of the global economy to other parts that are better equipped to bear that risk. Therefore, limiting the use of such important risk management tools would be counterproductive and highly disruptive.

However, Mr. Buffett may be half-right in that financial technology has become more complex over the last two decades, and because the derivatives and structured finance businesses have grown so rapidly, the expertise required to fully grasp the risk and return profiles of many new financial instruments and vehicles has increased just as rapidly. In some cases, even large financial institutions may not have had sufficient time to develop such expertise among their senior management and board members, and hiring the necessary expertise in booming business lines is always difficult by definition. To fully appreciate the intellectual challenges of recent

financial innovation, consider the case of the “hybrid” CDO first issued by HVB Asset Management in 2003, which is described by Bluhm (2003) in the following passage:

HVB Asset Management Asia (HVBAM) has brought to market the first ever hybrid collateralized debt obligation (CDO) managed by an Asian collateral manager. The deal, on which HVB Asia (formerly known as HypoVereinsbank Asia) acted as lead manager and underwriter, is backed by 120 million of asset-backed securitization bonds and 880 million of credit default swaps... Under the structure of the transaction, Artemus Strategic Asian Credit Fund Limited—a special purpose vehicle registered in the Cayman Islands—issued 200 million of bonds to purchase the 120 million of cash bonds and deposit 80 million into the guaranteed investment contract, provided by AIG Financial Products. In addition, the issuer enters into credit default swap agreements with three counterparties (BNP Paribas, Deutsche Bank and JPMorgan) with a notional value of 880 million. On each interest payment date, the issuer, after payments of certain senior fees and expenses and the super senior swap premium, will use the remaining interest collections from the GIC accounts, the cash ABS bonds, the hedge agreements, and the CDS premiums from the CDS to pay investors in the CDO transaction... The transaction was split into five tranches, including an unrated 20 million junior piece to be retained by HVBAM. The 127 million of A-class notes have triple-A ratings from Fitch, Moody’s and S&P, the 20 million B-notes were rated AA/Aa2/AA, the 20 million C bonds were rated A/A2/A, while the 13 million of D notes have ratings of BBB/Baa2 and BBB.

This new financial security is a claim on the ARTEMUS Strategic Asian Credit Fund, a Cayman Islands special purpose vehicle structured according to Figure 6. Note that this diagram is just an outline of the legal structure of the instrument! How many boards of directors of institutions managing these types of funds—of which there are many—truly understood the complexities of these investments?

Pricing such instruments is even more complex, involving a blend of mathematical, statistical, and financial models and computations, all of which are typically done under simplistic assumptions that rarely hold in practice, such as constant means, variances, and correlations that are measured without error. To develop an appreciation for the mathematical complexity of some of these pricing models, Figure 7 contains a technical appendix from Bluhm and Overbeck (2007) in which they describe one aspect of their proposed model for default probabilities, which is a critical element in evaluating the price of credit-sensitive instruments like credit default swaps. Models such as these are central to the current financial crisis, and their mis-calibration is one possible explanation for how so many firms under-estimated the risks of subprime-related securities so significantly. Unless senior management has the technical expertise to evaluate and challenge the calibrations of these models, they cannot manage their risks effectively.

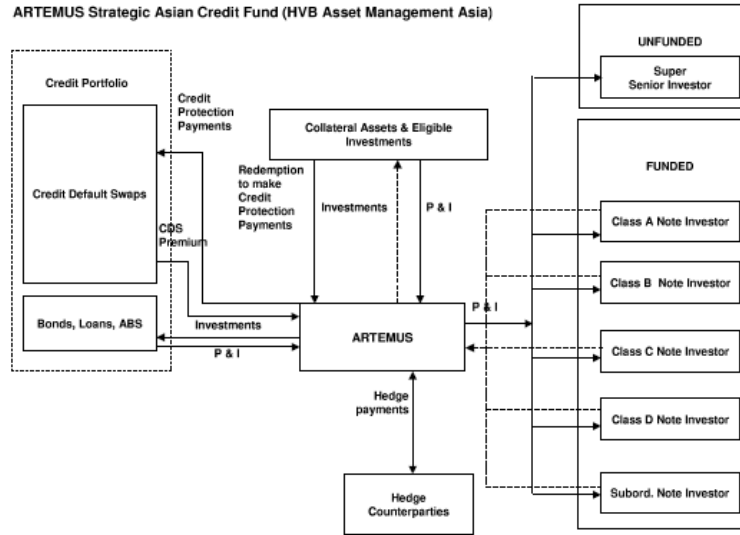


Figure 6. Structure of the ARTEMUS Strategic Asian Credit Fund, a Cayman Islands special purpose vehicle. Source: Bluhm (2003).

Now, we often take it for granted that large financial institutions capable of hiring dozens of “quants” each year must have the technical expertise to advise senior management, and senior management has the necessary business and markets expertise to guide the quantitative research process. However, in fast-growing businesses the realities of day-to-day market pressures make this idealized relationship between senior management and research a fantasy. Senior management typically has little time to review the research, much less guide it, and in recent years, many quants have been hired from technically sophisticated disciplines such as mathematics, physics, and computer science but without any formal training in finance or economics. While some on-the-job training is inevitable, the broad-based failure of the financial industry to fully appreciate the magnitude of the risk exposures in the CDO and CDS markets suggests that the problem was not too much knowledge of financial technology, but rather too little knowledge.

A case in point is the credit-rating agencies, who have been roundly criticized for their apparently overly optimistic ratings of the mortgage-backed securities and related instruments that lie at the epicenter of the current financial crisis. Some have argued that the inherent conflict of interest in the ratings business led to upward-biased ratings, others claim that the mathematical models on which ratings were based were too simplistic and static, and yet another set of critics blame the limited history that the rating agencies used to calibrate their models. Although it may be too early to draw any final conclusions about the ultimate origins of the breakdown in these credit ratings, one fact has emerged which seems uncontroversial: the clients of the rating agencies—hedge funds, commercial banks, investment banks, and mortgage companies—routinely hired away the raters’ most talented analysts. And given the business model of rating agencies, this was not hard to do, nor did the rating agencies object because it was both a compliment to the quality of their staff, and also a means for developing closer ties to their clients. But it did result in a continuous “brain drain” from the rating agencies to their clients, and even then, the demand for such talent continued to grow until the financial crisis hit.

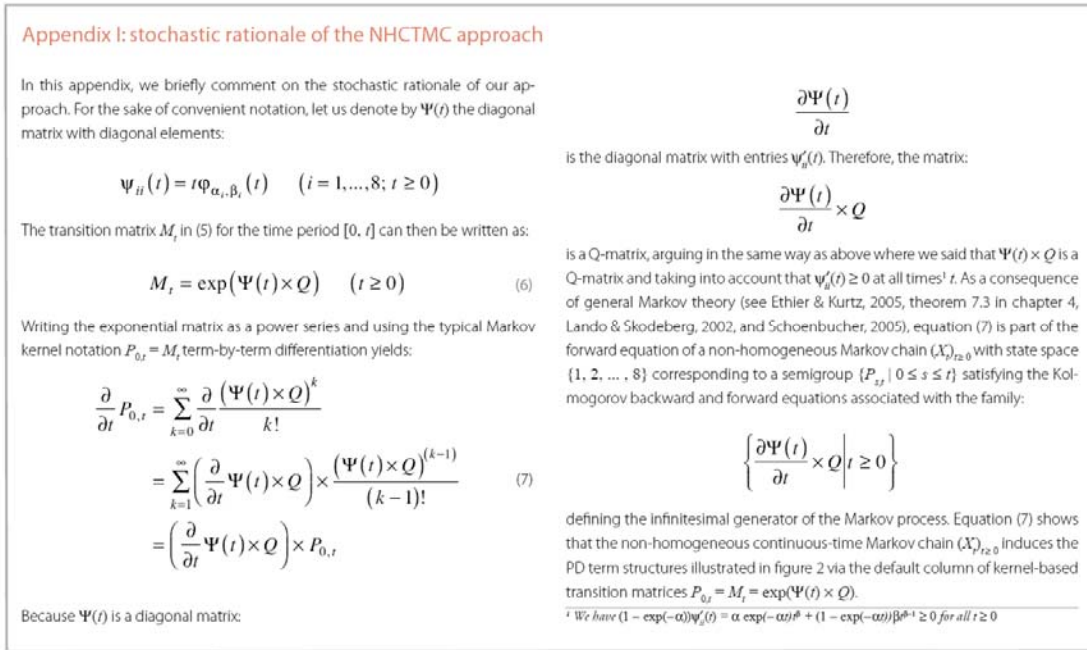


Figure 7. Appendix I from Bluhm and Overbeck (2007) providing details of their Markov default probability model. Source: Bluhm and Overbeck (2007).

Indirect evidence for an excess demand of finance expertise may also be found in Philippon and Reshef's (2007) comparison of the annual incomes of U.S. engineers and finance-trained graduates from 1967 to 2005. The comparison between finance and engineering students is a useful one because both are technical disciplines, and over the past 20 years, engineers have been making significant inroads into the finance labor market. Figure 8 shows that until the mid-1980's, college graduates in engineering enjoyed higher incomes than college graduates in finance, and post-graduates in engineering had about the same compensation as post-graduates in finance. However, since the 1980's, finance-trained college graduates have caught up to their engineering counterparts, and surpassed them in 2000 and every year thereafter. But the more impressive comparison is for post-graduates—since 1982, the annual income of finance post-graduates has exceeded that of engineers every year, and the gap has widened steadily over these two decades. This pattern suggests that the demand for financial expertise has grown considerably during this time.⁶

⁶ The increase in income can also be explained by a decline in the supply of finance graduates, but Philippon and Reshef (2007) show that the number of employees in this sector increased significantly since the 1980's, which suggests that a supply shock is not the source of the growth in income.

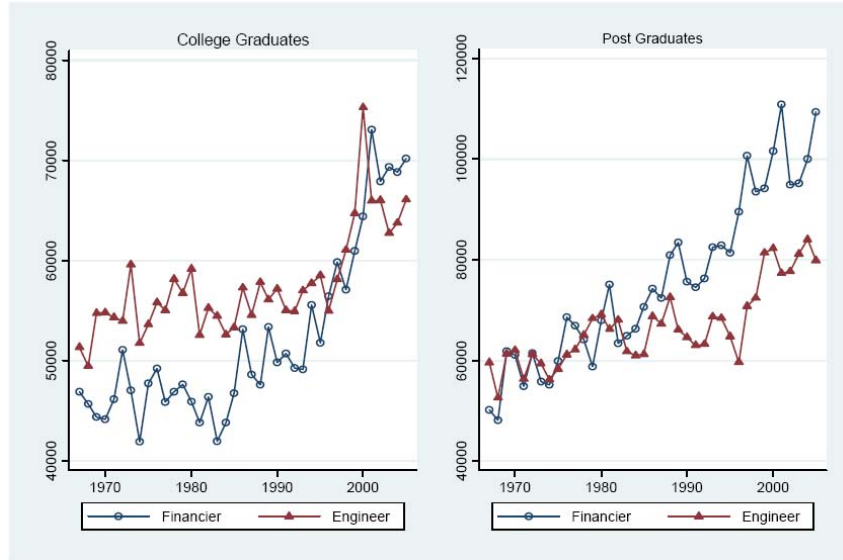


Figure 8. Comparison of finance and engineering graduates from 1967 to 2005, in 2001 U.S. dollars. Source: Philippon and Reshef (2007).

Table 1, which reports the number of MIT engineering and finance degrees awarded from 1999 to 2007, provides another perspective on the dearth of financial expertise. In 2007, MIT’s School of Engineering graduated 337 Ph.D.’s in engineering; in contrast, the MIT Sloan School of Management produced only 4 finance Ph.D.s. These figures are not unique to MIT, but are, in fact, typical among the top engineering and business schools. Now, it can be argued that the main focus of the Sloan School is its M.B.A. program, which graduates approximately 300 students each year, but most M.B.A. students at Sloan and other top business schools do not have the technical background to implement models such as the one described in Figure 7, nor does the standard M.B.A. curriculum include courses that cover such models in any depth. Such material—which requires advanced training in arcane subjects such as stochastic processes, stochastic calculus, and partial differential equations—is usually geared towards Ph.D. students. However, due to the growth of the derivatives business over the past decade, a number of universities have begun to offer specialized Master’s-level degree programs in financial engineering and mathematical finance to meet the growing demand for more technically advanced students trained in finance. Whether or not such students are sufficiently prepared to fill the current knowledge gap in financial technology remains to be seen.

The disparity between the number of Ph.D.s awarded in engineering and finance in Table 1 raises the question of why such a difference exists. One possible explanation may be the sources of funding. MIT engineering Ph.D. students are funded largely through government grants (DARPA, DOE, NIH, and NSF), whereas MIT Sloan Ph.D. students are funded exclusively through internal MIT funds. Given the importance of finance expertise, one proposal for regulatory reform is to provide comparable levels of government funding to support finance Ph.D. students, perhaps in conjunction with the research activities of the Capital Markets Safety Board (see Section 6). Alternatively, funding for finance Ph.D. students might be raised by

imposing a small surcharge on certain types of derivatives contracts, e.g., those that are particularly complex or illiquid and, therefore, contribute to systemic risk. This surcharge may be viewed as a means of correcting some the externalities associated with the impact of derivatives on systemic risk. A minuscule surcharge on, say, credit default swaps, could support enough finance Ph.D. students at every major university to have a noticeable and permanent impact on the level of financial expertise in both industry and government.

Year	MIT School of Engineering			MIT Sloan
	Bachelor's	Master's and MEng	PhD and ScD	Finance PhD
2007	578	710	337	4
2006	578	735	298	2
2005	593	798	286	1
2004	645	876	217	5
2003	679	817	210	7
2002	667	803	239	3
2001	660	860	248	1
2000	715	739	237	2
1999	684	811	208	4

Table 1. Number of MIT degrees awarded in engineering and in finance from 1999 to 2007. Source: MIT Annual Report of the President, 1999 to 2007.

In addition to providing support for finance Ph.D. students, another potential new role for government oversight is to mandate minimum levels of disclosure, “truth-in-labeling” laws, and financial expertise for those market participants involved in creating and selling complex financial securities to the public, much like the requirements imposed by the Food and Drug Administration on accurate and complete labeling of pharmaceuticals, and the educational and licensing requirements for pharmacists dispensing those products. Currently, a licensed pharmacist must earn a Pharm.D. degree from an accredited college or school of pharmacy, and then pass a series of examinations including the North American Pharmacist Licensure Exam (which tests for pharmacy skills and knowledge) and, in most states, the Multistate Pharmacy Jurisprudence Exam (which tests for knowledge of pharmacy law). The SEC already performs this function to some degree, but its focus is limited to a much narrower and simpler set of financial products than those at the center of the current crisis.⁷ Rather than setting up the infrastructure for administering such educational and licensing requirements, the government can partner with existing industry organizations such as the CFA Institute, the International Association of Financial Engineers, or the Global Association of Risk Professionals.

⁷ However, Macey, O’Hara, and Rosenberg (2008) make the following two bold claims: “First, we argue that the current subprime mortgage and credit crisis would have been avoided, or at least greatly mitigated, if existing securities laws had been properly applied to subprime mortgage brokers and originators. Second, we argue that under either of what we regard as two extremely reasonable interpretations of the securities laws, many of the problematic mortgages are actually under the SEC’s jurisdiction.”

9. Conclusion

While the current financial crisis is the most significant challenge in our lifetime, it is not unprecedented from a global historical perspective, nor is the magnitude unexpected given the excesses, growth, and financial liberalization of the past decade. While there are many factors that have contributed to the crisis, ultimately, we may conclude that the boom/bust pattern of economies is a natural consequence of human evolution and adaptation to a complex and dynamic economy. Recent research in the cognitive neurosciences confirms that fear and greed are hardwired into our decisionmaking processes, and the cyclical nature of economic growth is merely one manifestation of that hardwiring. Financial crises are an unfortunate but normal consequence of modern capitalism.

Although financial crises may be difficult to avoid, their devastating impact can be dramatically reduced with proper preparation. Financial losses are inevitable—in fact, they are a necessary consequence of innovation—but disruptions and dislocations are greatly magnified when risks have been incorrectly assessed and incorrectly assigned. For example, a money market fund investing in AAA-rated securities is not prepared for situations in which those securities exhibit one-year default rates of 5%, but a hedge fund investing in B-rated securities is prepared for considerably higher default rates. The most effective means for reducing the impact of any financial crisis is providing the public with greater transparency into the underlying risks of their investments.

The next several years will no doubt be extremely challenging for the U.S. economy. However, the likely contraction, rise in unemployment, and regulatory reforms can be viewed as the necessary restructuring costs for transitioning the existing economy to an even more robust one, a globally integrated economy in which labor and capital are more mobile, production is more efficient, and information is central to profitability and survival. And by implementing adaptive and functional regulatory changes, we will be creating the new infrastructure to support that growth and prosperity.

References

- Allen, F. and E. Carletti, 2008, “Mark-to-Market Accounting and Liquidity Pricing”, *Journal of Accounting and Economics* 45, 358–378.
- Bodie, Z., Gray, D., and R. Merton, 2007, “New Framework for Measuring and Managing Macrofinancial Risk and Financial Stability”, NBER Working Paper No. W13607.
- Bluhm, C., 2003, “CDO Modeling: Techniques, Examples and Applications”, unpublished working paper, http://www.defaultrisk.com/pp_crdrv_42.htm
- Bluhm, C. and L. Overbeck, 2007, “Calibration of PD Term Structures: To Be Markov or Not to Be”, *Risk* November, 98–103.
- Breiter, H., Aharon, I., Kahneman, D., Dale, A. and P. Shizga, 2001, “Functional Imaging of Neural Responses to Expectancy and Experience of Monetary Gains and Losses”, *Neuron* 30, 619–639.
- Chan, N., Getmansky, M., Haas, S., and A. Lo, 2004, “Systemic Risk and Hedge Funds”, Conference Paper, *NBER Conference on the Risks of Financial Institutions*, Woodstock, VT, October 22–23.
- Chan, N., Getmansky, M., Haas, S., and A. Lo, 2006, “Do Hedge Funds Increase Systemic Risk?”, *Federal Reserve Bank of Atlanta Economic Review* Q4, 49–80.
- Chan, N., Getmansky, M., Haas, S., and A. Lo, 2007, “Systemic Risk and Hedge Funds”, in M. Carey and R. Stulz, eds., *The Risks of Financial Institutions*. Chicago, IL: University of Chicago Press.
- Crane, D., Froot, K., Mason, S., Perold, A., Merton, R., Bodie, Z., Sirri, E. and P. Tufano, 1995, *The Global Financial System: A Functional Perspective*. Boston, MA: Harvard Business School Press.
- Damasio, A., 1994, *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Avon Books.
- Delgado, M., Nystrom, L., Fissell, C., Noll, D. and J. Fiez, 2000, “Tracking the Hemodynamic Response to Reward and Punishment in the Striatum”, *Journal of Neurophysiology* 84, 3072–3077.
- Easley, D. and M. O’Hara, 2008, “Liquidity and Valuation in an Uncertain World”, unpublished working paper, June.
- Getmansky, M., Lo, A. and I. Makarov, 2004, “An Econometric Analysis of Serial Correlation and Illiquidity in Hedge-Fund Returns”, *Journal of Financial Economics* 74, 529–609.

Getmansky, M., Lo, A., and S. Mei, 2004, “Sifting Through the Wreckage: Lessons from Recent Hedge-Fund Liquidations”, *Journal of Investment Management* 2, 6–38.

Gimein, M., 2005, “Is a Hedge Fund Shakeout Coming Soon?”, *New York Times*, Sunday, September 4.

Gray, D. and S. Malone, 2008, *Macrofinancial Risk Analysis*. New York: John Wiley & Sons.

Hogan, W. and I. Sharpe, 1997, “Prudential Regulation of the Financial System: A Functional Approach”, *Agenda* 4, 15–28.

Kaminsky, G. and C. Reinhart, 1999, “The Twin Crises: The Causes of Banking and Balance of Payments Problems”, *American Economic Review* 89, 473–500.

Khandani, A. and A. Lo, 2007, “What Happened to the Quants in August 2007?”, *Journal of Investment Management* 5, 5–54.

Khandani, A. and A. Lo, 2008, “What Happened to the Quants in August 2007?: Evidence from Factors and Transactions Data”, NBER Working Paper No. 14465.

Knutson, B. and R. Peterson, 2005, “Neurally Reconstructing Expected Utility”, *Games and Economic Behavior* 52, 305–315.

Lerner, J., 2002, “Where Does *State Street* Lead? A First Look at Finance Patents, 1971–2000”, *Journal of Finance* 57, 901–930.

Lo, A., 2004, “The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective”, *Journal of Portfolio Management* 30, 15–29.

Lo, A., 2005, “Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis”, *Journal of Investment Consulting* 7, 21–44.

Lo, A., 2008, *Hedge Funds: An Analytic Perspective*. Princeton, NJ: Princeton University Press.

Lo, A. and D. Repin, 2002, “The Psychophysiology of Real-Time Financial Risk Processing”, *Journal of Cognitive Neuroscience* 14, 323–339.

Lo, A., Repin, D. and B. Steenbarger, 2005, “Fear and Greed in Financial Markets: An Online Clinical Study”, *American Economic Review* 95, 352–359.

Macey, J., O’Hara, M. and G. Rosenberg, 2008, “Helping Law Catch Up to Markets: Applying Securities Law to Subprime Mortgages”, Unpublished Working Paper, October.

Merton, R. and Z. Bodie, 1993, “Deposit Insurance Reform: A Functional Approach”, *Carnegie-Rochester Conference Series on Public Policy* 38, 1–34.

Merton, R. and Z. Bodie, 1995, “Financial Infrastructure and Public Policy: A Functional Perspective”, in Crane, D., Froot, K., Mason, S., Perold, A., Merton, R., Bodie, Z., Sirri, E. and P. Tufano, 1995, *The Global Financial System: A Functional Perspective*, Chapter 8. Boston, MA: Harvard Business School Press.

Minsky, H., 2008, *Stabilizing an Unstable Economy*. New York: McGraw-Hill.

Mishkin, F., 1997, *The Economics of Money, Banking, and Financial Markets*, 5th edition. Reading, MA: Addison-Wesley.

Montague, R. and G. Berns, 2002, “Neural Economics and the Biological Substrates of Valuation”, *Neuron* 36, 265–284.

Perrow, C., 1984, *Normal Accidents: Living with High Risk Technologies*. New York: Basic Books.

Philippon, T. and A. Reshef, 2007, “Skill Biased Financial Development: Education, Wages and Occupations in the U.S. Financial Sector”, Working Paper, Stern School of Business, New York University.

Plantin, G., Sapra, H. and H. Shin, 2008, “Marking-to-Market: Panacea or Pandora's Box?”, *Journal of Accounting Research* 46, 435–460.

Plous, S., 1993, *The Psychology of Judgment and Decision Making*. New York: McGraw-Hill.

Reinhart, C. and K. Rogoff, 2008a, “This Time Is Different: A Panoramic View of Eight Centuries of Financial Crises”, NBER Working Paper No. 13882.

Reinhart, C. and K. Rogoff, 2008b, “Is the 2007 U.S. Subprime Crisis So Different? An International Historical Comparison”, *American Economic Review* 98, 339–344.

Sapra, H., 2008, “Do Accounting Measurement Regimes Matter? A Discussion of Mark-To-Market Accounting and Liquidity Pricing”, *Journal of Accounting and Economics* 45, 379–387.

Schultz, W., 2002, “Getting Formal with Dopamine and Reward”, *Neuron* 36, 241–263.

Slovic, P., 2000, *The Perception of Risk*. London: Earthscan Publications Ltd.

Soramäki, K., Bech, M., Arnold, J., Glass, R. and W. Beyeler, 2007, “The Topology of Interbank Payment Flows”, *Physica A* 379, 317–333.

Watts, D., 1999, *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton, NJ: Princeton University Press.

Watts, D. and S. Strogatz, 1998, “Collective Dynamics of ‘Small-World’ Networks”, *Nature* 393, 440–442.