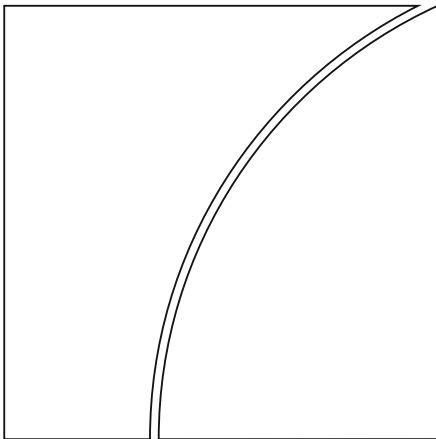




BANK FOR INTERNATIONAL SETTLEMENTS



BIS Working Papers No 564

Moore's Law vs. Murphy's Law in the financial system: who's winning?

by Andrew W Lo

Monetary and Economic Department

May 2016

JEL classification: G20, G28, G01

Keywords: Financial technology, systemic risk,
macroprudential policy, risk management

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2016. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Foreword

The 14th BIS Annual Conference took place in Lucerne, Switzerland, on 26 June 2015. The event brought together a distinguished group of central bank Governors, leading academics and former public officials to exchange views on the topic “Towards ‘a new normal’ in financial markets?”. The papers presented at the conference and the discussants’ comments are released as *BIS Working Papers* nos 561 to 564.

BIS Papers no 84 contains the opening address by Jaime Caruana (General Manager, BIS), the keynote address by John Kay (London School of Economics) and remarks by Paul Tucker (Harvard Kennedy School).

Moore's Law vs. Murphy's Law in the financial system: who's winning?¹

Andrew W Lo²

Abstract

Breakthroughs in computing hardware, software, telecommunications and data analytics have transformed the financial industry, enabling a host of new products and services such as automated trading algorithms, crypto-currencies, mobile banking, crowdfunding and robo-advisors . However, the unintended consequences of technology-leveraged finance include firesales, flash crashes, botched initial public offerings, cybersecurity breaches, catastrophic algorithmic trading errors and a technological arms race that has created new winners, losers and systemic risk in the financial ecosystem. These challenges are an unavoidable aspect of the growing importance of finance in an increasingly digital society. Rather than fighting this trend or forswearing technology, the ultimate solution is to develop more robust technology capable of adapting to the foibles in human behaviour so users can employ these tools safely, effectively and effortlessly. Examples of such technology are provided.

Keywords: Financial technology, systemic risk, macroprudential policy, risk management

JEL classification: G20, G28, G01

¹ Prepared for the 14th BIS Annual Conference, 26 June 2015. I thank John Gutttag, Andrei Kirilenko, Eric Rosenfeld, Hyun Song Shin and conference participants for helpful comments and discussion, and Jayna Cummings for editorial assistance . Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged. The views and opinions expressed in this article are those of the author only, and do not necessarily represent the views and opinions of any institution or agency, any of their affiliates or employees, or any of the individuals acknowledged above.


² Charles E and Susan T Harris Professor, MIT Sloan School of Management; director, MIT Laboratory for Financial Engineering; principal investigator, MIT Computer Science and Artificial Intelligence Laboratory; 100 Main Street, E62-618, Cambridge, MA 02142. E-mail: alo-admin@mit.edu. Other affiliations, including all commercial relationships, are disclosed at <http://alo.mit.edu/>.

1. Introduction

In 1965 – three years before he co-founded Intel, now the largest semiconductor chip manufacturer in the world – Gordon Moore published an article in *Electronics Magazine* in which he observed that the number of transistors that could be placed onto a chip seemed to double every year. This simple observation, implying a constant rate of growth, led Moore to extrapolate an increase in computing potential from 60 transistors per chip in 1965 to 60,000 in 1975. This number seemed absurd at the time, but it was realised on schedule a decade later. Later revised by Moore to a doubling every two years, “Moore’s Law” has forecasted the growth of the semiconductor industry over the last 40 years with remarkable accuracy, as [Figure 1](#) confirms.

Illustration of Moore’s Law via transistor counts on various semiconductor chips from 1971 to 2011, which seems to double every two years

Figure 1



Source: “Transistor Count and Moore’s Law – 2011 by Wgsimon”, CC BY-SA 3.0.

Technological change is often accompanied by unintended consequences. The Industrial Revolution of the 19th century greatly increased the standard of living, but it also increased air and water pollution. The introduction of chemical pesticides greatly increased the food supply, but it also increased the number of birth defects before we understood the risks. And the emergence of an interconnected global financial system greatly lowered the cost and increased the availability of capital to businesses and consumers around the world, but those same interconnections also served as vectors of financial contagion that facilitated the Financial Crisis of 2007–09. As a result, the financial industry must weigh Moore’s Law against Murphy’s Law, “whatever can go wrong, will go wrong”, as well as Kirilenko and Lo’s (2013)

technology-specific corollary, “whatever can go wrong, will go wrong faster and bigger when computers are involved.”

Some of the unintended consequences of financial technology include firesales, flash crashes, botched initial public offerings, cybersecurity breaches, catastrophic algorithmic trading errors and a technological arms race that has created new winners, losers and systemic risk in the financial ecosystem. The inherent paradox of modern financial markets is that technology is both the problem and, ultimately, the solution. Markets cannot forswear financial technology – the competitive advantages of algorithmic trading and electronic markets are simply too great for any firm to forgo – but rather must demand better, more robust technology, technology so advanced it becomes foolproof and invisible to the human operator. Every successful technology has gone through such a process of maturation: the rotary telephone versus the iPhone, paper road maps versus the voice-controlled touchscreen GPS and the kindly reference librarian versus Google and Wikipedia. Financial technology is no different. To resolve the paradox of Moore’s Law versus Murphy’s Law, we need version 2.0 of the financial system.

2. Moore’s Law and finance


Moore’s Law now influences a broad spectrum of modern life. It affects everything from household appliances to biomedicine to national defense, and its impact is no less evident in the financial industry. As computing has become faster, cheaper, and better at automating a variety of tasks, financial institutions have been able to greatly increase the scale and sophistication of their services. The emergence of automated algorithmic trading, online trading, mobile banking, crypto-currencies like Bitcoin, crowdfunding and robo-advisors are all consequences of Moore’s Law.

At the same time, the combination of population growth and the complexity of modern society has greatly increased the demand for financial services. In 1900, the total human population was an estimated 1.5 billion, but little more than a century later – a blink of an eye in the evolutionary timescale – the world’s population has grown to 7 billion (see [Figure 2](#)). The vast majority of these 7 billion individuals are born into this world without savings, income, housing, food, education or employment. All of these necessities today require financial transactions of one sort or another, well beyond the capacity of the financial industry in 1900. Therefore, it should come as no surprise that innovations in computer hardware, software, telecommunications and storage continue to shape Wall Street as a necessary part of its growth.

In fact, technological innovation has always been intimately interconnected with financial innovation. New stamping and printing processes, used to prevent clipping, counterfeiting and other forms of financial fraud, directly led to the modern system of paper banknotes and token coinage. The invention of the telegraph sparked a continent-spanning communications revolution that led the creation of the modern futures market in nineteenth-century Chicago. And improvements to the ticker tape machine – symbolic of Wall Street for over a century – made Thomas Edison his early fortune.

Semi-logarithmic plot of estimated world population from 10,000 B.C. to 2011 A.D.

Figure 2



Source: US Census Bureau (International Data Base) and author's calculation.

Technology and derivatives

Not very long ago, most trades were made through traders and specialists on the floors of the exchanges. The first electronic exchange, NASDAQ, opened in 1971, but it was originally only a quotation system for the slow-moving over-the-counter market. Most trades were placed over the telephone and executed on the trading floor well into the 1980s. Today, however, nearly all trades on the major financial exchanges are consummated electronically. Moore's Law made the floor specialist obsolete and trading volume increased exponentially to meet this increase in trading capacity. If the modern financial system had to rely on human specialists to manage even a fraction of this market volume, it would need a trading floor larger than a sports arena.

The symbiosis between technology and finance has accelerated the pace of the financial markets beyond mere human capacity at all levels of the financial system. One elegant example comes from the options market. The Chicago Board Options Exchange (CBOE), the first of its kind, opened just before Fischer Black, Myron Scholes and Robert Merton published their foundational papers in 1973.³ However, the rapid growth of the CBOE would have been impossible had financial professionals lacked an easy way to use the Black-Scholes/Merton option pricing formula. As luck would have it, in 1975 Texas Instruments introduced the SR-52, the first programmable handheld calculator, and one capable of handling the logarithmic and exponential functions of the Black-Scholes/Merton formula (see Figure 3).

³ Black and Scholes (1973); Merton (1973).

The Texas Instruments SR-52 programmable calculator, introduced in 1975 and used to compute the Black-Scholes/Merton option-pricing formula by CBOE floor traders.

Figure 3



At \$395 (\$1,737 in 2015 dollars), the SR-52 was a technological marvel that could store programs of up to 224 keystrokes on a thin magnetic strip that was fed through a motorised slot in the calculator. Shortly after the SR-52 debuted, one of the founders of the CBOE, a savvy options trader named Irwin Gutttag, purchased one for his teenage son and asked him to program the Black-Scholes/Merton formula for it.⁴ Within a year, many CBOE floor traders were sporting SR-52s of their own. By 1977, Texas Instruments introduced a new programmable TI-59 with a “Securities Analysis Module” that would automatically calculate prices using the Black-Scholes/Merton formula. When Scholes confronted Texas Instruments about their unauthorised use of the formula, they replied that it was in the public domain. When he asked for a calculator instead, Texas Instruments replied that he should buy one.⁵

A financial Moore’s Law

The Black, Scholes and Merton publications launched well over a thousand subsequent articles (eg Lim and Lo (2006)), becoming the intellectual foundation for three sectors of the derivatives industry: exchange-traded options, over-the-counter structured products and credit derivatives. As of September 2015, there were \$36 trillion of exchanged-traded options outstanding, and as of the first half of 2015, there were \$553 trillion in notional value of foreign exchange, interest rate, credit and other over-the-counter derivatives.⁶ As Lo (2013) observed, “In the modern history of all the social sciences, no other idea has had more impact on both theory and practice in such a short time span.” The reason cited is serendipity: the convergence of science, with the Black-Scholes/Merton formula; technology, with the formation of the CBOE; and need, for risk mitigation created by the economic turmoil of the mid-1970s.

⁴ Private communication with John V Gutttag – Irwin Gutttag’s son and SR-52 programmer – who became a computer scientist, eventually serving as chair of MIT’s Department of Electrical Engineering and Computer Science, and currently Dugald C Jackson Professor of Electrical Engineering and Computer Science at MIT.

⁵ Scholes (2006).

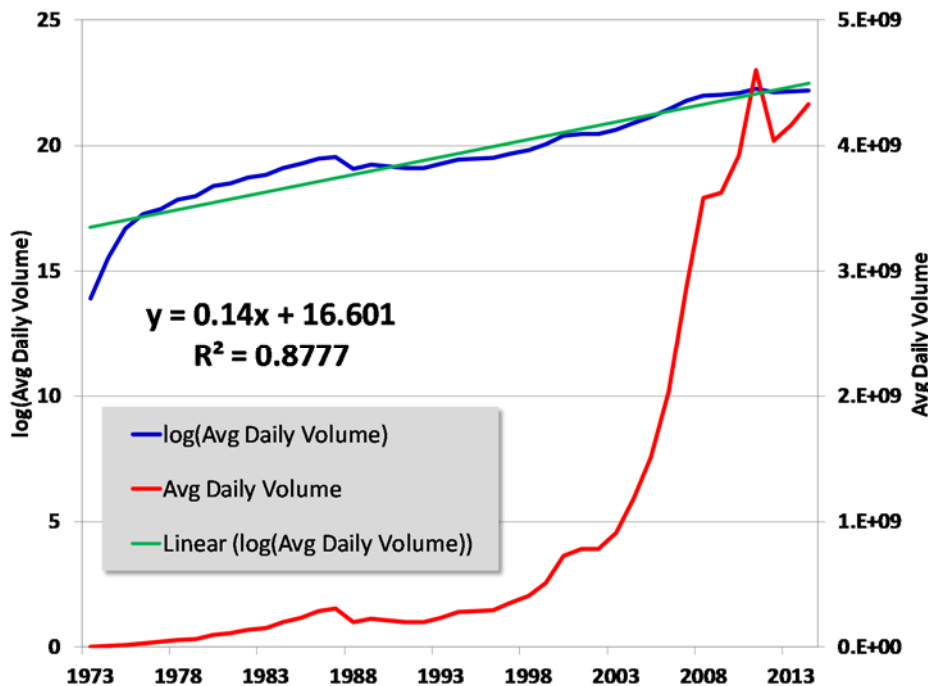
⁶ Bank for International Settlements (2015, Table D1 at <http://stats.bis.org/statx/srs/table/d1> and Table D5.1 at <http://stats.bis.org/statx/srs/table/d5.1> accessed 16 Jan 2016).

Figure 4 highlights just one consequence of such serendipity: the annual raw and log average daily trading volume of exchange listed options and futures from the Options Clearing Corporation from 1973 to 2014. A simple log-linear regression yields the financial equivalent of Moore’s Law: the volume of exchange-traded derivatives doubles approximately every five years. Even the Financial Crisis of 2007–09 could only temporarily halt this growth for a couple of years, after which the trend seems to continue unabated.

It should be clear from these examples that Moore’s Law and the exponential growth of computing power have utterly transformed the financial system.⁷ The collective intelligence of the market, dependent as it is on the rapid collection of accurate information, has been greatly magnified by the advances in telecommunications, processing power and data storage that Moore’s Law has made possible. The consequent easy access to financial services throughout the developed world has transformed the modern consumer lifestyle in a thousand small and large ways, from everyday purchases at the local coffee shop through a frictionless global electronic payment network, to investing for the life-changing events in one’s future with a combinatorially vast variety of individualised financial products.

Financial Moore’s Law: Raw and natural logarithm of average daily trading volume by year of exchange-listed options and futures from the Options Clearing Corporation, 1973–2014, and linear regression estimate of geometric growth rate, which implies a doubling every $\log_2(1/0.14) = 4.95$ years.

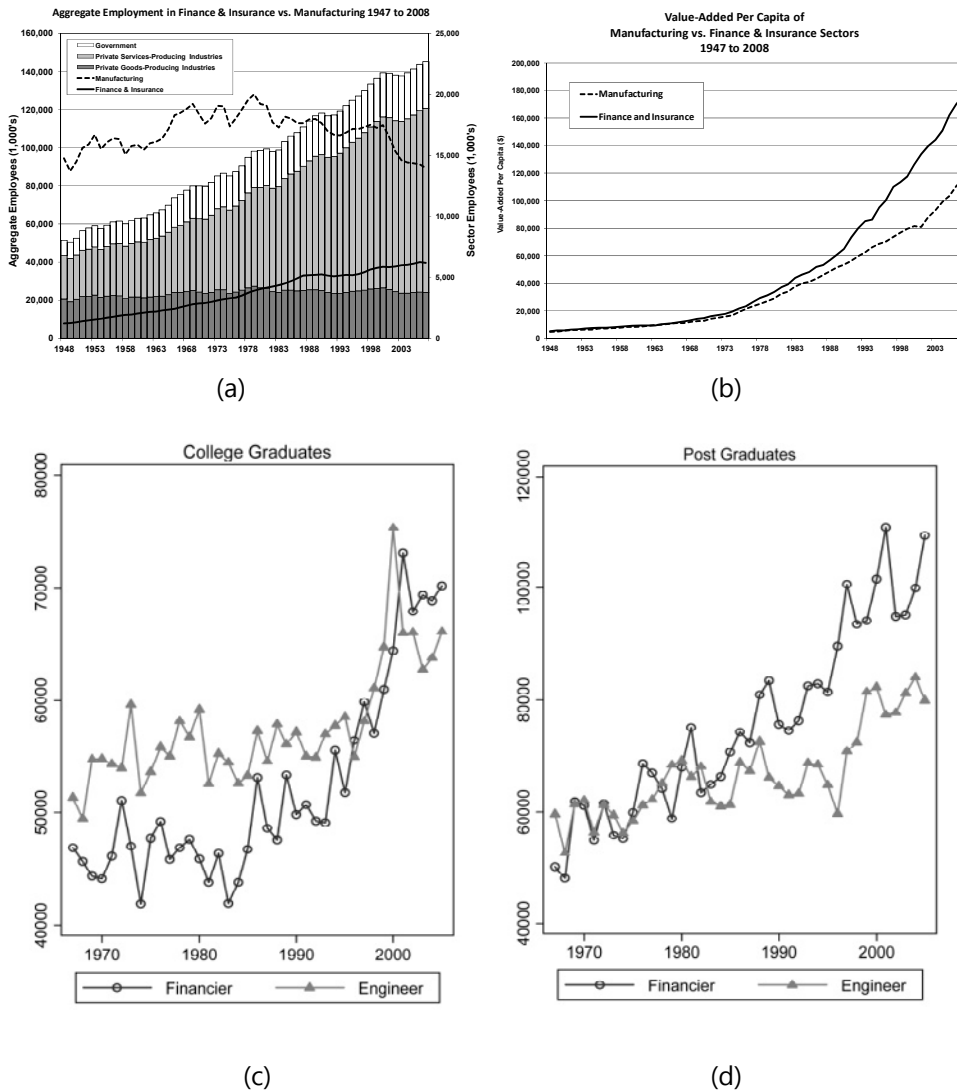
Figure 4



⁷ Kirilenko and Lo (2013).

Four illustrations of the growing importance of finance: (a) aggregate US employment in manufacturing vs finance and insurance sectors; (b) value-added per capita in manufacturing vs finance and insurance; (c,d) annual income of college-graduate and post-graduate engineers and financiers (all wages are in 2000 US dollars and are weighted using sampling weights), from Phillipon (2009, Figure 7).

Figure 5



The inexorable march of technological progress is part of a much broader trend of finance increasing its role in modern society. Figure 5 presents four illustrations of this trend. Figure 5(a) shows that aggregate employment in the finance and insurance sectors has been increasing steadily over time, unlike the manufacturing sector, which employs about the same number of workers today as it did in the 1940s. The manufacturing sector is able to produce a much greater GDP with the same amount of labour because of technological progress, especially automation. This explanation is confirmed in Figure 5(b), which shows an upward-sloping graph of the value-added per capita in the manufacturing sector, a clear measure of increasing productivity in manufacturing. However, Figure 5(b) also shows that the finance and insurance sectors have an even more steeply sloped productivity curve. This difference in value-

added per capita should translate into higher wages for finance and insurance professionals, and this prediction is confirmed by Figure 5(c) and 5(d), which contain plots comparing the average annual income of college graduate and post-graduate engineers and financiers. Finance is becoming more and more important. Therein lies the problem.

3. Moore's Law meets Murphy's Law

Moore's Law has an unspoken corollary. The rapid growth of financial innovation also means that much of this innovation is adopted without understanding the full risks. It follows, then, that financial innovation is peculiarly susceptible to Murphy's Law: "Anything that can go wrong, will go wrong." Murphy's Law originally comes from the postwar aviation industry, a time when aerospace engineers were finding ways to break the sound barrier and fly faster than the speed of sound, then a new and untested technology. Today, financial engineers are finding ways to move markets faster than the speed of thought. There is one important difference between the two industries, however. Aerospace engineers could test their designs through the efforts of brave test pilots before moving to production. Financial innovation necessarily relies on simulation and past market statistics before it is implemented into the financial system. As participants in the financial system, we ourselves are the test pilots for the accelerated pace of financial innovation.

From this perspective, perhaps the real surprise is that the financial system has not suffered more technological "prangs" in recent years, to borrow another term from the aerospace industry. But markets are resilient in a way that aircraft are not. Self-interest motivates the investor in a market to take advantage of any technological lapse in its functioning, and in doing so, the investor incorporates that information into market activity. It is only when the market innovation causes a system-wide malfunction that the market fails to compensate. Unfortunately, these breakdowns, although temporary, seem to be occurring at an accelerating rate. Moore's Law has apparently increased the risk of Murphy's Law in the financial system, and the following are some sobering examples.⁸

The quant meltdown of August 2007

Beginning on Monday 6 August 2007, and continuing through Thursday 9 August some of the most successful hedge funds in the industry suffered record losses. Despite the secretive nature of hedge funds and proprietary trading desks, the *Wall Street Journal* was able to report that some had lost 10–30% of their value in a single week.⁹ What made these losses even more extraordinary was the fact that they seemed to be concentrated almost exclusively among quantitatively managed equity market-neutral or "statistical arbitrage" hedge funds, giving rise to the event's nickname of the "quant quake" or "quant meltdown."

⁸ Some of these examples are drawn from Lo and Kirilenko (2013) with permission.

⁹ Zuckerman, Hagerty and Gauthier-Villars (2007), Sender, Kelly and Zuckerman (2007).

Although many outside observers were willing to speculate, no institution suffering such losses was willing to comment publicly on the causes of the quant meltdown. To address this lack of transparency, Khandani and Lo analysed the events of the meltdown by simulating the returns of the contrarian trading strategy of Lehmann, and Lo and MacKinlay, on the historical data.¹⁰ Their “Unwind Hypothesis” proposed that the losses during the second week of August 2007 were initially due to the forced liquidation of one or more large equity market-neutral portfolios. However, this large portfolio was not unique, but one of many portfolios that had converged on a similar selection, presumably as a result of a widely adopted financial innovation within the hedge fund industry. The price impact of this massive and sudden unwinding caused these similar but independent portfolios to experience losses. These losses in turn caused some funds to deleverage their portfolios, yielding an additional price impact that led to further losses and more deleveraging, and so on, in a deadly feedback loop. Many of the affected funds were considered to be at the vanguard of industry practice. The quant meltdown suggests that, for a time, they became victims of their financial innovation.

The flash crash

At approximately 1:32 pm Central Time, 6 May 2010, US financial markets experienced one of the most turbulent periods in their history. This period lasted all of about 33 minutes. The Dow Jones Industrial Average suffered its biggest one-day point decline on an intraday basis, at one point plunging 600 points in the space of five minutes. And the prices of some of the world’s largest companies traded at incomprehensible prices – Accenture traded at a penny a share, whereas Apple traded at \$100,000 per share. This remarkable event has been seared into the memories of investors and market-makers and, because of the speed with which it began and ended, is now known as the “flash crash.”

But the most disturbing aspect of the flash crash is that the subsequent investigation by the staffs of the Commodity Futures Trading Commission (CFTC) and the Securities and Exchange Commission (SEC) concluded that these events occurred not because of any single organisation’s failure, but rather as a result of seemingly unrelated activities across different parts of the financial system that fed on each other to generate a perfect financial storm.¹¹ In other words, there is no single “culprit” than can be punished for this debacle, nor any new regulation that can guarantee such an event will never happen again.

The joint CFTC/SEC report traced the event to an atypical automated sale of 75,000 E-mini S&P 500 June 2010 stock index futures contracts which occurred over an extremely short time period, creating a large order imbalance that apparently overwhelmed the small risk-bearing capacity of the high-frequency traders acting as market-makers. After accumulating E-mini contracts over a 10-minute interval, these high-frequency traders began to unwind their long positions, rapidly and aggressively passing contracts back and forth, like a “hot potato.” At the same time, cross-market arbitrage trading algorithms quickly propagated the price decline in E-mini futures to the markets for stock index exchange-traded funds like the Standard & Poor’s Depository Receipts S&P 500, individual stocks and listed stock options. In a manner

¹⁰ Khandani and Lo (2007, 2011), Lehmann (1990), Lo and MacKinlay (1990).

¹¹ CFTC/SEC (2010); Kirilenko, Kyle, Samadi and Tuzun (2014).

reminiscent of the 19 October 1987 stock market crash, sell orders in the futures market triggered by an automated selling program cascaded into a systemic event for the entire US financial market system. The difference is that the October 1987 crash took an entire day; the flash crash came and went in the space of a television sit-com episode.

This was the narrative as of 30 September 2010 when the joint CFTC/SEC report was published. However, the narrative has changed. On 21 April 2015, the US Department of Justice filed charges against Navinder Singh Sarao, a British national.¹² The criminal complaint, made with the CFTC, alleged that Sarao had attempted to manipulate the price of E-Mini S&P 500 futures contracts on the Chicago Mercantile Exchange, specifically using the tactic of “spoofing,” that is, transmitting orders that he intended to cancel. Sarao allegedly used a financial innovation called “dynamic layering,” reportedly convincing an automated trading software company to customise his software to submit orders to give the illusion of a deep market before they were cancelled. To quote from the Department of Justice’s affidavit, “SARAO’s activity created persistent downward pressure on the price of E-Minis. Indeed, during the dynamic layering cycle that ran from 11:17 a.m. to 1:40 p.m. [Central Time], SARAO’s offers comprised 20 to 29% of the CME’s entire E-Mini sell-side order book, significantly contributing to the order book imbalance. During that period of time alone, the E-Mini price fell by 361 basis points. In total, SARAO obtained approximately \$879,018 in net profits from trading E-Minis that day.”¹³

These charges have not yet been decided upon in a court of law, so they must necessarily remain hypothetical, but there is nothing prima facie implausible about these allegations as a possible component of an explanation for the flash crash. Even without fraudulent intent, adding to a large order imbalance in an exchange where market-makers were overwhelmed would make the conditions for a flash crash more likely. If these allegations hold, however, they show that the financial system also must be able to cope with innovations that are deliberately antagonistic to the wellbeing of the system.

The BATS and Facebook IPOs

On 23 March 2012, BATS Global Markets held its IPO. Founded in 2005 as a “Better Alternative Trading System” to the New York Stock Exchange and NASDAQ, BATS operated the third largest stock exchange in the United States at the time and did it from the suburbs of Kansas City. As one of most technologically sophisticated companies in the financial industry, BATS naturally decided to launch its IPO on its own exchange. That was a mistake. Shortly after its IPO debuted at an opening price of \$15.25, the price plunged to less than a tenth of a penny in a second and a half. Apparently, a software bug affecting stocks with ticker symbols from A to BFZZZ created an infinite loop that made these symbols inaccessible on the BATS system, including its own ticker symbol, BATS.¹⁴ No information about the glitch was made public during the day. Despite the quick deployment of a software patch by that

¹² Brush, Schoenberg and Ring (2015); CFTC (2015).

¹³ Department of Justice (2015).

¹⁴ Dornbrook (2012); Oran, Spicer, Mikolajczak and Mollenkamp (2012); Schapiro (2012).

afternoon, the confusion was so great that BATS suspended trading in its own stock, and ultimately cancelled its IPO altogether.

An even bigger glitch occurred on 18 May 2012 when the pioneering social network company, Facebook, launched the most highly anticipated IPO in recent financial history. As a company with over \$18 billion in projected sales, Facebook could have easily listed on the New York Stock Exchange alongside older blue-chip companies like IBM and Coca-Cola. Instead, Facebook chose to list on NASDAQ, quite a coup for that exchange in an era of increasingly fragmented markets. Although Facebook's opening was expected to generate huge order flows, NASDAQ prided itself on its ability to accommodate a high volume of trades so that capacity was not a concern. In fact, NASDAQ's IPO Cross software was reputed to be able to compute an opening price from a stock's initial bids and offers in less than 40 microseconds, approximately 10,000 times faster than the blink of an eye.

At the start of the Facebook IPO, demand was so heavy that it took NASDAQ's computers up to five milliseconds to calculate its opening price, more than 100 times slower than usual. As these computations were running, NASDAQ's order system allowed investors to change their orders up the moment the opening trade was printed on the tape. These few milliseconds before the print were more than enough for new orders and cancellations to enter NASDAQ's auction book, causing the IPO software to recalculate the opening trade price, during which time even more orders and cancellations entered its book, compounding the problem.¹⁵ Software engineers call this situation a "race condition"; a race between new orders and the print of an opening trade created an infinite loop that could only be broken by manual intervention, something that hundreds of hours of testing had apparently missed.

Although scheduled to begin at 11:00 am, Facebook's IPO opened a half hour late because of these delays. As of 10:50 am, traders had not yet received acknowledgments of pre-opening order cancellations or modifications. Even after NASDAQ formally opened the market, many traders still had not received these critical acknowledgements, creating more uncertainty and anxiety.¹⁶ By the time the system was reset, NASDAQ's programs were running 19 minutes behind real time. Seventy-five million shares changed hands during Facebook's opening auction, but orders totalling an additional 30 million shares took place during this 19-minute limbo. Many customer orders from both institutional and retail buyers went unfilled for hours, or were never filled at all, while other customers ended up buying more shares than they had intended.¹⁷ The SEC ultimately approved a plan for NASDAQ to pay its customers \$62 million for losses in its handling of Facebook's offering, eclipsing its achievement in handling the third largest IPO in US history.¹⁸

Knight Capital Group

At market open on 1 August 2012, the well known US broker-dealer Knight Capital Group – one of the largest equity traders in the industry at the time and among the

¹⁵ Benoit (2012); Schapiro (2012).

¹⁶ Strasburg, Ackerman and Lucchetti (2012).

¹⁷ Strasburg and Bunge (2012).

¹⁸ Strasburg and Bunge (2013).

most technologically sophisticated – issued a surge of unintended orders electronically. Many of these orders were executed, resulting in a rapid accumulation of positions “unrestricted by volume caps” that created significant swings in the prices of 148 stocks between 9:30 am and 10:00 am.¹⁹ Unable to void most of these unintentional trades, Knight Capital was forced to liquidate them at market prices, resulting in a \$457.6 million loss that wiped out its entire capital base. Its share price plunged 70% and Knight was forced to seek a rescuer; it was eventually acquired by competing broker-dealer GETCO in December 2012.

What could have caused this disaster? Knight subsequently attributed it to “a technology issue... related to a software installation that resulted in Knight sending erroneous orders into the market”. Apparently, the SEC later determined that this was the result of a program functionality called “Power Peg”, which had not been used since 2003, and had not been fully deleted from Knight’s systems.²⁰ The most surprising aspect of this incident was the fact that Knight was widely considered to be one of the best electronic market-makers in the industry, with telecommunications systems and trading algorithms far ahead of most of their competition.

The Treasury flash crash

Perhaps the most startling of the recent malfunctions took place in one of the cornerstones of the global financial system, the US Treasury market, on 15 October 2014. On that day, yields in benchmark 10-year Treasuries traded in a range of 35 basis points between market open and close, a seven-sigma intraday event with no obvious smoking gun. Like other flash crashes, much of this swing took place in a very brief interval. Market observers attributed an initial sharp decline at 8:30 am Eastern Time to sell-offs precipitated by poor US retail sales in September. Shortly after 9:33 am Eastern Time, however, the yield in 10-year Treasuries fell an additional 15 basis points to 1.86%, only to rebound to its former level, all in a space of 10 minutes.²¹

Coincidentally enough, the following day saw the monthly meeting of the Treasury Market Practices Group, sponsored by the Federal Reserve Bank of New York. The group had no immediate explanation for the flash crash in Treasuries, but hypothesised that it was driven by “large scale repositioning by leveraged investors, activities of electronic trading algorithms and dealer balance sheet and risk management constraints.”²² This was clearly a stopgap account before more information became available. Six months later, however, the New York Fed had not settled on a cause for that day’s events. Did automated trading firms “unplugging” their systems contribute to the plunge, or did they protect the market from greater volatility? Were regulatory changes that inhibited the traditional ability of dealers to buffer sudden changes in price a factor? Was there a liquidity crunch, or was the Treasuries market able to meet its orders despite the heavy trading volume? No clear answer yet exists to these and other important questions about that day’s events.²³

¹⁹ McCrank (2012); Telegraph (2012); Schapiro (2012).

²⁰ SEC (2013).

²¹ Alloway and MacKenzie (2014); Potter (2015).

²² TMPG (2014).

²³ Potter (2015).

The Bloomberg terminal outage

Financial markets increasingly rely on technologies that are not strictly part of the financial system, but whose failure may still have a systemic effect. Major exchanges now have uninterruptible power sources, multiple modes of communication and offsite backup storage in case of natural or manmade disaster. Despite this redundancy, however, when the Bloomberg terminal system was disrupted on 17 April 2015, for a period of two and a half hours, many of the system's over 300,000 subscribers were unable to function effectively, bringing transactions in some markets to a standstill. Bloomberg blamed "a combination of hardware and software failures in the network, which caused an excessive volume of network traffic. This led to customer disconnections as a result of the machines being overwhelmed."²⁴

The Bloomberg terminal system, a subscription-only data and communications network, is the financial information system of choice for many traders globally. Beginning shortly after market open in much of Europe, the terminal outage had little effect on traders in the United States, but affected markets throughout the eastern hemisphere, leading to the postponement of a multibillion buy-back of government debt by the UK Debt Management Office. Although alternative systems such as the Thomson Reuters network were available to many, these systems lacked the customised messaging capability of the Bloomberg terminal, which had developed into an important form of communication between traders. Some enterprising traders returned to making deals over the phone, a reversion to an earlier form of technology. To turn a potential financial tragedy into farce, one compelling but unverified rumour blamed the outage on a can of soda that was spilled onto a critical Bloomberg server.²⁵

4. Technology to the rescue

In her introductory speech to the SEC's Market Technology Roundtable in 2012, SEC chair Mary Schapiro condemned the increasing number of "Technology 101 issues" in the exchanges, while emphasising that contagion across markets and venues is still rare.²⁶ Solving these Technology 101 issues is a crucial first step in lowering this new form of financial risk. There is a saying in software engineering, "Given enough eyeballs, all bugs are shallow."²⁷ Presumably, more eyeballs could have prevented the simple software bug that confounded the BATS IPO and the race condition that disrupted Facebook's debut on NASDAQ.

However, as the flash crash and the quant meltdown demonstrate, it is entirely possible to cause financial disruption when all systems are operating normally. It is all too easy to imagine chains of financial contagion transmitting themselves through widely traded stocks such as Apple or Facebook, both of which have been badly disrupted by market glitches, or even more catastrophically, through the global cornerstone market for US Treasuries, whose plunge in October 2014 is still

²⁴ Cox (2015); Cox and Trivedi (2015).

²⁵ Brinded (2015).

²⁶ Schapiro (2012).

²⁷ Raymond (1999, p 30).

unexplained. Moreover, while solving Technology 101 issues is clearly important, a financial system that relies on all of its parts functioning at 100% efficiency is vulnerable to both accident and malice. The linkages made possible by technological innovation may have increased systemic financial risk in unforeseen ways, but to reduce this new form of systemic risk, the solution must be to make financial technology more robust, not to reach for an illusory perfection. Better software engineering in our financial system is analogous to improvements in our public health system to prevent the ill effects of bugs, but we also need a financial immune system that is able to adapt to circumstances to prevent system-wide catastrophes.

What do the financial failures in the preceding section have in common? The common hallmark is a coordinated response to unexpected loss. Under normal conditions, unanticipated financial losses affect market participants narrowly, eg the individual investor faced with a margin call they are unable to make. When the losers are sufficiently large in size or number, however, their responses can threaten the financial stability of the system as a whole. Unanticipated losses can cause widespread panic – in the form of flights to safety, rapid price declines and/or the evaporation of liquidity – that once triggered, is impossible to contain. Technological innovation changes the probability of these losses in unanticipated ways.

Adaptive regulation

One way that the financial system can adapt to these changing circumstances is to employ dynamic regulation in the financial markets. Consider an example from the private sector. The Chicago Mercantile Exchange, one of the world's largest organised financial exchanges, has developed dynamic margin requirements so as to protect both the exchange and market participants from default due to extreme losses.²⁸ To do this, it uses its in-house risk management system, Standard Portfolio Analysis of Risk (SPAN), a software suite originally developed in 1988, now in its fourth generation and widely adopted as an industry standard.²⁹ SPAN calculates the maximum market loss of a portfolio under multiple scenarios (typically 16; however, the number is user-defined) and then determines what the appropriate margin requirement should be.³⁰ Lo and Brennan have shown that the SPAN-calculated margin requirements for currency futures at the CME strongly correlate with recent volatility for US dollars in the euro market and other currencies, indicating that SPAN has many of the properties needed for dynamic regulation.

Risk management systems such as SPAN serve as a useful proof of concept for the importance of dynamic loss probabilities. The SPAN system is critical to the Chicago Mercantile Exchange for protecting its clearing house against defaults, and it incorporates the type of adaptive regulation that the financial system should also incorporate. However, the SPAN system is only concerned with mesoscale risks to the financial system. It is designed to protect individual clearing houses with highly liquid instruments, for which changes in volatility and price processes may be readily observed and incorporated into new margin requirements. It is not concerned with managing systemic risk, and it is difficult to see how it could be, given its

²⁸ Brennan and Lo (2014).

²⁹ CME Group (2015).

³⁰ CME Group (2010).

informational limitations. It is adaptive regulation, but at the level of the organ, not the organism.

Adaptive financial regulation needs to account for the macroeconomy. It should be informed by private sector examples like the Chicago Mercantile Exchange, and implement systems in the same spirit as SPAN, but the focus of macroprudential policies must necessarily be the entire financial system, the organism as a whole. The CME is able to treat activities outside its purview as exogenous events, while the financial system must address the endogenous nature of systemic risk and the impact of the regulatory requirements themselves.

Law is code

Therefore, to regulate the financial system as a whole, we need to better understand financial regulation as a whole. The US legal system is a working example of adaptive regulation, based on principles of common law that date back to the Middle Ages, and it incrementally changes in response to societal needs and political pressure. However, it was not designed for periods of rapid change, and many of the Founders saw a deliberative pace in legal change as a positive goal. Codification of federal law began startlingly late in American history (1926), and federal statutes are still poorly organised.

It is fruitful to think of the law as the software of the American operating system – yet if a team of software engineers were to analyse the corpus of federal law, they would see thousands of pages of poorly documented code, with a multitude of complex, spaghetti-like dependencies between individual modules.³¹ Using metrics for measuring the quality of software (see [Table 1](#)), Li, Azar, Larochelle, Hill and Lo (2015) analysed the entire text of the US legal code (all the permanent laws of the United States) and drew some sobering conclusions about its complexity and potential for unintended consequences.

One particularly informative measure is a network-based measure of complexity using the degree of “connectivity” across different sections of the US legal code, where a “connection” between two sections of the code is defined as a simple cross-reference of one section by another. Li et al (2015) cite the example of Section 37 USC § 329, which involves an incentive bonus for retired or former members of the military. This section cites exactly two other sections, 37 USC § 303a(e) (general provisions of special pay in the military), and 10 USC § 101(a)(16) (a definition of “congressional defense committees”), while 37 USC § 329 is cited by one other section, 10 USC § 641, which notes that other laws in Title 10 of the US legal code do not apply to the officers to whom the bonus in 37 USC § 329 applies. The interconnections are shown in [Figure 6](#). Now consider a much longer chain with multiple branches, some of which refer back to the section being modified. These chains may contain complex sequences of legal ramifications that even the most intelligent and knowledgeable human cannot fully grasp without some form of technological assistance.

³¹ Li et al (2015).


	Principle	Proposed Metric
1.	Conciseness: Good code should be as long as it needs to be, but no longer.	Number of words
2.	Cohesion: Modules in code should do one thing well, not multiple things badly.	Language perplexity
3.	Change: Code that exhibits large or frequent change may suggest defects.	Number of sections/subsections affected
4.	Coupling: Modular code is more robust and easier to maintain than code with unnecessary cross-dependencies.	Size of cross-reference network core versus periphery
5.	Complexity: Code with a large number of conditions, cases, and exceptions is difficult to understand and prone to error.	Number of condition statements in code (McCabe’s complexity)

Source: Li, Azar, Larochelle, Hill and Lo (2015).

The layout of these connections – often called the “network topology” in the jargon of mathematical graph theory – can also be used to construct quantitative measures of complexity. One such measure is the notion of a “strongly connected” set of nodes, defined to be a set of nodes in which there is a path from every node to every other node in the set. For example, in Figure 7 nodes B and E form a strongly connected set, but nodes (B, E, A) do not because there is no path from E to A within the subset of these three nodes.

Network representation of references to and from a section of the US legal code (37 USC § 329)

Figure 6



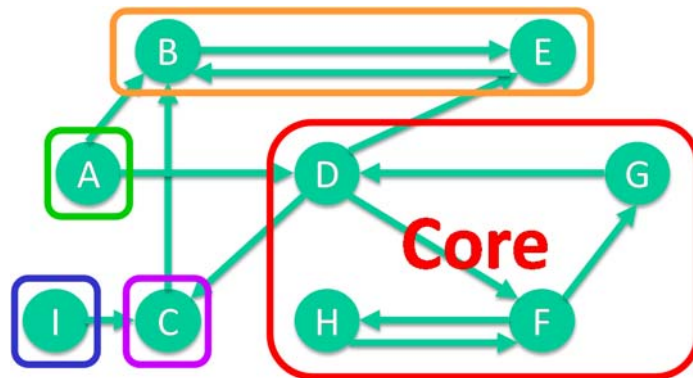
Source: Li, Azar, Larochelle, Hill and Lo (2015)

When applied to an entire network, it can be shown that the nodes can be partitioned into a finite number of disjoint subsets, each of which is strongly connected and the union of all these strongly connected subsets is the entire network.

A natural measure of complexity can then be defined as the size of the largest strongly connected subset, which is called the “core”. In Figure 7, the core is the subset (D, F, G, H) and its size is four nodes. The larger the size of the core, the more interconnected is the network; changes to one part of the core could affect every other part of the core (because there exists at least one path from every node to every other node). As the core increases in size, the possible interactions between different nodes grow exponentially.

Illustration of the concept of strong connectedness in a directed graph and the “core,” which is the largest strongly connected set of nodes

Figure 7



Source: Li, Azar, Larochele, Hill and Lo (2015).


Li et al (2015) apply this complexity measure to various parts of the US legal code and document some extraordinary levels of interconnectedness. Figure 8 displays three small examples from their more comprehensive analysis. In Figure 8(a), the network formed by the Omnibus Appropriations Act of 2009 is depicted, with blue nodes representing peripheral sections and red nodes representing the core. This network is relatively simple – a very small core surrounded by peripheral sections that are mostly isolated, indicating very few cross-references. This simplicity is not surprising in a bill that is largely a sequence of appropriations for a number of unrelated programmes.

However, Figure 8(b) shows a much more complex network representing the Dodd-Frank Wall Street Reform Act, a piece of legislation spanning 2,319 pages that was passed on 21 July 2010. Of the 390 rule-making requirements imposed by Dodd-Frank, only 267 have been satisfied by finalised rules as of end-December 2015.³² However, the complexity of Dodd-Frank does not compare with that of Title 12 of the US legal code, which governs the entire banking industry; its network structure is displayed in Figure 8(c). With an extremely large core and many connections between the core and the periphery, it is easy to see how small changes can lead to unpredictable and unintended consequences in other parts of the network.

³² Davis Polk (2015, p2).

Core-periphery network maps of: (a) sections of the US legal code modified by the Omnibus Appropriations Act of 2009; (b) sections of the US legal code modified by the Dodd-Frank Wall Street Reform Act; and (c) Title 12 of the US legal code (Banks and Banking). Blue dots indicate peripheral sections, red dots indicate the core.

Figure 8



Source: Li, Azar, Larochelle, Hill and Lo (2015)

These new tools provide an X-ray of the hidden structures within current banking regulation. It is perhaps unsurprising that the core sections on banking regulation have to do with the powers of the corporation, insurance funds and holding companies since that is where the vast majority of financial assets are organised. These sections of the law are of critical importance to the US financial system. To pursue the software analogy further, any effort to reform banking regulation should begin with a systematic “refactoring” and simplification of these sections, improving their internal structure without altering their external behaviour, rather than adding increasingly complicated patches to the law whose systemic effects are unknown.

Transparency vs privacy

One compelling concern about a systemic, macroprudential approach to financial regulation is financial privacy. Most of the financial industry relies on unpatentable business processes to make a living, as Myron Scholes discovered when he confronted Texas Instruments about its infringement on the Black-Scholes formula. As a result, the financial industry necessarily practises security through obscurity, preferring to use trade secrets to protect its intellectual property. Hedge funds and proprietary trading desks take this to an extreme, essentially serving as “black boxes” for investors, as opaque as the law will allow. However, even the average financial institution has a need to limit disclosure of their business processes, methods and data, if only to protect the privacy of their clients. Accordingly, government policy has tread carefully on the financial industry’s disclosure requirements.

How can financial institutions provide the information that adaptive regulation requires, without feeling burdened or threatened by regulatory intrusion? One solution is to make the interactions between financial institutions and regulators secret. However, this fails to provide the public with the transparency about systemic risk it increasingly wants from the financial system, while putting an enormous burden on regulators.

Fortunately, developments in cryptography, made possible by the acceleration in computing power under Moore’s Law, show a way to solve to this dilemma. A well known technique from the computer science literature called “secure multi-party computation” provides an elegant solution to the need for sharing certain types of

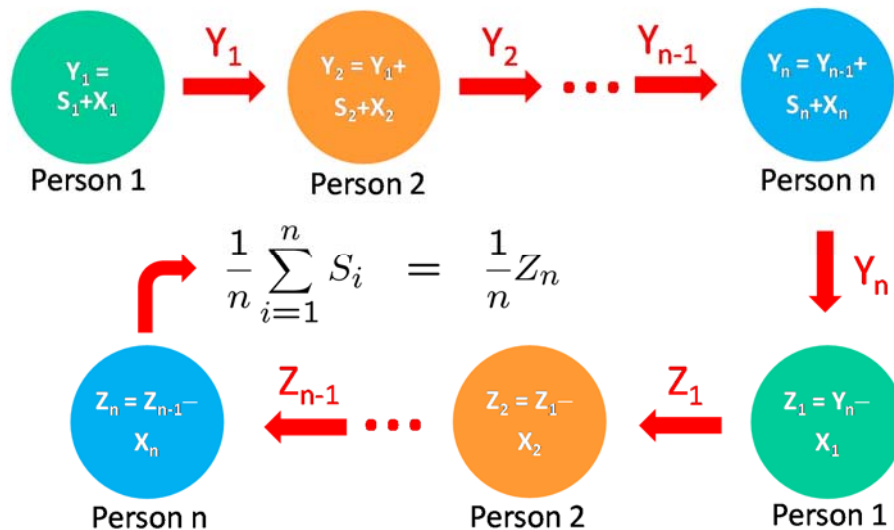
information while preserving the confidentiality of each party's data. A simple illustration of this technique involves the indelicate task of computing the average salary of a roomful of n conference attendees, a very intrusive computation given the sensitive nature of individual salary figures.

Suppose person 1 takes his salary S_1 and adds to it a random number of his choosing X_1 to obtain the sum $Y_1 = S_1 + X_1$ and then shares this sum (but not the components) with person 2. Person 2 then performs the same calculation, adding a random number of her choosing, X_2 , to her salary S_2 and then adding these two values to person 1's information to obtain $Y_2 = Y_1 + S_2 + X_2$. She then passes Y_2 to person 3 who adds his random number and salary to it before passing it to the next person, and so on. This process continues from one person to the next until the last person, n , adds his salary and random number to it, yielding $Y_n = S_1 + S_2 + \dots + S_n + X_1 + X_2 + \dots + X_n$.

Now suppose person n passes this sum to person 1 and asks him to subtract his random number X_1 from it before passing it to person 2. Person 2 does the same operation, subtracting her random number X_2 from the cumulative sum before passing the value to person 3, and so on. Once the process returns to person n , who subtracts his random number, X_n , from the cumulative sum, the value remaining is the sum of all the salaries $S_1 + S_2 + \dots + S_n$, which, when divided by the number n which is observable, yields the average salary in the room. Figure 9 summarises this simple algorithm. At no point during this process did anyone have to reveal his or her private information, yet by the end of the process, the average salary was computed. Such algorithms are the essence of secure multi-party computation.

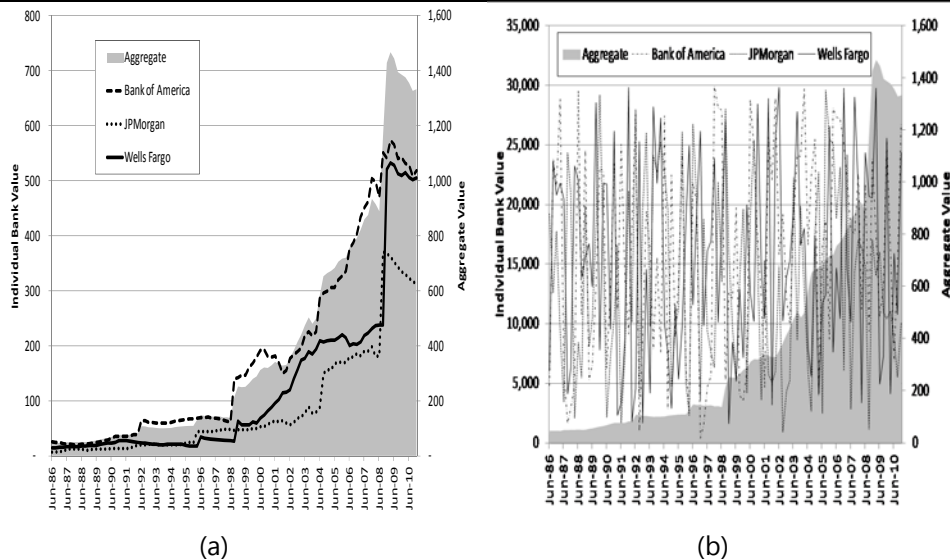
Illustration of a simple secure multi-party computation algorithm for computing the average salary of a group of n individuals without requiring any individual to reveal his or her salary.

Figure 9



Example of secure multi-party computation of the aggregate size of the real-estate loan portfolios of Bank of America, JP Morgan and Wells Fargo. Graph (a) contains the raw time series for the three individual banks as well as the aggregate sum; graph (b) contains three encrypted time series which, when summed, yields the same aggregate sum as the unencrypted data.

Figure 10



Source: Abbe, Khandani and Lo (2012)

Now, of course, two participants could easily collude so as to infer the salary of a third individual. For example, if persons 1 and 3 compared their cumulative sums before and after person 2 subtracted her random number, they could infer her random number and deduce her salary. However, there are simple ways of constructing cheat-proof algorithms that allow all parties to share certain kinds of information while keeping their raw data confidential. In Abbe, Khandani and Lo (2012), we construct secure multi-party computation algorithms that can be used to encrypt proprietary information from banks, broker-dealers and other financial institutions while still allowing regulators to compute aggregate risk measures such as sums, averages, value-at-risk, loss probabilities and Herfindahl indexes.

Figure 10 contains a concrete illustration of this technology applied to the sizes of the real-estate loan portfolios of Bank of America, JP Morgan and Wells Fargo. Figure 10(a) contains the individual time series for these three institutions (the line graphs), which are the proprietary information of each institution and only publicly disclosed with a lag. From a systemic risk perspective, the individual values are of less importance than the aggregate sum, depicted by the area graph in Figure 10(a). Using a particular algorithm designed just for this purpose, Abbe et al (2012) show that the individual time series can be encrypted, as in the line graphs in Figure 10(b), yet the sum of the encrypted time series yields the very same bar graph as in Figure 10(a). Aggregate sums can be shared by financial institutions while maintaining the privacy of each institution.

Using secure multi-party computation tools, it is possible to construct mathematical protocols that allow aggregate measures to be computed without

revealing any of the individual components of that aggregate.³³ Thus, the aggregate risk exposures of a group of financial institutions can be calculated and made public, while preserving the privacy of any individual financial institution. This method is ideal for use in macroprudential regulation. Furthermore, since the cost-benefit ratio to financial institutions is so low, there is even reason to believe that the financial industry may adopt such disclosures voluntarily, if informational incentives are structured correctly.

Of course, techniques like secure multi-party computation certainly do not eliminate the need for regulations or regulators – for example, there is no way to ensure that institutions report truthfully other than through periodic examination – but they can lower the economic cost of sharing certain types of information and provide incentives for the private sector to do so voluntarily. If financial institutions can maintain the privacy of their trade secrets while simultaneously sharing information that leads to more accurate measures of threats to financial stability, they stand to benefit as much as the regulators and the public.

5. Conclusion

These examples show how technology can reduce the additional systemic financial risk brought about by technological innovation. This is not a paradox. Rather, it is a consequence of the symbiotic relationship between finance and technology. Not very long ago, the financial markets were the most informationally intensive places on Earth, the collective intelligence of the markets incorporating the world's data into prices faster than any computer of the time. Today, the financial markets are one informationally intensive system among many, in a symbiotic relationship with search engines, social networks, messaging systems and the growing colossus of Big Data.

In this brave new networked world, we will need to adopt a more advanced systems approach to financial technology. No financial engineer or programmer or designer of exchange servers should assume that a new product will function in isolation, but should rather imagine a changing financial environment where past statistics almost certainly will not apply. Similarly, no financial regulator should assume that an innovator will not find a way to circumvent a regulation, perhaps in a more disruptive way than what the regulation originally intended to prevent. To return to the analogy of software engineering, perhaps we should be assembling tools for financial system administrators to monitor and troubleshoot problems in the markets, similar to the way a sysadmin monitors and troubleshoots problems in a computer system.

To do this effectively, however, we need more and better information about the operation of financial markets. Going back to the example of the flash crash, the CFTC investigators were unable to find signs of Sarao's alleged activities because they were only given a list of completed transactions. "Spoofing," however, cancels the transactions before they are executed, leaving no evidence in the market print. All important market failures and events need to be analysed scrupulously and no data must be withheld from investigators.

³³ Abbe, Khandani and Lo (2012).

One potential model for this scrupulous form of analysis already exists.³⁴ The National Transportation Safety Board (NTSB) has an excellent track record in analysing and determining the causes of transportation accidents in the United States. The NTSB has no regulatory authority, freeing the agency to criticise regulations and regulators that it believes may have contributed to the cause of an accident. In addition, the NTSB has subpoena power to obtain the information it needs to make a full analysis of an accident. The NTSB's accident report is not admissible as evidence in lawsuits for civil damages, which allows the stakeholders to be much more candid about their role in an accident. As a result, the NTSB can address the systemic causes of an accident, as it did in its report on USAir Flight 405, which put the ultimate cause of that flight's crash in 1993 on a system-wide failure in de-icing procedures.³⁵ Under an NTSB-like system, stakeholders in financial system failures would have less reason not to be candid about their possible shortcomings – but if this is still insufficient, secure multi-party methods may allow financial information to be observed without identifying specific financial institutions, in a form of cryptographic redaction.

Better information about financial system failures will require better tools to remedy those failures. Here, mention must be made of the Food and Drug Administration's (FDA's) call for greater "regulatory science".³⁶ Like the financial system, the human body is also an immensely complicated and hyperconnected assemblage of disparate parts. The FDA's mission for over a century has been to protect that body by prohibiting certain dangerous or fraudulent products and by testing the efficacy of others. To continue to do so effectively in the future, the FDA has proposed a broad strategy to harness science to serve regulation. For example, many models and assays currently used in toxicology are of limited accuracy in predicting adverse events in human beings. They are still in use, however, because they are still considered best practice – a state of affairs that should be uncomfortably familiar to many financial regulators. The FDA's regulatory science proposal would clearly define the reliability of these tests and their limitations – also something that should be familiar to financial regulators.

The global financial system has experienced exponential growth as a result of its intimate, symbiotic relationship with Moore's Law and new technologies. This has resulted in an unfortunate expansion of new forms of systemic risk, as new linkages made possible by these new technologies changed previously well understood probabilities of risks in unexpected ways. However, the same technologies that created these linkages also allow us to monitor and supervise the financial system in ways that would have been unthinkable in earlier years. Because of Moore's Law, it is now possible to regulate margin requirements dynamically, analyse financial regulation as though it were a recalcitrant piece of computer code and oversee aggregate financial data publicly without violating financial privacy or confidentiality requirements. Although it is too soon to tell, it may be that the past few years have been a temporary blip in the symbiotic Red Queen's Race between finance and technology. Just as technology can add risk to a system, technology can remove it as well.

³⁴ Fielding et al (2011).

³⁵ NTSB (1993).

³⁶ FDA (2011).

References

Abbe, E, A Khandani and A W Lo (2012): "Privacy-preserving methods for sharing financial risk exposures", *American Economic Review: Papers and Proceedings*, vol 102, pp 65–70.

Alloway, T and M MacKenzie (2014): "Bonds: anatomy of a market meltdown", *Financial Times*, 17 November.

Benoit, D (2012): "Nasdaq's blow-by-blow on what happened to Facebook", *Wall Street Journal*, Deal Journal, 21 May, <http://blogs.wsj.com/deals/2012/05/21/nasdaqs-blow-by-blow-on-what-happened-to-facebook/>.

Black, F and M Scholes (1973): "The pricing of options and corporate liabilities", *Journal of Political Economy*, vol 81, pp 637–54.

Brinded, L (2015): "RUMOUR: the Bloomberg outage was caused by a spilled can of Coke" *Business Insider*, 17 April, www.businessinsider.com/reason-for-the-global-bloomberg-outage-2015-4.

Brennan, T and A W Lo (2014): "Dynamic loss probabilities and implications for financial regulation", *Yale Journal on Regulation*, no 31, pp 667–94.

Brush, S, T Schoenburg and S Ring (2015): "How a mystery trader with an algorithm may have caused the flash crash", *Bloomberg*, 21 April, www.bloomberg.com/news/articles/2015-04-22/mystery-trader-armed-with-algorithms-rewrites-flash-crash-story.

CFTC (Commodity Futures Trading Commission) (2015): *US Commodity Futures Trading Commission v Nav Sarao Futures Limited PLC and Navinder Singh Sarao*, www.cftc.gov/ucm/groups/public/@lrenforcementactions/documents/legalpleading/enfsaraocomplaint041715.pdf.

CFTC/SEC (Commodity Futures Trading Commission/Securities and Exchange Commission) (2010): *Preliminary Findings Regarding the Market Events of May 6, 2010*, Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues, 18 May, www.sec.gov/sec-cftc-prelimreport.pdf

CME Group (2010): CME SPAN: Standard Portfolio Analysis of Risk, www.cmegroup.com/clearing/files/span-methodology.pdf.

——— (2015): Risk Management Overview. Available at: www.cmegroup.com/clearing/risk-management

Cox, J (2015): "Deals postponed after Bloomberg terminals go down globally", *Wall Street Journal*, 17 April, <http://blogs.wsj.com/moneybeat/2015/04/17/deals-postponed-after-bloomberg-terminals-go-down-in-europe-and-asia/>.

Cox, Josie and Anjani Trivedi. 2015. "Bloomberg Terminals Go Down Globally." *Wall Street Journal*. April 17. <http://www.wsj.com/articles/bloomberg-terminals-go-down-globally-1429262782>.

Davis P, (2015): Dodd Frank Progress Report: Fourth Quarter 2015, Davis Polk, http://prod.davispolk.com/sites/default/files/Q4_2015_Dodd-Frank_Progress_Report.pdf (accessed 16 January 2016).

Department of Justice (2015): *Criminal complaint, United States v Navinder Singh Sarao*, 11 February, www.justice.gov/sites/default/files/opa/press-releases/attachments/2015/04/21/sarao_criminal_complaint.pdf.

Dornbrook, J (2012): "BATS CEO: We solved the problem, weren't dependent on the IPO", *Kansas City Business Journal*, 29 March, www.bizjournals.com/kansascity/news/2012/03/29/bats-ceo-we-solved-the-problem.html.

FDA (Food and Drug Administration) (2011): *Advancing Regulatory Science at FDA*, www.fda.gov/downloads/ScienceResearch/SpecialTopics/RegulatoryScience/UCM268225.pdf.

Fielding, E, A W Lo, and J Yang (2011): "The National Transportation Safety Board: A model for systemic risk management", *Journal Of Investment Management*, vol 9, pp 17–49.

Khandani, A and A W Lo (2007): "What happened to the quants in August 2007?", *Journal of Investment Management*, vol 5, pp 29–78.

——— (2011): "What happened to the quants in August 2007?: Evidence from factors and transactions data", *Journal of Financial Markets*, vol 14, pp 1—46.

Kirilenko, A and A W Lo (2013): "Moore's Law versus Murphy's Law: Algorithmic trading and its discontents", *Journal of Economic Perspectives*, vol 27, pp 51–72.

Kirilenko, A, A Kyle, M Samadi and T Tuzun (2014): "The flash crash: the impact of high frequency trading on an electronic market", http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1686004.

Lehmann, B (1990): "Fads, martingales, and market efficiency", *Quarterly Journal of Economics*, vol 105, pp 1–28.

Li, W, P Azar, D Larochelle, P Hill and A W Lo (2015): "Law is code: a software engineering approach to analyzing the United States code", *Journal of Business & Technology Law*, vol 10, no 2.

Lim, T and A W. Lo, (2006): *The Derivatives Sourcebook*, Now Publishers.

Lo, A W, (2013): "Introduction to Volume 5 of the Annual Review of Financial Economics", *Annual Review of Financial Economics*, vol 5, pp 1–7.

Lo, A W and C MacKinlay (1990): "When are contrarian profits due to stock market overreaction?", *Review of Financial Studies*, vol 3, pp 175–205.

Merton, R (1973): "Theory of rational option pricing", *Bell Journal of Economics and Management Science*, vol 4, pp 141–83.

McCrack, J (2012): "Knight glitch likely to lead to regulatory changes: CEO", *Reuters* 11 September, www.reuters.com/article/2012/09/11/us-knight-ceo-idUSBRE88A1GW20120911.

Moore, G (1965): "Cramming more components onto integrated circuits", *Electronics*, vol 38, pp 114–7.

NTSB (National Transportation Safety Board) (1993): "Takeoff stall in icing conditions, USAir Flight 405 Fokker F-28, N485US, LaGuardia Airport, Flushing, New York, March 22, 1992", NTSB/AAR-93/02.

Oran, O, J Spicer, C Mikolajczak and C Mollenkamp (2012): "BATS Exchange withdraws IPO after stumbles" *Reuters* March 24 www.reuters.com/article/2012/03/24/us-bats-trading-idUSBRE82M0W020120324.

Philippon, T and A Reshef (2009): "Wages and human capital in the US financial industry: 1909–2006", *NBER Working Paper*, no 14655, www.nber.org/papers/w14644.

Potter, S (2015): "Challenges posed by the evolution of the Treasury market", remarks at the 2015 Primary Dealer Meeting, New York City, 13 April www.newyorkfed.org/newsevents/speeches/2015/pot150413.html.

Raymond, E (1999): *The cathedral & the bazaar: musings on Linux and open source by an accidental revolutionary*, O'Reilly.

Schapiro, M (2012): "Introductory remarks at SEC's market technology roundtable", 2 October, www.sec.gov/news/speech/2012/spch100212mls.htm.

Scholes, M (2006): "Derivatives in a dynamic environment", in T Lim, A W Lo, R Merton, M Scholes and M Haugh (eds), 2006, *The Derivatives sourcebook*, Now Publishers.

SEC (Securities and Exchange Commission) (2013): Administrative proceeding in the matter of Knight Capital Americas LLC *Order Instituting Administrative And Cease-And-Desist Proceedings, Pursuant To Sections 15(B) and 21c Of The Securities Exchange Act Of 1934, Making Findings, And Imposing Remedial Sanctions And A Cease-And-Desist Order*, 16 October, www.sec.gov/litigation/admin/2013/34-70694.pdf.

Sender, H, K Kelly, and G Zuckerman (2007): "Goldman wagers on cash infusion to show resolve", *Wall Street Journal*, 14 August, p A1.

Strasburg, J, A Ackerman and A Lucchetti (2012): "Nasdaq CEO lost touch amid Facebook chaos", *Wall Street Journal*, 11 June, p A1.

Strasburg, J and J Bunge (2012): "Social network's debut on Nasdaq disrupted by technical glitches, trader confusion", *Wall Street Journal*, 19 May, p A2.

——— (2013): "Nasdaq is still on hook as SEC backs payout for Facebook IPO", *Wall Street Journal*, 25 March.

The Daily Telegraph (2012): "Knight Capital's \$440m trading loss 'caused by disused software'", 14 August, www.telegraph.co.uk/finance/newsbysector/banksandfinance/9475292/Knight-Capitals-440m-trading-loss-caused-by-disused-software.html.

Treasury Market Practices Group (2014): TMPG Meeting Minutes, 16 October, www.newyorkfed.org/tmpg/october_minutes_2014.pdf.

Zuckerman, G, J Hagerty and D Gauthier-Villars (2007): "Impact of mortgage crisis spreads, Dow tumbles 2.8% as fallout intensifies; moves by central banks", *Wall Street Journal*, 10 August, p A1.

Comments by Darrell Duffie¹

Unfinished work on CCP failure resolution planning

A central counterparty (CCP) is a financial market utility that lowers counterparty default risk on each trade that it clears by becoming the buyer to each seller, and the seller to each buyer. The main regulatory goal for this central clearing technology is to enhance financial stability by setting uniform, high, and transparent standards for counterparty risk mitigation, and by lowering the risk of propagation of systemically important defaults. The “new normal” for the regulation of financial markets, globally, is that standard over-the-counter (OTC) derivatives are now centrally cleared by CCPs. Central clearing is also heavily used in exchange-based markets, and is increasingly applied in other financial markets, such as those for repurchase agreements.

A key concern is how to manage the risk of a CCP failure. When it becomes unable to meet its obligations, a CCP could be forced into a normal insolvency process such as bankruptcy, or alternatively into a government-administered failure resolution process. The ongoing global regulatory reform of financial markets envisions effective failure resolution processes for all systemically important financial firms. Yet, despite this intent, and despite the reality that central clearing is now mandated for a vast amount of derivatives under the G20 Pittsburgh Declaration of 2009,² there is still slim progress on the design of specific government-administered failure resolution plans for CCPs. In some jurisdictions, even though the principles have been established, government entities responsible for administering the actual failures of CCPs have yet to provide specific implementing operational plans or procedures. These plans and procedures should include methods for (1) deciding when to override the contractual loss allocation process, (2) how to obtain the liquidity required during the failure resolution process, (3) how to allocate final losses, and (3) how to provide for continuity of clearing services (whether at the same CCP, or at an alternative venue).

How CCPs recover from clearing member defaults

A CCP’s balance sheet is quite different from those of other major types of systemically important financial institution such as banks, broker-dealers and insurance companies. The bulk of the financial risk of a CCP is not represented by its own legal assets and liabilities. Rather, a CCP is essentially a nexus of contracts by which its clearing members net and mutualise their counterparty default risk. In the

¹ Dean Witter Distinguished Professor of Finance, Graduate School of Business, Stanford University.

This note summarises the discussion of CCP failure resolution that I provided during the session: “Has technology changed the nature of risks in financial markets?” The author for this session was Professor Andrew Lo. In addition to the topic of this note, this session covered a range of other technology issues. This note is based in part on D Duffie (2016). For other related work, see www.stanford.edu/~duffie.

² See G20 (2009).

normal course of business, the daily payment obligations of a CCP automatically sum to zero. Because of this, a CCP tends to have tiny amounts of equity and conventional debt relative to its largest potential clearing obligations. Most of the tail risk of a CCP is allocated to its clearing members. The clearing members tend to be systemically important. Moreover, clearing members tend to be members of many different CCPs.

When the market value of a centrally cleared derivatives contract increases on a given day, any clearing member who is a buyer of that contract type collects a variation margin payment from its CCP on the next day, equal to the assessed change in market value of the position. Any seller is likewise required to make a variation margin payment to the CCP. Because the total amounts of cleared bought and sold contracts are identical, the CCP's positions are exactly balanced, long against short, leaving the CCP with zero net payment obligations. If, however, one or more clearing members fail to meet their payment obligations, the CCP has unbalanced exposures and must find the resources necessary to liquidate the failed positions and rebalance itself. If it cannot, its failure must be resolved.

The greatest associated risks are (a) contagion, by which the failure of a clearing member could cause the CCP to fail to meet its obligations to other systemically clearing members, (b) firesales of collateral or derivatives contracts, exacerbating broad market volatility, and (c) loss of continuity of critical clearing services on which the financial system has come to depend.

As an example of the significant counterparty risk managed by a single CCP, SwapClear, a cross-border central counterparty operated by LCH.Clearnet, currently has a total notional amount of cleared interest-rate swaps of approximately \$360 trillion.³ Notional positions do not translate directly to risk. More commensurate with the amount of the risk managed by a CCP is the total amount of balances held at the CCP by its clearing members, which constitute mainly the funds available to cover the defaults of clearing members. At the end of 2014, this amount for LCH Clearnet Group was €407 billion.⁴ Large CCPs are systemically important, and are becoming ever more important as the implementation of new regulations forces more and more positions into CCPs.

The failure of a major CCP would probably come at an extremely stressful moment because it is most likely to be precipitated by the failure of one or more systemically important clearing members, who would probably have also failed to meet their payment obligations to many other major financial firms, including other CCPs. Ironically, the better is the quality and depth of the risk-management resources of a CCP, the more likely it is that its failure could only have been caused by the collapse of extremely large clearing members, and probably by more than one of them. Under CPSS-IOSCO principles, a globally systemically important CCP must have the resources necessary to cover the failures of its two largest clearing members.⁵

At the failure of one or more clearing members of a CCP, if the cost of liquidating the failed members' positions exceeds the margin and default guarantee funds

³ See www.lchclearnet.com/swaps/swapclear_for_clearing_members/.

⁴ See Consolidated Financial Statements of LCH Clearnet Group, p 13, www.lchclearnet.com/documents/731485/762550/lch+group+financial+statements+2014+signed.pdf/0bb5f0db-8c4b-4120-9cdc-439eea1129c9.

⁵ See Principle 4 of Committee on Payment and Settlement Systems-International Organization of Securities Commissions, "Principles for Financial Market Infrastructures", April 2012.

provided by the failed members, the surviving members and the CCP operator absorb the remaining cost of liquidating the failed positions, if that is actually possible with the available resources. While the contribution of CCP capital to the default management resources is not negligible, it is typically most important as a means of giving the CCP operator some “skin-in-the-game” incentive to design and manage the CCP safely.

In a recovery process, a CCP might assign losses to its surviving members (and perhaps their clients) in a manner that causes significant distress costs. The existence of a contractual recovery approach that avoids the insolvency of a CCP does not imply that the contractual recovery approach should be followed to its end regardless of the situation. There may be unforeseen circumstances in which total distress costs can be lowered by winding down or restructuring a CCP with a procedure that overrides contracts, such as bankruptcy or a government-administered resolution process.

The “waterfall” of recovery resources of some CCPs extends beyond the guarantee fund by permitting the CCP to contractually restructure its clearing payment obligations to clearing members.

One such procedure is “variation margin gains haircutting” (VMGH). By this approach, the CCP can conserve or accumulate cash by cancelling or reducing the variation margin payments that it would otherwise have been required to make to clearing members. At the same time, the CCP collects 100% of the variation margin payments that it is due to receive from clearing members.⁶ Beyond its role as a short-term liquidity backstop, VMGH could in some cases continue until the CCP has enough resources to pay for the liquidation of failed positions. There is no assurance that VMGH would be sufficient to entirely rebalance the CCP, although experts believe that VMGH would suffice in most scenarios.⁷ Those clearing members suffering losses from VMGH could in principle be given compensating claims, for example equity or debt issued by the CCP.

Another potential contractual restructuring approach is a “tear-up”, by which the CCP could cancel some or all of its outstanding notional derivatives positions with selected clearing members. For example, suppose the failure of a clearing member has left the CCP with a net short position in some specific class of derivatives that is 90% of the total of its outstanding long positions. In this case, the CCP could, assuming that it has the necessary contractual right, rebalance its exposure by stipulating that all long positions shall henceforth be 90% of their former notional size. An alternative is to simply tear up 100% of all outstanding positions of the affected type.

Variation margin gain haircuts and partial tear-ups have the disadvantage of sharing losses unpredictably, given that it would be difficult to predict much more than a day in advance whether it would be long or short position holders that would be allocated the losses. This is unlike the situation facing normal creditors, who know they are line for losses at the borrower’s default, and know the priority order in which they will take losses. In terms of sharing distress costs, one would prefer to have losses

⁶ A version of this approach has been adopted by the Japanese Securities Clearing Association. See Japanese Securities Association, www.jscc.co.jp/en/cash/irs/loss.html.

⁷ For more details on VMGH, see Elliott (2013), ISDA (2013), LCH (2015), and Singh (2014).

borne by all CCP members, perhaps pro rata with some measure of the expected amount of potential loss that a clearing member would impose on the CCP, conditional on both the failure of the clearing member and the failure of the CCP.⁸

On a risk-corrected basis, this might suggest end-of-waterfall loss-sharing that is proportional to total initial margins. In the same way as VMGH and tear-ups, loss-sharing in proportion to total initial margins would also encourage clearing members to reduce their positions with weak CCPs. Initial margins, moreover, are held by the CCP in the form of "paid-in" assets, usually in a liquid form. Unfortunately, there is no obvious method for exploiting assets of this sort during the recovery process. Initial margin funds are the property of clearing members, and are not legally accessible to CCPs, absent voluntary contracting that would make them available at the end of the default-management waterfall.⁹

It is not clear to me why CCPs and their clearing members prefer to use VMGH or tear-ups rather than to adjust their clearing agreements so as to allow legal end-of-waterfall access to initial margin funds. In any case, more predictable loss-sharing is normally more efficient. There are no clear incentive benefits associated with disproportionate and unpredictable loss-sharing by clearing members who happen to be buyers, or who happen to be sellers. Moreover, economic principles suggest that it is better for a clearing member to suffer a moderate loss with certainty when a CCP fails to meet its clearing obligations, than to "flip a coin" to determine whether the size of the loss is zero or not. The marginal cost to a clearing member of bearing an incremental unit of unexpected loss is normally increasing in the total amount of loss, a "convexity effect" that suggests losses should be shared across all clearing members, pro rata to the loss exposures they impose on the CCP. This is one of the reasons that CCPs are supposed to ensure adequately sized paid-in guarantee funds, which do share losses predictably and broadly. When the default guarantee fund is revealed to be inadequate, and when it is deemed appropriate to attempt recovery through further contractual loss-sharing rather than failure resolution, I have seen persuasive arguments for switching to a form of unequal and unpredictable loss-sharing.

Principles for government-administered CCP failure resolution

Failure-resolution administrators may one day be faced with a decision of whether and when to halt the contractual recovery process of a distressed or failed CCP, forcing the CCP into an administered resolution. The right to do so exists in principle in the United Kingdom, and arguably in the United States,¹⁰ and will likely soon exist

⁸ Both conditioning events are relevant, as explained by Dembo, Deuschel and Duffie (2004).

⁹ See European Parliament (2012), which states that "A CCP shall have a right of use relating to the margins or default fund contributions collected via a security financial collateral arrangement, within the meaning of Article 2(1)(c) of Directive 2002/47/EC of the European Parliament and of the Council of 6 June 2002 on financial collateral arrangements provided that the use of such arrangements is provided for in its operating rules."

¹⁰ In the United States, it seems likely that Title II of the Dodd-Frank Act assigns the administration of the failure resolution process of systemically important CCPs to the Federal Deposit Insurance

throughout the European Union, whenever financial stability is threatened, even if the CCP is currently meeting its contractual obligations.

Failure resolution procedures should be designed so as to minimise the total expected distress costs of all market participants, including clearing members, CCP operators, as well as unrelated market participants and taxpayers that could suffer from failure spillover costs.

While the government bailout of a systemically important CCP should not be legally impossible, reliance on government capital should not be part of the failure resolution design, given the attendant moral hazard. In order to align incentives in a socially efficient manner, the CCP operator and its clearing members should expect that they are on the hook for all of the losses, one way or another.

The key questions facing the failure resolution administrator are: (a) how to efficiently allocate the CCP's losses, (b) how to mitigate firesales, and (c) how to arrange for the prompt continuation of clearing services.

At the point of resolution of a CCP, most or all of its waterfall of contractually available resources has likely been exhausted. As the resolution procedure begins, the CCP may therefore have limited remaining resources with which to restructure its obligations to clearing members.

In principle, a resolution authority addressing a CCP in this financial condition could simply declare that the CCP will discontinue clearing and return any remaining assets to its clearing members, pro rata to unmet clearing obligations. This liquidation approach is more easily contemplated when continuity of clearing services can be provided by an alternative CCP handling the same classes of trades. During a wind-down, a CCP may need to haircut variation margin gains as a cash management strategy.

The alternatives to liquidating an insolvent CCP are:

1. Reorganise the CCP through some combination of new capital injections and restructuring of its clearing obligations. The debt of the CCP can also be restructured, but in practice CCPs do not usually have much debt.
2. Transfer the clearing obligations of the CCP, if necessary after some restructuring, to another existing CCP or to a "bridge" CCP.

There is a critical decision as to when to trigger this form of resolution, balancing the harm caused by lack of access of the original CCP to additional default fund contributions from clearing members that could prevent a CCP failure, relative to the harm caused by draining capital from systemically important clearing members without necessarily the prospect of emerging with a viable CCP.

Even if pre-funded and escrowed funds are available to set up a new guarantee fund,¹¹ the CCP (or its successor or bridge) may need additional capital to cover the

Corporation (FDIC). (Whether this is in fact the case, however, is not a completely settled matter, as explained by Steigerwald and DeCarlo (2014).) Assuming that Title II applies, the FDIC can become the receiver of a CCP in the event that the Secretary of the Treasury, the Federal Reserve Board, and the FDIC find that there would be otherwise be a risk of financial instability. In that case, the FDIC could liquidate the CCP, or alternatively could assign its assets and obligations to another CCP or to a "bridge", which in principle could become a successor CCP. In Europe, directives, principles, and plans or enabling legislation are described in European Commission (2012), European Securities and Markets Authority (2012), European Union (2012), European Union (2014) and HM Treasury (2014).

¹¹ For example, see JPMorgan Chase (2014).

cost of liquidating failed positions. Several approaches for this have been suggested, including tear-ups and VMGH, along the same lines that could be applied contractually in a recovery process.

Until the completion of the resolution process, the failure administrator needs access to liquidity to manage cash payments. In addition to its immediately available cash, a CCP could obtain liquidity through financing that is secured by non-cash assets. These include assets held in the default-management waterfall as initial margin (to the extent legally permitted) or paid-in guarantee fund contributions, and claims to future contributions to the default guarantee fund. A further potential source of liquidity is VMGH, to the extent legally permitted. Some government sources of liquidity may be available,¹² including in the form of central bank lending of last resort. In the United Kingdom, central bank liquidity to financial market infrastructure is explicitly available.¹³

If restructuring is an option, failure resolution administration authorities could be given the legal authority to apply VMGH or tear-ups, even if that option is not contractually recognised in clearing agreements. Under a US Title II failure administration procedure, the FDIC has the legal right to reject contracts, provided that rejection is not applied selectively across contracts with the same counterparty. It is not clear whether or how the FDIC could conduct VMGH or partial tear-ups. To this point, the FDIC has not described the failure-resolution strategies that it would use in the case of CCPs.

Outline of a CCP failure resolution process

Based on our discussion, an administrative CCP failure resolution process could have the following basic steps.

1. Verify the conditions for initiating a failure resolution process and initiate the process. Consult with relevant foreign authorities.
2. Stay the termination of clearing agreements and other contracts, with the likely exception of interoperability agreements with other CCPs.
3. Replace the senior CCP management if that is deemed appropriate, while taking steps to retain key personnel.
4. Assess the immediate cash needs of the CCP and the available sources of liquidity. Make a plan to access liquidity in priority order. Obtain the necessary cash, whether for orderly wind down or for continuity of clearing.
5. In the event of insufficient cash, interrupt payments to clearing members as legally feasible under contracts or stays, and as appropriate to mitigating the aggregate losses of all parties, including unrelated market participants.
6. Enter claims on the estates of failed clearing members.
7. In a restructuring aimed at the continuation of clearing services:

¹² In the United States, restrictions on government sources of liquidity are outlined by Skeel (2014).

¹³ See Bank of England (2014a,b).

- a. If the CCP undergoing resolution is not suitable for restructuring and continuation as a single entity, then transfer unrejected clearing agreements and other CCP property and agreements to a bridge or other successor CCP.
 - b. Replenish the default guarantee fund, using pre-funded assets as available and additional replenishment contributions from clearing members to the extent permitted by contract and judged systemically safe from the viewpoint of contagion risk.
 - c. Rebalance the derivatives positions of the CCP. For example, conduct tear-ups or allocate failed derivatives positions to surviving members, for example by auction.
 - d. Assign the equity and any debt claims of the recapitalised or bridge CCP.
 - e. Resume clearing new trades.
 - f. Make appropriate changes to the CCP's rules, clearing agreements, and risk management procedures.
 - g. Permit clearing member resignations after a "cooling-off" period.
8. In a liquidation and wind-down:
- a. Tear up remaining positions, or novate them to other CCPs.
 - b. Evaluate claims against the assets of the CCP held by clearing members and other creditors.
 - c. Liquidate the CCP's remaining assets.
 - d. Assign the liquidated assets of the CCP to claimants.

It is not clear whether a CCP should continue clearing new derivatives trades while undergoing resolution. This should presumably be determined by the circumstances at the time, with the objective of mitigating the total distress costs of clearing members and other market participants. An inability to clear new trades could present some difficulty to market participants who have come to rely on straight through processing of trades (including clearing), and are attempting to quickly add or replace hedges. Moreover, US regulations may require that a "designated clearing organization" (DCO) continues to provide clearing services, and require (subject to exemptions) that market participants continue centrally clearing designated "standardised" derivatives. Whether waivers of these regulations can be obtained, and under what circumstances, is not clear. If a CCP is unable to clear during its reorganisation, then regulatory clearing requirements should, if possible, be temporarily waived for those types of derivatives for which there are no alternative CCPs.

In closing, I would say that much work remains to be done by failure administrative authorities. So far, most of the progress has been in the form of international standards and coordination of planning activities among international bodies, as described by the Financial Stability Board (2014, 2015). In particular, it was agreed in 2015 by the chairs of the FSB Standing Committee on Supervisory and Regulatory Cooperation, the FSB Resolution Steering Group, the Basel Committee on Banking Supervision, the CPMI, and IOSCO, that the FSB Resolution Steering Group "should serve as the focal point for work on the resolution of CCPs, working in close cooperation with CPMI and IOSCO". At the national level, particularly in some

jurisdictions, progress on the development of plans for the failure resolution of CCPs should be better disclosed for public consideration.

References

Bank of England (2014a): "The Bank of England's Sterling Monetary Framework", (also known as "The Red Book"), November, www.bankofengland.co.uk/markets/Documents/money/publications/redbook.pdf.

——— (2014b): "The Bank of England's Supervision of Financial Market Infrastructures", in *Bank of England Annual Report*, March 2014, www.bankofengland.co.uk/publications/Documents/fmi/fmiap1403.pdf.

CDS Default Management Working Group (2011): "Principles and best practices for managing a defaulting clearing member's remaining portfolio and shortfall in available funds", January.

The Clearinghouse (2012): "Central counterparties: recommendations to promote financial stability and resilience", December 2012. <http://www.theclearinghouse.org/index.html?f=074643>

Committee on Payment and Settlement Systems, Board of the International Organization of Securities Commissions (2012a): "Principles for financial market infrastructure", Bank for International Settlements, April 2012, www.bis.org/cpmi/publ/d101a.pdf.

——— (2012b): "Recovery and resolution of financial market infrastructures", Consultative Report, July 2012, www.bis.org/publ/cpss103.pdf, Comments at www.bis.org/publ/cpss103/comments.htm.

——— (2014): "Recovery of financial market infrastructures", Consultative Report, October, 2014.

Davies, G and M Dobler (2011): "Bank resolution and safeguarding the creditors left behind", Research and Analysis, *Bank of England Quarterly Bulletin*, Q3 2011, pp 213–23.

Dembo, A, J-D Deuschel, and D Duffie (2004): "Large portfolio losses", *Finance and Stochastics*, vol 8, pp 3–16.

Duffie, D (2016): "Resolution of failing central counterparties", in K Scott and J Taylor (eds), *Making failure feasible: how bankruptcy reform can end 'too big to fail'*, Hoover Institution Press.

Elliott, D (2013): "Central counterparty loss-allocation rules", Bank of England, *Financial Stability Paper*, no 20, April 2013, www.bankofengland.co.uk/publications/Documents/fsr/fs_paper20.pdf.

European Commission (2012): "Consultation on a possible recovery and resolution framework for financial institutions other than banks", Directorate General, Internal Markets and Services.

European Securities and Markets Authority (2012): "Draft technical standards under the Regulation (EU): No 648/2012 of the European Parliament and of the Council of 4 July 2012 on OTC Derivatives, CCPs and Trade Repositories, Final Report", September, ESMA/2012/600, www.esma.europa.eu/system/files/2012-600_0.pdf.

European Union (2012): "Regulation (EU) No 648/2012 of the European Parliament and of the Council of 4 July 2012 on OTC Derivatives, Central Counterparties and Trade Repositories", *Official Journal of the European Union*, 12 June 2012, <http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2012:201:0001:0059:EN:PDF>.

——— (2014): "Directive 2014/59/EU of The European Parliament and of the Council of 15 May 2014, Establishing a Framework for the Recovery and Resolution of Credit Institutions and Investment Firms", *Official Journal of the European Union*, 12 June 12 2014, http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2014.173.01.0190.01.ENG.

Financial Stability Board (2014): "Key attributes of effective resolution regimes for financial institutions", Basel, 15 October 2014, www.financialstabilityboard.org/2014/10/r_141015/?page_moved=1.

——— (2015): "2015 CCP Workplan", Financial Stability Board, 15 April 2015, www.fsb.org/wp-content/uploads/Joint-CCP-Workplan-for-2015-For-Publication.pdf

G20 (2009): "Pittsburgh Declaration", G20. https://g20.org/wp-content/uploads/2014/12/Pittsburgh_Declaration_0.pdf.

HM Treasury (2014): "Secondary legislation for non-bank resolution regimes", Consultation Outcome, 9 June 2014, www.gov.uk/government/consultations/secondary-legislation-for-non-bank-resolution-regimes/secondary-legislation-for-non-bank-resolution-regimes.

ISDA (2013): "CCP loss allocation at the end of the waterfall", August 2013, www2.isda.org/attachment/NTc5Nw==/CCP_loss_allocation_waterfall_0807.pdf.

JPMorgan Chase (2014): "What is the resolution plan for CCPs?", JPMorgan Chase Office of Regulatory Affairs, September.

LCH Clearnet (2015): "CCP risk management, recovery, and resolution", *LCH Clearnet White Paper*.

Singh, M (2014): "Limiting taxpayer puts – an example from CCPs", *IMF Working Paper*, October 2014.

Skeel, D (2014): "Financing systemically important financial institutions in bankruptcy", University of Pennsylvania Law School, this volume.

Steigerwald, R, and D DeCarlo (2014): "Resolving central counterparties after Dodd-Frank: Are CCPs eligible for 'orderly liquidation'", working paper, Federal Reserve Bank of Chicago, November 2014.

Written contribution by Benoît Cœuré¹

Has technology changed the nature of risks in financial markets?

Information technology has undoubtedly transformed financial markets over the last few decades. While both the scope and the pace of change vary greatly across different market segments, it is hard to think of any financial activity that still works in the way that it did when computers were large, slow and expensive.

There are many ways in which technology affects the way markets work. I will focus on four different areas, highlighting both the benefits and the concerns associated with each.

Automation

Powerful computers have transformed securities trading. Today's modern exchanges are fully electronic, and algorithmic and high-frequency trading have grown to dominate the marketplace. Computers have proven to be particularly useful for the automation of repetitive tasks such as submitting new quotes, comparing prices across different trading platforms and monitoring newswires. This naturally entails efficiency gains, as traditional activities (market-making, arbitrage, optimal order execution etc) are now carried out more quickly and cheaply than ever before.

However, this rapid transformation also entails some risks. Markets now function almost at the speed of light, which requires instantaneous reactions when things go wrong. Incidents such as the 2010 "Flash Crash" show that reliance on technology is not a one-way street, and new safeguards need to be adopted in order to keep up with the changing environment. Computers now perform tasks that are critical to the functioning of the market, and the effects of operational failures have to be kept in check in order to mitigate the resulting tail risks.

One example of such an operational failure was the so-called "Knightmare", during which the system used by New York Stock Exchange market-maker Knight Capital sent erroneous orders into the market for around 30 minutes.² As Sağlam (2015)³ explains, this trading glitch led to an increase in institutional trading costs for stocks where Knight Capital was acting as the designated market-maker. Such incidents speak to some of the concerns voiced by critics of high-frequency trading and ultimately pose a threat to the integrity of the markets in question. It is clear that

¹ Executive Board member of the European Central Bank and Chairman of the Committee on Payments and Market Infrastructures (CPMI).

² This incident resulted in a loss of more than USD 400 million for the firm (see "Knight Capital glitch loss hits \$461m", *Financial Times* (online edition), 17 October 2012). It was subsequently acquired by Getco, another electronic trading firm.

³ M Sağlam, "The rogue algorithm and its discontents: evidence from a major trading glitch", unpublished working paper, 2015.

the ongoing arms race in the financial industry presents a serious challenge for regulators, as they need to upgrade their human capital and the tools they have available in order to preserve an efficient and reliable market structure.

Decentralisation

Technology also enables the decentralisation of markets, as it obviates the need for economic agents to be physically located together in a single marketplace in order to reap the benefits of liquidity externalities. Today, IT systems enable market participants to virtually consolidate fragmented markets. This development, with the support of regulators on both sides of the Atlantic, has broken long-standing monopolies and fostered competition among financial market infrastructure providers (exchanges, central counterparties etc).⁴ This has undoubtedly benefited investors through increases in both diversity and innovation and reductions in fees. However, this new landscape also poses some challenges. Infrastructure is becoming more interconnected, which allows local shocks to spread throughout the system, and market fragmentation can lead to opacity and increased complexity, which can hamper efficient regulatory oversight. Moreover, the combination of reliance on technology and competitive pressure creates the potential for operational failures and resulting market breakdowns. It also increases vulnerability to cyber threats, which now feature prominently on the agendas of financial regulators. The industrial organisation of financial market infrastructures deserves close attention from both regulators and academics, especially because the provision of public goods such as price discovery can run counter to the private incentives of profit-maximising firms.⁵ This trade-off was at the heart of the discussion surrounding the updated rules for markets in financial instruments ("MiFID II") in Europe.

Shift of activities

The use of technology also allows certain activities to move away from the traditional banking sector. For example, hedge funds and other shadow banks frequently rely on quantitative investment strategies, which are greatly facilitated by the use of IT systems. More generally, technological innovations such as the internet reduce search costs, thereby lowering barriers to entry for innovative adaptations of long-standing business models across the entire spectrum of financial services, from consumer finance to securities trading. While this, again, entails significant benefits for investors and consumers through increased competition, it also implies certain risks. In particular, some activities may not be subject to efficient regulation and supervision – an issue that is well known to all of us, not least since the subprime mortgage crisis. The Financial Stability Board has been addressing these risks in its workstream on

⁴ Regulation National Market System in the United States and the Markets in Financial Instruments Directive (MiFID) in Europe paved the way for fierce competition among stock exchanges.

⁵ For example, the joint pricing of trading fees and market data by exchanges can distort price discovery (see G Cespa and T Foucault, "Sale of price information by exchanges: does it promote price discovery?", *Management Science*, vol 60, 2015, pp 148–65).

shadow banking. In addition, sufficient transparency is also important for investor and consumer protection, including as regards privacy issues.

New activities

Finally, technological progress inevitably leads to the emergence of new financial activities. A prime example of this is the emergence of virtual currency schemes such as Bitcoin. While these schemes have many of the features of a regular currency – being, for example, a unit of account and a means of payment – they are not legal tender and are not backed by a state or government. Moreover, their value tends to be rather volatile, meaning that they could, instead, be considered to be a commodity.

Most virtual currency schemes include a payment system, such as Bitcoin's distributed ledger, the blockchain. This is a beneficial feature, which can exert competitive pressure on existing payment systems (particularly those used for costly transactions, such as international money transfers to and from developing countries) and could foster the introduction of faster and cheaper payment systems. This feature is also helpful in terms of financial inclusion.

However, there are also several risks. First, Bitcoin is opaque; that is to say, there is no supervision or oversight, as there is with regular currencies. Moreover, there is no investor/depositor protection, so that users have no recourse to a legal framework or judicial enforcement (for example, in cases of hacking and other such crimes). Finally, Bitcoin is entirely based on trust. While this is, in principle, also the case for a fiat currency, Bitcoin is not backed by a sovereign state with the right to levy taxes, and it has no legal/political basis. This can be attractive for small communities, but it cannot support economic activity in larger jurisdictions.

Finally, the rise of digital finance poses the same challenges for the financial industry that digitalisation and disintermediation (in more mundane terms, "Uberisation") pose for other parts of the economy: will established players be capable of reaping the benefits of exploiting consumer and investor data, or will these benefits be captured by new, less well regulated actors operating at the retail end of the system? Do financial intermediaries need to grow even bigger to reap the benefits of scale, potentially running counter to regulators' desire to avoid any organisation being "too big to fail"? I very much look forward to financial and industry organisations providing answers to these questions in the years to come.

Summary

Technology has a fundamental impact on financial markets, as it changes the way in which transactions are handled and information is processed. Overall, technological progress leads to considerable efficiency gains and reduced costs for investors and consumers through automation, innovation and competition. However, it also results in new risks, and it is important that regulators keep up to speed in order to ensure that the benefits are reaped by society as a whole.

Previous volumes in this series

No	Title	Author
563 May 2016	Who supplies liquidity, how and when?	Bruno Biais, Fany Declerck and Sophie Moinas
562 May 2016	Expectations and investment	Nicola Gennaioli, Yueran Ma and Andrei Shleifer
561 May 2016	Mobile collateral versus immobile collateral	Gary Gorton and Tyler Muir
560 May 2016	Has the pricing of stocks become more global?	Ivan Petzev, Andreas Schrimpf and Alexander F Wagner
559 April 2016	A comparative analysis of developments in central bank balance sheet composition	Christiaan Pattipeilohy
558 April 2016	Why bank capital matters for monetary policy	Leonardo Gambacorta and Hyun Song Shin
557 April 2016	How Does Bank Capital Affect the Supply of Mortgages? Evidence from a Randomized Experiment	Valentina Michelangeli and Enrico Sette
556 April 2016	Threat of Entry and Debt Maturity: Evidence from Airlines	Gianpaolo Parise
555 March 2016	The causal effect of house prices on mortgage demand and mortgage supply: evidence from Switzerland	Christoph Basten and Catherine Koch
554 March 2016	Can a Bank Run Be Stopped? Government Guarantees and the Run on Continental Illinois	Mark Carlson and Jonathan Rose
553 March 2016	What drives the short-run costs of fiscal consolidation? Evidence from OECD countries	Ryan Banerjee and Fabrizio Zampolli
552 March 2016	Fiscal sustainability and the financial cycle	Claudio Borio, Marco Lombardi and Fabrizio Zampolli
551 March 2016	When the Walk is not Random: Commodity Prices and Exchange Rates	Emanuel Kohlscheen, Fernando H. Avalos and Andreas Schrimpf
550 March 2016	A new dimension to currency mismatches in the emerging markets: non-financial companies	Michael Chui, Emese Kuruc and Philip Turner
549 March 2016	Monetary policy spillovers and currency networks in cross-border bank lending	Stefan Avdjiev and Előd Takáts
548 March 2016	Moving in tandem: bank provisioning in emerging market economies	Andres Murcia and Emanuel Kohlscheen
547 March 2016	When pegging ties your hands Intertemporal considerations in currency crises	Nikola Tarashev and Anna Zabai

All volumes are available on our website www.bis.org.