JOIM

MOORE'S LAW VS. MURPHY'S LAW IN THE FINANCIAL SYSTEM: WHO'S WINNING?

Andrew W. Lo[a]

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 behavior so users can employ these tools safely, effectively, and effortlessly. Examples of such technology are provided.

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 (Moore, 1965). This simple observation, implying a constant rate of growth, led Moore to extrapolate an increase in computing potential from sixty transistors per chip in 1965 to sixty thousand in 1975. This number seemed absurd at the time, but it was realized on schedule a decade later. Later revised by Moore to a doubling every two years, "Moore's Law" has been a remarkably prescient forecast of the growth of the semiconductor industry over the last 40 years, as Figure 1 confirms.

Technological change is often accompanied by unintended consequences. The Industrial

[a]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, USA; alo-admin@mit.edu. Other affiliations, including all commercial relationships, are disclosed at http://alo.mit.edu/.

Figure 1   Illustration of Moore's Law via transistor counts on various semiconductor chips from 1971 to 2011, which seems to double every two years (*Source*: "*Transistor Count and Moore's Law—2011 by Wgsimon*", CC BY-SA 3.0).

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 increased a number of birth defects before we understood their properties. 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 2007–2009. 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 (IPOs), 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 iPhone, paper road maps versus

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 estimated to be 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 data 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—led directly to the modern system of paper banknotes and token coinage. The invention of the telegraph sparked

Figure 2   Semi-logarithmic plot of estimated world population from 10,000 B.C. to 2011 A.D. *Source*: U.S. Census Bureau (International Data Base) and author's calculation.

a continent-spanning communications revolution that led to the creation of the modern futures market in 19th-century Chicago. And improvements to the ticker tape machine—symbolic of Wall Street for over a century—made Thomas Edison his early fortune.

2.1   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, NAS-DAQ, 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.[1] 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).

At \$395 (\$1,767 in 2016 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 motorized slot in the calculator. Shortly after the SR-52 debuted, one of the founders of the CBOE, a savvy options trader named Irwin Guttag, purchased one for his teenage son and asked him to program the Black–Scholes/Merton formula for it.[2] Within a year, many CBOE floor traders were sporting SR-52's

Figure 3   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.
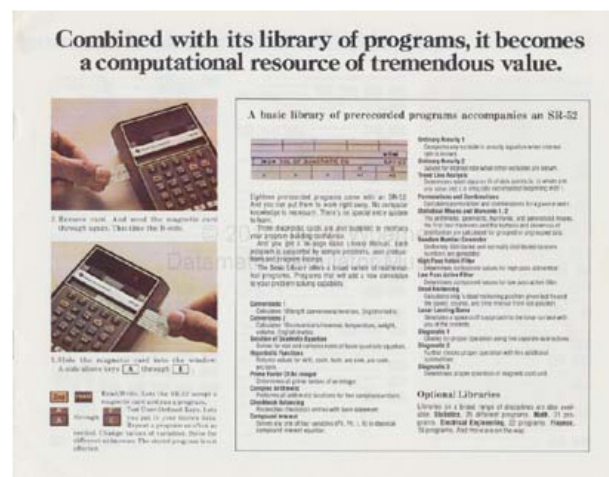
of their own. By 1977, Texas Instruments introduced a new programmable TI-59 with a "Securities Analysis Module" that would automatically calculate prices using the Black–Scholes/Merton formula. When Scholes confronted Texas Instruments about their unauthorized use of the formula, they replied that it was in the public domain. When he asked for a calculator instead, Texas Instruments replied that he should buy one.[3]

2.2 A financial Moore's Law

The Black, Scholes, and Merton publications launched well over a thousand subsequent articles (see Lim and Lo, 2006, for example), 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.[4] 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.

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 5 years. Even the Financial Crisis of 2007–2009 could only temporarily halt this growth for a couple of
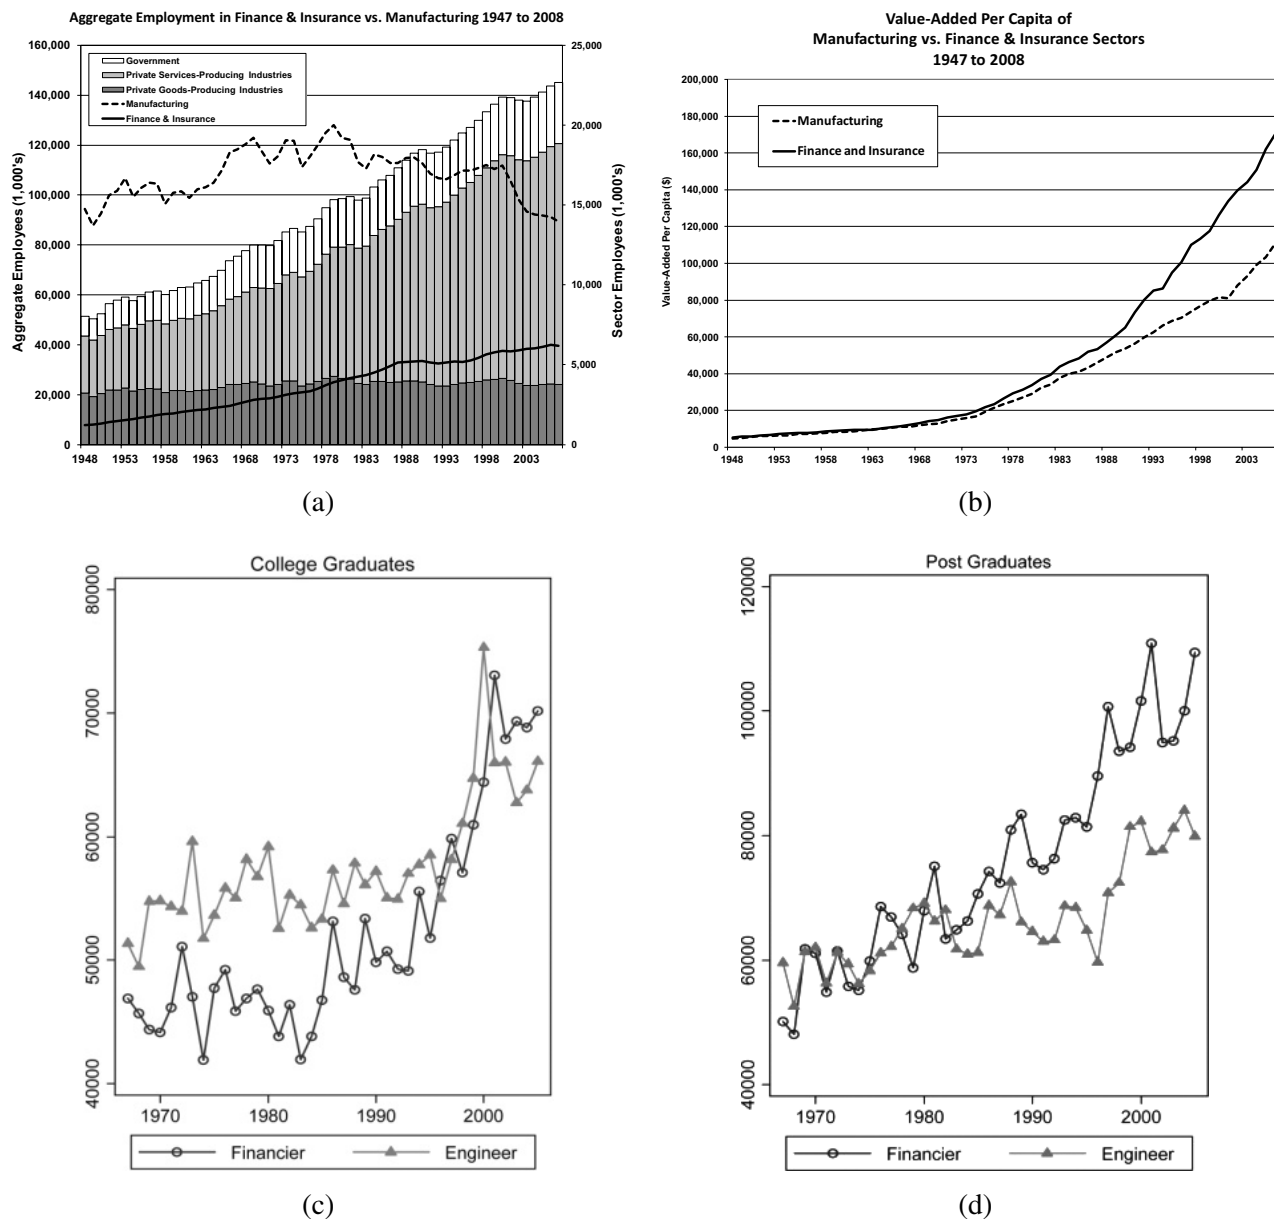

Figure 4 Financial Moore's Law: Raw and natural logarithm of average daily trading volume by year of exchange-listed options and futures from the Options Clearing Corporation, 1973–2014, and linear regression estimate of geometric growth rate, which implies a doubling every $\log(2)/0.14 = 4.95$ years.

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.[5] 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 individualized financial products.

The inexorable march of technological progress is part of a much broader trend of finance increasing its role in modern society. Figure 5 presents

Aggregate Employment in Finance & Insurance vs. Manufacturing 1947 to 2008

(a)

Value-Added Per Capita of
Manufacturing vs. Finance & Insurance Sectors
1947 to 2008

(b)

College Graduates

(c)

Post Graduates

(d)

Figure 5 Four illustrations of the growing importance of finance: (a) aggregate U.S. employment in manufacturing vs. finance & insurance sectors; (b) value-added per capita in manufacturing vs. finance & insurance; (c, d) annual income of college-graduate and post-graduate engineers and financiers (all wages are in 2000 U.S. dollars and are weighted using sampling weights), from Philippon and Reshef (2009, Figure 7).

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 labor 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 Figures 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 particularly 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.[6]

3.1 The Quant Meltdown of August 2007

Beginning on Monday, August 6, 2007, and continuing through Thursday, August 9, 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.[7]

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."

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 (2007, 2011) analyzed the events of the meltdown by simulating the returns of the contrarian trading strategy of Lehmann (1990), and Lo and MacKinlay (1990), on the historical data. 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 the 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.

3.2   The Flash Crash

At approximately 1:32 P.M. Central Time, May 6, 2010, U.S. 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 organization'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.[9] In other words, there is no single "culprit" that 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 reminiscent of the October 19, 1987 stock market crash, sell orders in the futures market triggered by an automated selling program cascaded into a systemic event for the entire U.S. 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 September 30, 2010 when the joint CFTC/SEC report was published. However, the narrative has changed. On April 21, 2015, the U.S. Department of Justice filed charges against Navinder Singh Sarao, a British national.[10] 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 (CME), 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 customize his software to submit orders to give the illusion of a deep market before they were canceled. To quote from the Department of Justice's (2015) 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."

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.

3.3   The BATS and Facebook IPOs

On March 23, 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 U.S. at the time and did it from the suburbs of Kansas City. As one of the 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.[12] No information about the glitch was made public during the day. Despite the quick deployment of a software patch by that afternoon, the confusion was so great that BATS suspended trading in its own stock, and ultimately canceled its IPO altogether.

An even bigger glitch occurred on May 18, 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 in 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 high volume of trades so 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.[13] 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 a.m., Facebook's IPO opened a half hour late because of these delays. As of 10:50 a.m., 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 acknowledgments, creating more uncertainty and anxiety.[14] 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 totaling 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.[15] 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 U.S. history.[16]

3.4   Knight Capital Group

At market open on August 1, 2012, the well-known U.S. broker-dealer Knight Capital Group—one of the largest equity traders in the industry at the time and among the 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 a.m. and 10:00 a.m.[17] 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.[18] 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 telecommunication systems and trading algorithms far ahead of most of their competition.

3.5   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 U.S. Treasury market, on October 15, 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 a.m. Eastern Time to selloffs precipitated by poor U.S. retail sales in September. Shortly after 9:33 a.m. 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.[19]

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 hypothesized that it was driven by "large scale repositioning by leveraged investors, activities of electronic trading algorithms, and dealer balance sheet and risk management constraints."[20] 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.[21]

3.6 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 man-made disaster. Despite this redundancy, however, when the Bloomberg terminal system was disrupted on April 17, 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."[22]

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 U.S., but affected markets throughout the Eastern Hemisphere, leading to the postponement of a multibillion buy-back of government debt by the U.K. Debt Management Office. Although alternative systems such as the Thomson Reuters network were available to many, these systems lacked the customized 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 rumor blamed the outage on a spilled can of soda on a critical Bloomberg server.[23]

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 emphasizing that contagion across markets and venues is still rare.[24] 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."[25] 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 of U.S. Treasuries, whose plunge in October 2014 is still 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 brittle to accident and deliberate bad intent. The linkages made possible by technological innovation may have increased systemic financial risk in unforeseen ways, but to lower 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, e.g., 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.

4.1   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 CME, one of the world's largest organized financial exchanges, has developed dynamic margin requirements so as to protect both the exchange and market participants from default due to extreme losses.[26] 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.[27] SPAN calculates the maximum market loss of a portfolio under multiple scenarios (typically sixteen; however, the number is user-defined) and then determines what the appropriate margin requirement should be.[28] Brennan and Lo (2014) have shown that the SPAN-calculated margin requirements for currency futures at the CME strongly correlate with recent volatility for U.S. 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 CME for protecting its clearinghouse 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 clearinghouses 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 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 CME, 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.

4.2 Law is code

Therefore, to regulate the financial system as a whole, we need to better understand financial regulation as a whole. The U.S. 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 organized.

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 analyze 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.[29] Using metrics for measuring the quality of software (see Table 1), Li *et al.* (2015) analyzed the entire text of the U.S. 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 U.S. 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 U.S.C. § 329, which involves an incentive bonus for retired or former members of the military. This section cites exactly two other sections, 37 U.S.C. § 303a(e) (general provisions of special pay in

Table 1 Principles of good software design and corresponding metric. Source: 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)

Figure 6   Network representation of references to and from a section of the U.S. legal code (37 U.S.C. § 329). *Source*: Li *et al.* (2015).

the military), and 10 U.S.C. § 101(a)(16) (a definition of "congressional defense committees"), while 37 U.S.C. § 329 is cited by one other section, 10 U.S.C. § 641, which notes that other laws in Title 10 of the U.S. legal code do not apply to the officers to whom the bonus in 37 U.S.C. § 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.

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.

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 4 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.

Li *et al.* (2015) apply this complexity measure to various parts of the U.S. 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
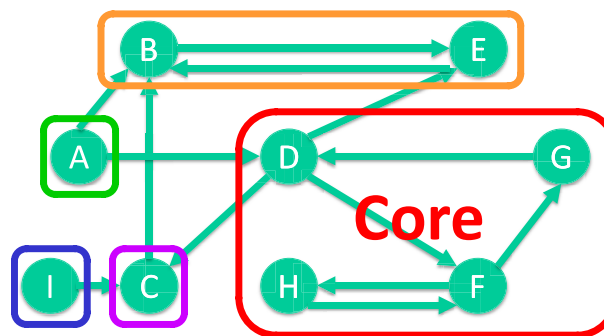
Figure 7   Illustration of the concept of strong connectedness in a directed graph and the "core," which is the largest strongly connected set of nodes. *Source*: Li *et al.* (2015).

| (a) | (b) | (c) |

Figure 8 Core–periphery network maps of: (a) sections of the U.S. legal code modified by the Omnibus Appropriations Act of 2009; (b) sections of the U.S. legal code modified by the Dodd–Frank Wall Street Reform Act; and (c) Title 12 of the U.S. legal code (Banks and Banking). Blue dots indicate peripheral sections, red dots indicate the core. *Source*: Li *et al.* (2015).

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 programs.

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 July 21, 2010. Of the 390 rulemaking requirements imposed by Dodd–Frank, only 267 have been satisfied by finalized rules as of December 31, 2015.[30] However, the complexity of Dodd–Frank does not compare with that of Title 12 of the U.S. 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.

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 organized. These sections of the law are of critical importance to the U.S. 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 behavior, rather than adding increasingly complicated patches to the law whose systemic effects are unknown.

4.3 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 practices 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 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 $Y_n = S_1 + S_2 + \cdots + S_n + X_1 + X_2 + \cdots + X_n$.

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 $S_1 + S_2 + \cdots + S_n$, which, when divided by the number n which is observable, yields the average salary in the room. Figure 9 summarizes 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.

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

Figure 9 Illustration of a simple secure multi-party computation algorithm for computing the average salary of a group of a individuals without requiring any individual to reveal his or her salary.

confidential. In Abbe *et al.* (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.[31] 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

(a)   (b)
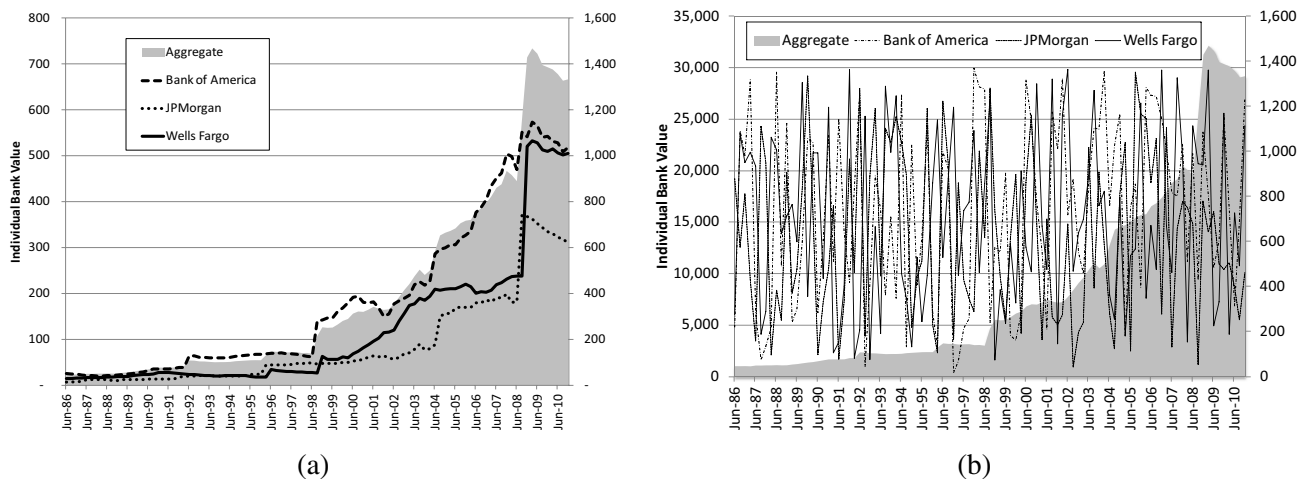
Figure 10 Example of secure multi-party computation of the aggregate size of the real-estate loan portfolios of Bank of America, JP Morgan, and Wells Fargo. Graph (a) contains the raw time series for the three individual banks as well as the aggregate sum; graph (b) contains three encrypted time series which, when summed, yields the same aggregate sum as the unencrypted data. Source: Abbe *et al.* (2012).

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 worse way than what the regulation originally intended to ban. 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 analyzed scrupulously, and no data must be withheld from investigators.

One potential model for this scrupulous form of analysis already exists.[32] The National Transportation Safety Board (NTSB) has an excellent track record in analyzing and determining the causes of transportation accidents in the U.S. The NTSB has no regulatory authority, freeing the agency to criticize 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, an NTSB report is able to 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.[33] 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."[34] Like the financial system, the human body is also an immensely complicated and hyper-connected 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 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, analyze 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.

Acknowledgment

Notes

[1] Black and Scholes (1973) and Merton (1973).
[2] Private communication with John V. Guttag—Irwin Guttag'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.
[3] Scholes (2006).
[4] 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).
[5] Kirilenko and Lo (2013).
[6] Some of these examples are drawn from Kirilenko and Lo (2013) with permission.
[7] Zuckerman *et al.* (2007) and Sender *et al.* (2007).
[8] Khandani and Lo (2007, 2011), Lehmann (1990) and Lo and MacKinlay (1990).
[9] CFTC/SEC (2010) and Kirilenko *et al.* (2014).
[10] Brush *et al.* (2015) and CFTC (2015).
[11] Department of Justice (2015).
[12] Dornbrook (2012), Oran *et al.* (2012), and Schapiro (2012).
[13] Benoit (2012) and Schapiro (2012).
[14] Strasburg *et al.* (2012).
[15] Strasburg and Bunge (2012).

[16]   Strasburg and Bunge (2013).
[17]   McCrank (2012), Telegraph (2012), and Schapiro (2012).
[18]   SEC (2013).
[19]   Alloway and MacKenzie (2014) and Potter (2015).
[20]   TMPG (2014).
[21]   Potter (2015).
[22]   Cox (2015) and Cox and Trivedi (2015).
[23]   Brinded (2015).
[24]   Schapiro (2012).
[25]   Raymond (1999, p. 30).
[26]   Brennan and Lo (2014).
[27]   CME Group (2015).
[28]   CME Group (2010).
[29]   Li *et al.* (2015).
[30]   Davis Polk (2015, p. 2).
[31]   Abbe *et al.* (2012).
[32]   Fielding *et al.* (2011).
[33]   NTSB (1993).
[34]   FDA (2011).

References

Abbe, Emmanuel A., Khandani, Amir E., and Lo, Andrew W. (2012). "Privacy-Preserving Methods for Sharing Financial Risk Exposures," *American Economic Review*: *Papers and Proceedings* **102**, 65–70.

Alloway, Tracy and MacKenzie, Michael (2014). "Bonds: Anatomy of a Market Meltdown," *Financial Times*. November 17.

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

Black, Fischer and Scholes, Myron (1973). "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* **81**, 637–654.

Brinded, Lianna (2015). "RUMOUR: The Bloomberg Outage was Caused by a Spilled Can of Coke," Business Insider. April 17. http://www.businessinsider.com/reason-for-the-global-bloomberg-outage-2015-4.

Brennan, Thomas J. and Lo, Andrew W. (2014). "Dynamic Loss Probabilities and Implications for Financial Regulation," *Yale Journal on Regulation* **31**, 667–694.

Brush, Silla, Schoenburg Tom, and Ring Suzi (2015). "How a Mystery Trader with an Algorithm may have Caused the Flash Crash," *Bloomberg*. April 21. http://www.bloomberg.com/news/articles/2015-04-22/mystery-trader-armed-with-algorithms-rewrites-flash-crash-story.

CFTC (Commodity Futures Trading Commission) (2015). *U.S. Commodity Futures Trading Commission v. Nav Sarao Futures Limited PLC and Navinder Singh Sarao.* http://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. May 18. http://www.sec.gov/sec-cftc-prelimreport.pdf

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

CME Group (2015). "Risk Management Overview," Available at: http://www.cmegroup.com/clearing/risk-management.

Cox, Josie (2015). "Deals Postponed After Bloomberg Terminals Go Down Globally," *Wall Street Journal*. April 17. http://blogs.wsj.com/moneybeat/2015/04/17/deals-postponed-after-bloomberg-terminals-go-down-in-europe-and-asia/.

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

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

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

Dornbrook, James (2012). "BATS CEO: We Solved the Problem, Weren't Dependent on the IPO," *Kansas City Business Journal. March 29*. http://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*. http://www.fda.gov/downloads/ScienceResearch/SpecialTopics/RegulatoryScience/UCM268225.pdf

Fielding, Eric, Lo, Andrew W., and Yang Jian Helen (2011). "The National Transportation Safety Board: A Model for Systemic Risk Management," *Journal of Investment Management* **9**, 17–49.

Khandani, Amir E. and Lo, Andrew W. (2007). "What Happened to the Quants in August 2007?" *Journal of Investment Management* **5**, 29–78.

Khandani, Amir E. and Lo, Andrew W. (2011). "What Happened to the Quants in August 2007?: Evidence from Factors and Transactions Data," *Journal of Financial Markets* **14**, 1–46.

Kirilenko, Andrei A. and Lo, Andrew W. (2013). "Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents," *Journal of Economic Perspectives* **27**, 51–72.

Kirilenko, Andrei, Kyle, Albert S., Samadi, Mehrdad, and Tuzun Tugkan (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, Bruce N. (1990). "Fads, Martingales, and Market Efficiency," *Quarterly Journal of Economics* **105**, 1–28.

Li, William, Azar Pablo, Larochelle David, Hill Phil, and Lo, Andrew W. (2015). "Law is Code: A Software Engineering Approach to Analyzing the United States Code," 10 J. Bus. & Tech. L. 297 (2015).

Lim, Terence and Lo, Andrew W. (2006). *The Derivatives Sourcebook*. Boston, MA: NOW Publishers.

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

Lo, Andrew, and MacKinlay, Craig (1990). "When Are Contrarian Profits Due to Stock Market Overreaction?" *Review of Financial Studies* **3**, 175–205.

Merton, Robert (1973). "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science* **4**, 141–183.

McCrank, John (2012). "Knight Glitch Likely to Lead to Regulatory Changes: CEO," *Reuters*. September 11. http://www.reuters.com/article/2012/09/11/us-knight-ceo-idUSBRE88A1GW20120911.

Moore, Gordon E. (1965). "Cramming More Components onto Integrated Circuits," *Electronics* **38**, 114–117.

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, Olivia, Spicer, Jonathan, Mikolajczak, Chuck, and Mollenkamp, Carrick (2012). "BATS Exchange Withdraws IPO after Stumbles," *Reuters*. March 24. http://www.reuters.com/article/2012/03/24/us-bats-trading-idUSBRE82M0W020120324.

Philippon, Thomas and Reshef, Ariell (2009). "Wages and Human Capital in the U.S. Financial Industry: 1909–2006," NBER Working Paper No. 14655, http://www.nber.org/papers/w14644.

Potter, Simon (2015). "Challenges Posed by the Evolution of the Treasury Market," Remarks at the 2015 Primary Dealer Meeting, New York City, April 13. http://www.newyorkfed.org/newsevents/speeches/2015/pot150413.html.

Raymond, Eric S. (1999). *The Cathedral & the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary* Cambridge, Mass: O'Reilly.

Schapiro, Mary (2012). "Introductory Remarks at SEC's Market Technology Roundtable." October, 2. http://www.sec.gov/news/speech/2012/spch100212mls.htm.

Scholes, Myron S. (2006). "Derivatives in a Dynamic Environment," In Lim, Terence, Lo, Andrew W., Robert C. Merton, Myron S. Scholes, and Martin B. Haugh. 2006. *The Derivatives Sourcebook*. Boston: 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*. October 16. https://www.sec.gov/litigation/admin/2013/34-70694.pdf.

Sender, Henny, Kelly, Kate, and Zuckerman, Gregory (2007). "Goldman Wagers on Cash Infusion to Show Resolve," *Wall Street Journal.* August 14, p. A.1.

Strasburg, Jenny, Ackerman, Andrew, and Lucchetti, Aaron (2012). "Nasdaq CEO Lost Touch Amid Facebook Chaos," *Wall Street Journal.* June 11, p. A.1.

Strasburg, Jenny and Bunge, Jacob (2012). "Social Network's Debut on Nasdaq Disrupted by Technical Glitches, Trader Confusion," *Wall Street Journal.* May 19, p. A.2.

Strasburg, Jenny and Bunge, Jacob (2013). "Nasdaq is Still on Hook as SEC Backs Payout for Facebook IPO," *Wall Street Journal.* March 25.

Telegraph (2012). "Knight Capital's $440m Trading Loss" Caused by Disused Software," August 14. http://www.telegraph.co.uk/finance/newsbysector/banksandfi-

nance/9475292/Knight-Capitals-440m-trading-loss-cau-sed-by-disused-software.html.

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

Zuckerman, Gregory, Hagerty, James, and Gauthier-Villars, David (2007). "Impact of Mortgage Crisis Spreads, Dow Tumbles 2.8% as Fallout Intensifies; Moves by Central Banks," *Wall Street Journal.* August 10, p. A.1.