

The Global Financial Markets: An Ultra-Large-Scale Systems Perspective

Dave Cliff¹ and Linda Northrop²

¹ UK Large Scale Complex IT System Initiative,
Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.
dc@cs.bris.ac.uk

² Research, Technology, and System Solutions Program,
Software Engineering Institute, Carnegie-Mellon University, Pittsburgh PA 15213, USA
lmn@sei.cmu.edu

Abstract. We argue here that, in recent years, the world's financial markets have become a globally interconnected complex adaptive ultra-large-scale socio-technical system-of-systems, and that this has important consequences for how the financial markets should be engineered and managed in future. Major failures in the financial markets can now occur at super-human speeds, as was witnessed in the “Flash Crash” of May 6th 2010. Events such as the Flash Crash may become more commonplace in future, unless lessons are learned from other fields where complex adaptive socio-technical systems-of-systems have to be engineered for high-integrity, safety-critical applications. In this document we review the literature on failures in risky technology and high-integrity approaches to safety-critical SoS engineering. We conclude with an argument that, in the specific case of the global financial markets, there is an urgent need to develop major national strategic modeling and predictive simulation capabilities, comparable to national-scale meteorological monitoring and modeling capabilities. The intent here is not to predict the price-movements of particular financial instruments or asset classes, but rather to provide test-rigs for principled evaluation of systemic risk, estimating probability density functions over spaces of possible outcomes, and thereby identifying potentially high-consequence failure modes in the simulations, before they occur in real life, by which time it is typically too late.

Keywords: Large-scale complex IT systems, ultra-large-scale systems, financial markets, algorithmic trading, high-frequency trading, normalization of deviance, flash crash.

1 Introduction

For what events will the date of May 6th, 2010 be remembered? In Britain, there was a general election that day, which ousted the ruling Labour Party after 13 years and led to the formation of the UK's first coalition government since 1945. Nevertheless, it seems likely that in financial circles at least, May 6th will instead long be remembered for dramatic and unprecedented events that took place on the other side of the Atlantic, in the US capital markets. May 6th is the date of what is now widely known as the “Flash Crash”.

On that day, in a period lasting roughly 30 minutes from approximately 2:30pm to 3:00pm EST, the US equity markets underwent an extraordinary upheaval: a sudden

catastrophic collapse followed by an equally unprecedented meteoric rise. In the space of only a few minutes, the Dow Jones Industrial Average dropped by over 600 points, representing the disappearance of more than 850 billion dollars of market value. In the course of this sudden downturn, the share-prices of several blue-chip multinational companies went haywire, with shares in companies that had previously been trading at a few tens of dollars plummeting to \$0.01 in some instances, and rocketing to values of \$100,000 in others. Seeing prices quoted by some major exchanges suddenly going crazy, other major exchange-operators declared “self-help” (that is, they invoked a regulation allowing them to no longer treat the price-feeds from the other exchanges as valid), thereby decoupling the trading on multiple venues that had previously been unified by the real-time exchange of reference price data.

Then as suddenly as this downturn occurred it reversed, and over the course of another few minutes most of the 600-point loss in the Dow was recovered, and share prices returned to levels within a few percentage points of the values they had held before the crash. That recovery took less than twenty minutes.

Two weeks after the Flash Crash, the US Securities and Exchange Commission (SEC) and the US Commodity Futures Trading Commission (CFTC) jointly released an interim report into the events of May 6th (CFTC&SEC, 2010a) that established very little, other than dispelling rumours of the flash crash having been caused by a “fat-finger” error (where a trader mis-keys an order) or terrorist action. After that, for more than four months there was open speculation on the cause of the Flash Crash, and senior figures in the markets voiced their growing exasperation at the lack of a straightforward explanation. Identifying the cause of the crash was made difficult by the “fragmentation of liquidity” (trading taking place simultaneously on a number of independent but interconnected exchange-venues); by the lack of a single unifying “consolidated audit trail” showing synchronized timestamps for all events in all the markets with identifiers of the originators of those events; and by the widespread use of algorithmic trading systems: autonomous adaptive software systems that automate trading jobs previously performed by human traders, many operating at super-human speeds. Various theories were discussed in the five months that it took the SEC and CFTC to produce their joint final report on the events of May 6th. Many speculated on the role of high-frequency trading (HFT) by investment banks and hedge funds, where algorithmic traders buy and sell blocks of financial instruments on very short timescales, sometimes holding a position for only a few seconds or less.

When the SEC/CFTC final report on the Flash Crash was eventually published on September 30th, nearly five months after the event, (CFTC&SEC, 2010b), it made no mention of a “bug” anywhere in the system being a causal factor. Instead, the story it told was that the trigger-event for the crash was a single block-sale of \$4.1bn worth of futures contracts, executed with uncommon urgency on behalf of a traditional fund-management company. It was argued that the consequences of that trigger event interacting with HFT systems rippled out to cause the system-level failures just described. The SEC/CFTC report was met with very mixed responses. Many readers concluded that it left more questions unanswered than resolved, and a subsequent much more detailed analysis of the time-series “tapes” of market event data conducted by Nanex Corp.¹ offered an alternative story that many market

¹ See www.nanex.net

practitioners found more plausible: see Meerman *et al.* (2010) and Easley *et al.* (2011) for further details of the extent to which the CFTC/SEC version of events is disputed.

Ten months after the event, in February 2011, a specially convened panel of regulators and economists, the Joint CFTC-SEC Advisory Committee on Emerging Regulatory Issues, released a report (CFTC&SEC, 2011) urging a number of rule changes, some of them fundamental and potentially far-reaching. At the time of writing this Foresight review, the extent to which the report's recommendations will be acted upon is unclear (see, e.g., Demos, 2011a, 2011b, 2011c).

Now the fact that there was such a rapid recovery immediately after the down-spike meant that, by the close of business on May 6th the overall inter-day price change on the previous day was nothing particularly dramatic. To someone focused only on daily market-close prices, this may look like just another day of a downward-trending market in a time of significant political and economic uncertainty: on that day, the Greek national debt crisis was threatening to destabilize the entire Euro-zone single-currency economic union; and the indeterminate outcome of the UK general election was a further distraction. For sure, the intra-day dynamics on May 6th were unlike anything ever seen before, but the market pulled back, so what is there to worry about?

We contend that there are two significant reasons to be worried by the Flash Crash. The first worry is that at the micro-level there was a clear market failure: whether a trader was richer or poorer by the end of that disorderly day was in many cases not much more than a lottery. The second worry is the macro-level observation that, with only a very slight change in the sequence of events, the global financial markets could plausibly have gone into meltdown, with May 7th 2010 (i.e. the *next* day) becoming the date of a global collapse that dwarfed any previous stock-market crash. We'll expand on these two worries in the next two paragraphs.

The first worry, on the micro-level, is that while some equity spot and derivatives trades that took place at the height of the mayhem were subsequently "busted" (declared to be invalid on the basis that they were clearly made on the basis of erroneous data) by the exchanges, the means by which trades were selected for busting was argued by many to be arbitrary, after-the-fact rule-making. Some traders who had lost large amounts of money did not have their trades busted; some who had made handsome profits found their gains taken away. The flash-crash chaos had rippled beyond the equity markets into the foreign exchange (FX) markets where certain currency exchange rates swung wildly on the afternoon of May 6th as the markets attempted to hedge the huge volatility and risk that they were suddenly seeing explode in equities. There is no provision to retrospectively bust trades in FX, and so those deals were left to stand. Sizeable fortunes were made, and sizeable fortunes were lost, by those caught in the storm; the issue of who lost and who gained was in too many cases almost random.

The second worry is a much more significant concern: the Flash Crash could have occurred any time that day. Certainly the specific time-period during which the Flash Crash occurred, roughly 2:30pm to 3:00pm, was not cited as a causal factor in the official CFTC/SEC report on the events of May 6th, nor in the much more detailed analysis performed by Nanex Corp. This is a point recently explored in public statements by Bart Chilton, head of the CFTC, who said the following in a public lecture given in March 2011: "...Think about it. There are stocks and futures, which are arbitrated internationally. If the Flash Crash had taken place in the morning on

May 6th, when E.U. markets were open, it could have instigated a global economic event. Since it took place in the mid-afternoon, it was primarily limited to U.S. markets...” (Chilton, 2011). Although we respect Commissioner Chilton’s view, we think that in fact the much, much bigger worry is not what would have happened if the Flash Crash had occurred in the morning of May 6th, but instead what would have happened if it had occurred a couple of hours or so *later* that day. Specifically, we think that the true nightmare scenario would have been if the crash’s 600-point down-spike, the trillion-dollar write-off, had occurred immediately before market close: that is, if the markets had closed just after the steep drop, before the equally fast recovery had a chance to start. Faced with New York showing its biggest ever one-day drop in the final 15 minutes before close of business on May 6th, and in the absence of any plausible public-domain reason for that happening, combined with the growing nervousness that the Greek government would default on its sovereign debt and throw the entire Euro-zone economic union into chaos, traders in Tokyo would have had only one rational reaction: sell. The likelihood is that Tokyo would have seen one of its biggest ever one-day losses. Following this, as the mainland European bourses and the London markets opened on the morning of May 7th, seeing the unprecedented sell-offs that had afflicted first New York and then Tokyo, European markets would have followed into precipitous freefall. None of this would have been particularly useful in strengthening confidence in the Greek debt crisis or the future of the Euro, either. And, as far as we can tell, the only reason that this sequence of events was not triggered was down to mere lucky timing. Put simply, on the afternoon of May 6th 2010, the world’s financial system dodged a bullet.

Although the Flash Crash was a particularly extreme event, similar phenomena have been witnessed in various markets in the period since May 2010. Some notable examples are listed here, but this is by no means an exhaustive list:

- On 28th September 2010, share-prices of major technology stocks Apple, Dell, Hewlett-Packard, IBM, Microsoft, and Oracle all experienced sudden severe spike-transitions before returning to normal price ranges.²
- On May 2nd, 2011, the market price of gold spiked sharply downwards by \$20 and then immediately recovered more than \$15 of that loss. The graph of price against time for this event is strongly reminiscent of the graph of the Dow Jones Industrial Index during the Flash Crash. Unlike the Flash Crash, which unfolded over a period of roughly 30 minutes, this down-spike and recovery in the price of gold took less than 10 minutes.³
- The next day, on May 3rd, 2011, the price of silver dropped dramatically in after-hours trading, an event that was again attributed to algorithmic trading systems.⁴
- On June 8th 2011 the price of natural gas in the USA commodity markets had been trending flat (i.e. showing neither a rise or a fall for the day) over a period of

² <http://ftalphaville.ft.com/blog/2010/09/28/355081/market-on-edge-after-apple-drops-like-a-stone/>

³ See <http://www.zerohedge.com/article/golden-flash-crash>

⁴ See <http://www.zerohedge.com/article/and-now-todays-mini-silver-flash-crash-same-time-same-place>

several hours, when the price suddenly started to oscillate in a pattern strongly reminiscent of a smooth sine-wave, with the amplitude of the oscillations (the height of the peaks and the depths of the troughs) growing steadily in a short space of time, and then the price crashed dramatically. This event was also attributed to an erroneously programmed algorithmic trading system.⁵

- On July 7th 2011, there were sizeable swings in the price of crude oil futures on the New York Mercantile Exchange (one of the world's primary exchanges for trading of commodities and commodity derivatives). According to the analysis published by Nanex Corp, these swings appear to have been the result of a "massive arbitrage algorithm" running at a significant speed advantage for a period of around five seconds.⁶

- On March 23rd 2012, the initial public offering (IPO) of shares in the company *BATS Global Markets* was marred by an astonishingly fast collapse in its share price. BATS (an acronym for "Better Alternative Trading Systems") was founded in 2005 and is the owner and operator of popular electronic trading venues in major economies such as the USA and UK, that are alternatives to the traditional, longer-established stock exchanges. At 11:14 am on the day of the IPO, BATS' shares commenced trading at a price of \$15.25. Within 0.9 seconds of the start of trading, the price had dropped to \$0.2848; and at +1.5 seconds, the price was \$0.0002. A total of 567 orders had been executed by the time trading in BATS was halted, and BATS subsequently cancelled its IPO, thereby depriving the company of the capital that would otherwise have been raised by the sale of its shares that day. Subsequent reports in the financial media quoted BATS as blaming this collapse on "problems" and "bugs" in their own exchange software,⁷ but analysis of market data subsequently released by Nanex Corp⁸ indicated that the collapse in BATS' share price had been driven primarily by so-called *intermarket sweep orders* originating from a trader or traders operating on the Nasdaq exchange. Nanex's analysis also demonstrated that the collapse in BATS' price traced an almost-perfect logarithmic decay curve, strong evidence that a computerized system was driving the price down by a fixed percentage with each successive trade. In the days that followed, there was open speculation that BATS' crash had been caused by a rogue trading system, and that perhaps someone had deliberately programmed a system to inflict this grief on BATS.

In each of these cases, there is evidence to suggest that computer-based trading systems were involved in the transactions that played causal roles in these events, but unlike the events of May 6th 2010, there have been no official investigations launched by regulatory bodies such as the CFTC and the SEC. There is frequent open discussion among market practitioners (especially on the anonymously-sourced but very well-informed website www.zerohedge.com) that current markets are too often

⁵ See <http://www.zerohedge.com/article/story-berserk-nat-gas-algo-just-got-really-strange> and for the supporting data analysis see <http://www.nanex.net/StrangeDays/06082011.html>

⁶ See <http://www.nanex.net/StrangeDays/07072011.html>

⁷ See T. Demos (2012), "IPO Software Behind BATS' Failure", *The Financial Times*, March 26th 2012.

⁸ See <http://www.nanex.net/aqck/2970.html>

showing price movements for which the only plausible explanation is that a computerised system is operating in an unexpected or unanticipated fashion, giving market dynamics that deviate from expectations based on experience of markets populated by human traders, or from rational economic argument.

We argue here that market disorder such as the May 6th 2010 Flash Crash, and the mini-flash-crashes that have been recorded since then in various markets, are *deviant events* that are best understood as evidence of “*normal failure*” in an *ultra-large-scale complex adaptive socio-technical system-of-systems*.

Unpacking that assertion requires some care, so in the following sections we’ll start first with a discussion of notable technology failures, then bring the conversation back to discussion of failures of the financial markets.

Systems, such as the financial markets, that are themselves composed of constituent stand-alone systems that are each operationally and managerially independent, are very often the result of incremental, sporadic, organic growth and unplanned accretion rather than clean-sheet engineering design. They thereby involve or acquire significant degrees of variability in components and heterogeneity of constituent systems, and their make-up changes dynamically over multiple timescales. For this reason traditional engineering techniques, which are predicated on very different assumptions, cannot necessarily be trusted to deliver acceptable solutions. And, therefore, new approaches are required: new engineering tools and techniques, new management perspectives and practices.

In the main text and the appendices of this review, we survey some recently developed approaches that look likely to grow into promising new engineering techniques in the coming decade and beyond, better suited to current and future systems than our traditional engineering practices, which owe more to the mid-twentieth-century than they can offer the early-twenty-first.

2 Background: Failures in Risky Technology

The global financial markets are not the only area in which the application of new technologies has led to failures. Although operator error can be attributed to many failures, as technological systems grow in complexity the prospect of failure-modes being inadvertently designed-in also grows. Take, for example, bridge building. As an engineering activity this is something that dates at least as far back as ancient Rome (c.150BC) and so probably doesn’t figure as a risky technology for many people.

Yet for decades, engineering students have been taught the story of the Tacoma Narrows suspension bridge, opened in July 1940, which collapsed four months later, where the designers did not anticipate the prospect of wind-flows over the bridge deck reinforcing the deck’s natural mode of vibrations, leading to the bridge shaking itself apart. Presumably, current and future students will also be taught the story of the London Millennium Bridge, which opened in June 2000 and two days later was closed for two years to remedy destabilizing swaying motions induced when groups of people walked over it. A significant difference between Tacoma Narrows and London Millennium is that in the latter case, it was the interaction of people, the users, with the engineered system that caused the problem. The Millennium Bridge on its own, as a piece of engineering, was a fine and stable structure; but when we

consider the interaction dynamics of the larger system made up of the bridge *and* its many simultaneous users, there were serious unforeseen problems in those dynamics that only came to light when it was too late.

As engineered systems become more complex, it becomes more reasonable to argue that no one person or group of users is responsible for failures, but rather that the failures are inherent, latent, in the system; this seems especially so in the case of *socio-technical systems*, i.e. systems (like the Millennium Bridge, when in use) whose dynamics and behaviour can only be properly understood by including human agents (such as operators and/or users) within the system boundary.⁹

This is perhaps most clear in some of the more famous technology failures of the past 40 years. The oxygen-tank explosion that crippled the *Apollo 13* Lunar Service Module as it was en route to the moon in 1970, and subsequent safe return of her crew, has been rightly popularized as a major triumph of bravery, skill, teamwork, and engineering ingenuity. Nevertheless, the fact remains that NASA very nearly suffered the loss of *Apollo 13* and her crew, due to the compounding effect of several independent small failures of process rather than malign intent or major error from one or more individuals. The successful return of *Apollo 13*'s crew owed an awful lot to the availability of accurate simulation models, physical replicas on the ground of key components of the spacecraft, where emergency procedures could be invented and rehearsed and refined before being relayed to the astronauts. The value of simulation models is something that we will return to in depth, later in this paper.

While loss of a space vehicle is undoubtedly a tragedy for those concerned, the number of fatalities is small in comparison to the potential losses in other high-consequence systems, such as petrochemical plants and nuclear power stations. The release of toxic gas at the Union Carbide plant in Bhopal in December 1984 immediately killed over 2,000 people, with estimates of the subsequent delayed fatalities running at 6,000-8,000. The partial meltdown at the Three Mile Island nuclear plant in 1979 was successfully contained, but the reactor-core fire at Chernobyl in 1986 was not, and estimates of the number of deaths resulting from that event range from many hundreds to several thousand.

High-risk technology failures including *Apollo 13* and Three Mile Island were the subject of serious scholarly analysis in Charles Perrow's seminal work *Normal Accidents* (Perrow, 1984). Perrow argued that in tightly-coupled systems with sufficiently complex internal interactions, accidents and failures, including catastrophic disasters of high-risk systems with the potential to end or threaten many lives, are essentially inevitable – such accidents are, in that sense, to be expected as “normal”, regardless of whether they are common or rare.

In Perrow's terms, the losses of the NASA space shuttles *Challenger* in January 1986 and *Columbia* in February 2003 were also normal accidents. However, the sociologist Diane Vaughan argued for a more sophisticated analysis in her classic study *The Challenger Launch Decision* (1997), in which she presented a detailed analysis of transcripts, covering the hours immediately preceding *Challenger*'s launch, of interactions between NASA staff and the staff of Morton Thiokol, manufacturers of the shuttle's solid-fuel rocket booster (SRB) that failed leading to loss of the vehicle and her crew.

⁹ For an early, but very insightful, discussion of the dynamics of socio-technical systems, see Bonen (1979).

The transcripts had been released as part of the official Presidential Commission on the Space Shuttle *Challenger* Accident, led by William Rogers. A shocking finding of the Rogers investigation was that the specific failure-mode (burn-through of rubber O-ring seals in a critical joint on the SRB) had been known since 1977 and the consequent potential for catastrophic loss of the vehicle had been discussed at length by NASA and Thiokol, but the shuttle had not been grounded. Vaughan concluded that while the *proximal* cause of disaster was the SRB O-ring failure, the *ultimate* cause was a social process that Vaughan named *normalization of deviance*.

Put simply, normalization of deviance occurs when the safe-operating envelope of a complex system is not completely known in advance, and where events that were *a priori* thought to be outside the envelope, but which do not then result in failures, are taken after the fact as evidence that the safe envelope should be extended to include those events. In this way, deviant events become normalized: the absence of a catastrophe thus far is taken as evidence that in future catastrophes are less likely than had previously been thought. The flaw in this line of reasoning is starkly revealed when a catastrophe then ensues. In Vaughan's analysis, the loss of *Challenger* was not a purely technical issue but rather was an organizational failure in the *socio-technical system* comprised of the (technical) shuttle hardware systems and the (social) human individuals, teams, and organizations that had to interact appropriately to ensure safe launch and return of the shuttle.

Vaughan's analysis of the *Challenger* accident came more than a decade after the official inquiry into that 1986 event. In contrast, because of her work on *Challenger*, following the loss of *Columbia* in 2003 Vaughan was immediately invited onto the Columbia Accident Investigation Board (CAIB) and subsequently authored a chapter of the CAIB official report. It was argued that once again an organizational failure at NASA had resulted in loss of a vehicle, once again via a long-standing process of normalization of deviance.

For *Columbia*, the *proximal* cause of disaster was a lump of insulating foam that broke away from the external fuel tank and struck the leading edge of the orbiter's left wing, damaging its thermal insulation: on re-entry, this damage allowed atmospheric gases, compressed in the bow-wave at the wing edge and hence heated to more than 1,500 degrees Celsius, to penetrate the wing; and the vehicle then broke up at high speed. But the *ultimate* cause was an organizational culture that had once again engaged in normalization of deviance, despite the warnings from Vaughan's analysis of the *Challenger* disaster.

Prior to the loss of *Columbia*, sixty-four previous missions had suffered strikes from insulating material breaking away during launch and hitting the orbiter, and yet each such strike was technically a violation of the shuttle's design requirements: the shuttle had simply not been designed to withstand impacts from breakaway insulating material. Most notably, in 1988 on mission STS-27, insulation broke away from an SRB during launch and damaged 700 of the heat-insulating tiles on shuttle *Atlantis*, and the crew on board believed they would very likely be killed on re-entry; nevertheless, they weren't, and post-mission repairs to the shuttle's damage from insulation strikes became increasingly seen as nothing more than a routine maintenance issue (Mullane, 2006).

Vaughan discussed the similarities between the *Challenger* and *Columbia* losses in a book chapter (Vaughan, 2005) and has documented her experience on the CAIB and

her subsequent interactions with NASA in a 40-page journal article (Vaughan, 2006). The CAIB report is probably the first major US government accident investigation that explicitly states the cause of the disaster to be a socio-technical system failure.

The approaches exemplified by the writings of Perrow and Vaughan are not the only ones. Studies of what are known technically as High-Reliability Organizations (such as emergency rooms in hospitals, firefighter teams, and the flight-deck operations crews on aircraft carriers) have revealed that there are social and organizational, as well as technical, solutions to creating resilient socio-technical systems: see, for example, Roberts (1990); Weick & Sutcliffe (2007); and Reason (2008). The results from these studies indicate that there is no traditional, “pure” engineering approach that is suitable for ultra-large-scale systems. Multi-disciplinary approaches, that integrate the social with the technical, need to be developed: so-called *socio-technical systems engineering*.

But what does this academic literature on the study of technology failures offer to teach us about the events of May 6th, 2010?

Of course, the Flash Crash was by no means the first failure in a major financial market. As anyone reading this paper must surely be aware, in July 2007 the investment bank Bear Stearns was the first in what turned out to be a sequence of major financial institutions to signal that it had suffered significant losses on subprime hedge funds, triggering a sudden dramatic reassessment of counterparty risk in most major financial institutions around the world which led, *inter alia*, to the UK’s Northern Rock consumer bank being the first to suffer a full-scale public bank run in 150 years; and to the US government bailing out insurance giant AIG, mortgage providers Freddie Mac and Fannie Mae, and yet famously not extending a lifeline to Lehman Brothers, which turned out not to be too big to fail, and duly went bust.

Taking a longer historical perspective, the crisis of 2007-08 was just one in a sequence that stretches back through the collapse of the LTCM hedge-fund in 1998; the “black Monday” crash of October 1987; the US savings-and-loan crisis of the mid-1980’s; the Wall Street Crash of October 1929; the South-Sea Bubble of the 1720s; and the Tulip Mania of the 1630s.

This history of financial crises has been documented in a popular text by Kindleberger (2001), and with more academic rigour by Gorton (2010). The events of 2007-08 have been recounted from a number of journalistic perspectives, of which Lewis’s (2010) and Tett’s (2009) are notably thorough and well written. Tett’s perspective is particularly insightful: she is a senior executive journalist for the *Financial Times* but has a PhD in social anthropology, and this clearly influences her analysis. Tett was one of the few journalists to warn of the impending crisis before it came to pass, and notes various events that are clear instances of normalization of deviance. Lewis’s brilliant book tells the story of the few individuals who recognized that deviance, and bet on the markets failing. For more scholarly, academic, studies of the sociology of the financial markets, see the works of Edinburgh sociologist Donald MacKenzie and his colleagues (MacKenzie 2008a, 2008b; MacKenzie *et al.* 2008), although all of those pre-date the turmoil of the subprime crisis.

One significant difference between previous financial crises and the Flash Crash is the speed at which they played out. In the past quarter of a century, financial-market

trading has shifted from being a largely human, face-to-face activity, to being phone-and-screen-based rather than face-to-face, but still largely requiring a human at each end of the phone or screen. But within the past decade a fundamental technology-led shift has occurred. Increasingly, the counterparties at either end of the trade, at each end of the telecoms cable, are pieces of software rather than humans. Algorithmic trading systems are increasingly trusted to do trading jobs that were previously done by human traders, and to do jobs that would require super-human data-integration abilities in a person.¹⁰ As was seen on May 6th, the system-wide interaction between multiple independently-engineered, independently operated, and independently managed automated trading systems had at least one unknown catastrophic failure mode. A major proportion of traders in the markets are still human, but to understand today's markets it is necessary to study the interaction of these human traders with their automated counterparts; that is, we need to study the socio-technical system.

The danger that normalization of deviance posed in high-frequency automated trading systems in the global financial markets, and the possibility of major catastrophe happening within very short time-scales, was discussed in a strategic briefing paper written by one of us for the UK Government's Office of Science, first draft of which was submitted in January 2010 and the final version of which (Cliff, 2010) was submitted to the government nine days *before* the Flash Crash. Similarly, in the US at least one academic was repeatedly warning the SEC of the likelihood of a Flash Crash type of event in the year running up to May 6th 2010 (Angel, 2009a, 2009b, 2009c; Angel *et al.*, 2010; Angel 2010a, 2010b).

We think it is reasonable to argue that the Flash Crash was, at least in part, a result of normalization of deviance. For many years, long before May 6th 2010, concerns about systemic effects of rapid increases in the price volatility of various instruments had led several exchanges to implement "circuit breaker" rules, requiring that trading in a security be suspended for some period of time if the price of that security moved by more than some percentage within a sufficiently short time-period. For instance, the London Stock Exchange first adopted circuit-breakers, now known there as Automated Execution Suspension Periods (AESPs) and Price Monitoring Extensions (PMEs), shortly after the 1987 Black Monday crash; and Chi-X Europe enforces "order-entry controls" that prevent orders being entered that are more than 20% away from the current price (Flinders, 2007; Grant, 2010). In response to the Flash Crash, the USA's SEC has now enforced similar mechanisms in the US markets with the aim of preventing such an event re-occurring. In fact the move toward introducing circuit-breakers in the US pre-dates the Flash Crash by more than two years: it had been proposed in an influential report on the sub-prime crisis from the Institute of International Finance (IIF, 2008) but seems to have been actively resisted until the events of May 2010. Thus, it seems plausible to argue that before the Flash Crash occurred there had been some significant degree of normalization of deviance: high-speed changes in the prices of equities had been observed, market participants were well aware that that could lead to a high speed crash, but these warning signals were ignored and the introduction of safety measures that could have prevented them was resisted.

¹⁰ The history of the spread of technology innovations in the financial markets, and some likely future developments, are discussed in a recent review by Cliff, Brown, & Treleaven (2011).

Moreover, it could plausibly be argued that normalization of deviance has continued to take place in the markets *since* the Flash Crash. The SEC's introduction of circuit breakers seems to have been offered, and largely accepted, as the one necessary solution for preventing another similar event; and (so the story goes) all is now well. We are told that adding circuit breakers firmly shuts the stable door. Admittedly, this was done only after the Flash Crash horse had bolted, but at least the door is now shut.

Now, for sure, the introduction of circuit breakers means that the US markets today are not the same markets as they were on May 6th 2010. With circuit breakers added, those markets, and the other markets around the world that they are coupled to (i.e., the entire global financial market system) should be in a new dynamic regime – that is, their market dynamics are different now. But the new dynamics are still not entirely known, and so the new regime is certainly not yet guaranteed to be safe. Despite the circuit breakers, the next Flash Crash could be lurking just around the corner.

There are anecdotal stories that the speed of price fluctuations occurring *within* the limits of circuit breaker thresholds seems to be increasing in some markets (See, e.g., Blas, 2011); and there is evidence to suggest that another Flash Crash was “dodged” on September 1st 2010, in a similarly bizarre period when quote volumes exceeded even those seen at peak activity on May 6th 2010 (Steiner, 2010), but no official investigation was commissioned to understand that latter event. Furthermore, the circuit-breaker mechanisms in each of the world's major trading hubs are not harmonized, exposing arbitrage opportunities for exploiting differences; computer and telecoms systems can still fail, or be taken down by enemies of the system, and the systemic effects of those failures may not have been fully thought through.

Of course, the next Flash Crash won't be exactly the same as the last one, the SEC's circuit breakers will probably see to that. But there are no guarantees that another event, just as unprecedented, just as severe, and just as fast (or faster) than the Flash Crash cannot happen in future. Normalization of deviance can be a very deep-running, pernicious process. After *Challenger*, NASA addressed the issue with the SRB O-ring seals, and believed the Shuttle to be safe. That was no help to the crew of *Columbia*.

Reassurances from regulators that all is now well are likely to sound somewhat hollow for as long as people can remember the near-total failure of the regulatory bodies to have anything useful to say about the subprime crisis until shortly after its severity was clear to even the most casual of observers. Light-touch regulation and its consequence for financial markets in the UK were discussed in the 2009 Turner Review¹¹, and the parallel systemic failure of the economics profession is discussed at length by Colander *et al.* (2009) and by Bootle (2009). The next market failure may well be a failure of risky technology that, like the Flash Crash, has no clear precedent.

The global financial markets, considered as a single ultra-large-scale super-system, is made up of components, of constituent systems. These constituents include the human traders and their trading procedures; the various electronic exchanges; the automated trading systems operated by the various investment banks and hedge funds; and their associated clearing, settlement and risk-management systems. All of these constituent systems have been developed, procured, operated and managed independently, although for some of them the development and procurement processes were informal, organic

¹¹ http://www.fsa.gov.uk/pubs/other/turner_review.pdf

growth rather than pre-specified projects. That is, the current global financial markets are, from a technology perspective, *systems of systems* (SoS). We explore the definition of “system of systems” in some detail in Appendix A.2.

A key issue with SoS is that the effects of failure in one or more of the constituents may be contained, or may ripple out in a domino-effect chain reaction, analogous to the crowd-psychology of contagion. Furthermore, SoS are often used in unanticipated circumstances and by unanticipated users. In such situations, the response of the constituent systems may not result in local failure but rather the combined local responses can trigger a global failure: this seems to be what happened in the Flash Crash. In this very definite sense, the global financial markets have become high-consequence socio-technical systems of systems, and with that comes the risk of problems occurring that are simply not anticipated until they occur, by which time it is typically too late, and in which minor crises can escalate to become major catastrophes at timescales too fast for humans to be able to deal with them. The extent to which the SEC/CFTC report attributes cause to a single rushed block-sale as a \$4.1bn hedge as the trigger-event in the Flash Crash seems comparable to the way in which the *Challenger* accident investigation report identified failed SRB O-rings: there is a wider socio-technical perspective that should not be ignored, and which was already being pointed to by some authors prior to the events of May 6th 2010 (Haldane, 2009; Cliff, 2010).

That the global financial markets have become ultra-large-scale complex IT-centric socio-technical systems is perhaps no surprise, given the wider context that IT systems have moved from back-office support (for payroll processing, say) firmly onto the critical path for very many enterprises and organizations, to the point where failure of the IT system can incapacitate an organization. For example, ten years ago a failure of the IT servers in a hospital would not have a major negative effect; whereas in the near future, once all data is digitized at the point of capture and integrated with patient’s historical data before delivery in an appropriate form to a healthcare practitioner, then when a hospital’s servers go down it will cease to be a functioning hospital and instead be a big building full of sick people, with highly trained professionals frantically tapping the touch screens on their PDAs/tablet-computers, wondering where the data went. Similar stories can be told, or are already plausibly foreseeable, in very many private-sector, public-sector, and defence organizations in most industrialized economies.

Most notably, such issues have for some time been a growing, major concern in those areas of systems engineering where system failures can result in hundreds or thousands of fatalities or where, in the limit, system failures pose existential threats to entire nations: the engineering research literature in aerospace, nuclear, and defence systems may well be a source of experiences and new tools and techniques that could be applicable to the financial markets, although it is doubtful that any techniques yet exist that address the unique characteristics of ultra-large-scale systems. The manifestly dire consequences of failure in aerospace, nuclear, and defence systems, and also of course in automotive systems, has led to the development of engineering teaching and practices specific to the development and maintenance of safety-critical, high-integrity systems: a field known as high-integrity systems engineering (HISE), which we briefly review in Appendix A.1 of this document.

So, the concerns expressed here about modern computer-based trading in the global financial markets are really just a detailed instance of a more general story: it seems likely, or at least plausible, that major advanced economies are becoming increasingly reliant on large-scale complex IT systems (LSCITS): the complexity of these LSCITS is increasing rapidly; their socio-economic criticality is also increasing rapidly; our ability to manage them, and to predict their failures before it is too late, may not be keeping up. That is, we may be becoming critically dependent on LSCITS that we simply do not understand and hence are simply not capable of managing. This is something that we illustrate, purely notionally, as a single three-line graph, shown in Figure 1.

We, the authors of this review, each work for major national strategic initiatives intended to address these issues. In the UK, the National Research and Training Initiative in the Science and Engineering of LSCITS was started in 2007 as a strategic investment with the primary aim being to foster the formation of a new community of researchers and practitioners with the training and experience appropriate for dealing with future software-intensive systems engineering dominated by LSCITS issues (Cliff *et al.* 2006). At pretty much exactly the same time as the UK LSCITS Initiative was being planned and set up, entirely independently, in the USA the US Army




Fig. 1. The Complexity Crossover Crisis. The complexity of information and communications technology (ICT) socio-technical systems of systems (SoS) has increased dramatically since ICT was first commercialized in the 1950s, and in recent years the socio-economic criticality of ICT SoS has also sharply increased, as very many enterprises and organizations in advanced economies have become dependent on the availability of ICT functionality as a key component on the critical paths of their operations. Over the same period, there is increasing concern (and growing evidence) that our ability to manage and predict the behavior of these critical ICT SoS is not increasing at the same pace, and so at some point in time there is the potential for crisis, where major socio-economic systems are critically dependent on ICT SoS whose complexity is beyond that which we can manage. We are deliberately non-committal on the precise timing of this crossover point: for some companies or industrial sectors it could be a decade or more away, for others it could have happened already.

commissioned a team of world-class researchers led by the Software Engineering Institute (SEI) at Carnegie Mellon University to conduct a study of ultra-large-scale systems software. The study resulted in a major report that argued the necessity for the USA to invest in ultra-large-scale systems engineering research, to safeguard its international dominance in information systems; this authoritative report marked the first major output from the SEI Ultra-Large-Scale (ULS) Systems Project (Northrop *et al.*, 2006). For a brief overview of the ULS Systems project, the UK LSCITS Initiative, and other related projects, see Goth (2008).

3 Where Next for the Financial Markets?

One criticism that is sometimes leveled at the academic study of technology failures is that there is perhaps a tendency to be wise after the event. That is, a large amount of the work is descriptive (saying what happened) but not sufficiently predictive (saying what could happen next) or prescriptive (saying what should be done differently in future, to predict or prevent such failures from re-occurring).

One possible approach, which side-steps the need for specific predictions, is to accept that technology failures are simply to be expected every now and again as part of the Darwinian arms-race dynamics at the leading edge of technology-dependent institutions, comparable to natural “failures” such as the species-extinctions that occur relatively routinely in biological ecosystems, when viewed over evolutionary timescales, and which also seem to follow a power-law distribution (small failures being common, big failures being rare: see e.g. Ormerod, 2006). Such a perspective may be well-aligned with the new schools of thought in economics and the study of technology innovation that are influenced by complexity science and autopoiesis (e.g. Ormerod, 1998; Blume & Durlaf, 2005; Beinhocker, 2007; Arthur, 2009), but taking a Darwinian, *laissez-faire*, “stuff happens” approach isn’t particularly helpful in the quest for new engineering practices, for predictive and preventative tools and techniques.

Recently, there has been growing recognition within the engineering community that the engineering of systems in which failures are expected, and where the systems are resilient to those failures, may require a fundamental reassessment of established engineering teaching (see, e.g., Hollnagel *et al.* 2006). Similar views have also been expressed, earlier, in the business administration literature dealing with the management of large-scale technology-driven projects (Collingridge, 1992). It seems reasonable to suggest that changes are necessary both in engineering practices, and in the coordination, incentivization, and management of projects, for all LSCITS including those underlying the global financial markets. But such changes are likely to take time, and while we wait for them to take effect it would be good to have a viable near-term strategy, one that would potentially offer major payoff within five to seven years (seven years is long enough to achieve quite a lot, given enough resources: the US Apollo programme took seven years, from John F. Kennedy’s famous speech to Neil Armstrong’s famous small step.) In the following pages, we outline one such strategy. It will require national-scale investment, to create a national-scale strategic resource (or, perhaps, international collaboration to create a shared multinational resource, rather like the CERN Large Hadron Collider or the European Space Agency’s *Ariane* space rocket).

The proposed strategy is simple enough to state: build a predictive computer simulation of the global financial markets, as a national-scale or multinational-scale resource for assessing systemic risk. Use this simulation to explore the “operational envelope” of the current state of the markets, as a hypothesis generator, searching for scenarios and failure modes such as those witnessed in the Flash Crash, identifying the potential risks before they become reality. Such a simulator could also be used to address issues of regulation and certification. Doing this well will not be easy and will certainly not be cheap, but the significant expense involved can be a help to the project rather than a hindrance.

Explaining and justifying all that was written in that last paragraph will take up the next several pages.

For most engineering and scientific domains, in recent years it has become increasingly commonplace to rely on high-precision computer simulation as a means of studying real-world systems. Such simulations offer the possibility of evaluating the performance of proposed systems that have not yet been physically constructed, and of exploring the response of existing real-world systems to different operating-environment conditions, and to alterations of the system itself, allowing “test-to-destruction” without actually destroying anything. Engineers interested in aerodynamic flows over aeroplanes and cars, or around buildings, or hydrodynamic flows around a ship’s hull, can routinely call upon highly accurate computational fluid dynamics (CFD) models to evaluate these flows in simulation, rather than building physical models to test in wind-tunnels or test-tanks. Almost all silicon chip designs are evaluated in microelectronics circuit-simulators such as SPICE (e.g. Tuinenga, 1988) before the chip-producers make the final (and most expensive) step of committing their design to fabrication. Fissile nuclear reactions can be simulated with sufficient accuracy that designs for nuclear power stations, and for nuclear weapons, can be evaluated in simulation without splitting a single atom. In most advanced economies, weather forecasts are produced by national agencies on the basis of detailed sensor readings, and advanced computer simulations, that allow for accurate short-term and medium-term predictions of the future. Similar stories can be told in computational drug design, computational systems biology, and so on. Advocates of the use of predictive computer simulations in science and engineering have argued that this approach now represents a well-established third paradigm within science, in addition to the two long-established paradigms of empirical observation and theoretical modeling/generalization (see e.g. Gray, 2009, p.xviii).¹²

It’s important to be clear about the nature of the predictive simulation models that we are advocating here. Meteorological simulations are predictive in the sense that they make weather-forecasts, specific projections about the likely future state or states that the real-world weather system may find itself in; that is, they say what is about to happen, or what would be likely to happen under specific circumstances. This is the most familiar practical use of simulation modeling. But there is a second use to which simulation modeling can be put: simulating a model of some system allows the model itself to be explored; in this sense, the model is an embodiment, an implementation in computer-code, of a theory of how the thing being modeled works. This second type of simulation

¹² The use of predictive simulations in engineering safety-critical complex systems-of-systems is discussed further in Appendix A.4.

modeling often starts out as essentially exploratory, with the intention of delivering explanatory insights that would otherwise have been difficult or impossible to come by.

One illustrative example of this kind of simulation-as-explanation is Schelling's (1971) model of racial segregation, where a very simple iterative process (i.e., an algorithm) operating on black or white markers positioned on a grid of square cells arranged chessboard-like over a two-dimensional space (i.e., an abstract model of the real world) was used to explore and explain how groups of people expressing only very weak preferences for wanting to live near to neighbours of the same race could lead over time to total segregation with large spatial clusters all of one race or the other. That is, the Schelling model, when simulated, showed in its dynamics an emergent behavior at the system-level that was unexpected and difficult to predict from mere inspection of the set of rules that the simulated people (the "agents" in the model) were specified to follow; Schelling was subsequently awarded the 2005 Nobel Memorial Prize in Economic Sciences. For a recent collection surveying such exploratory and explanatory simulation modeling in social sciences research, an approach now widely known as agent-based modeling, see Epstein (2007); and for a review of foundational work in agent-based computational finance, see LeBaron (2000).

Of course, computational simulations are currently also routinely used by financial institutions: Monte-Carlo techniques are used to solve and explore options-pricing models, to evaluate value at risk, to back-test trading algorithms on historical data, and to perform stress-tests on individual financial instruments or on portfolios of such instruments. But historically it has been much less commonplace to simulate entire markets at a fine-grained level to study issues in overall system behaviour in an exploratory fashion.

In an excellent book, Darley & Outkin (1997) give a detailed description of how they used complex adaptive systems (CAS)¹³ agent-based simulation-modeling techniques to explore the consequences of the Nasdaq exchange's move from quoting prices expressed as multiples of sixteenths of a dollar to fully decimalized prices, expressed as multiples of one hundredth of a dollar (i.e., as dollars and cents). In the language of the markets, this was exploring the effects of a reduction in the Nasdaq "tick size" from \$0.0625 to \$0.01. Nasdaq had previously changed its tick-size from $\$1/8^{\text{th}}$ to $\$1/16^{\text{th}}$ in 1997, and there was evidence to suggest that at the same time there had been a change of strategies among the market participants trading on Nasdaq. Nasdaq commissioned Darley & Outkin to construct a detailed simulation model to evaluate possible effects of changing the tick-size to \$0.01, in advance of the actual decimalization which was completed in April 2001; that is, Darley & Outkin were dealing in predictions, not postdictions. Darley & Outkin's book recounting this predictive-simulation CAS work was published several years later. In it, they state:

"While building the simulated model of the market, we interacted extensively with many market participants: market-makers, brokers, traders, large investors, etc. We found this interaction invaluable – as a source of information in particular on often subtle details of market operations, as a venue for verifying our assumptions and simulations

¹³ The definition of a "complex adaptive system" is explored in more depth in Appendix A.3.

results, and at times as a source of constructive criticism. One conversation with a market maker still stays clear in our minds. He was supportive, but skeptical. The core of his skepticism lay in this question: how one can model the fear and greed often ruling the market behavior? This is a valid point: while fear and greed affect markets immensely, as has been time and again demonstrated by numerous booms and busts, understanding of underlying individual and mass psychology is lacking.

“In our approach we address this problem by explicitly modeling strategies of individual market participants, by allowing those strategies to evolve over time due to individual learning or evolutionary selection, and by allowing to [*sic*] investigate various what-if scenarios by using user-defined strategies.”

(Darley & Outkin, 1997, pp.5-6)

Darley & Outkin report that the results from their CAS simulations led them to make six substantive predictions before decimalization was enacted, and that events subsequent to the actual decimalization largely supported all of those predictions, except one (concerning the upper bound on the increase in trading volume, which had not yet been reached by the time that Darley & Outkin published their book).

Darley & Outkin’s book describes a simulation model of one specific real-world exchange, and was the first to do so in such detail. For other studies of using CAS simulation-modeling techniques to explore how the collective behaviour of individual trader-agents can give rise to certain market-level phenomena, see e.g. Palmer *et al.*, 1994; Cliff & Bruten, 1999; LeBaron, 1999; Levy *et al.*, 2000; and Tesfatsion & Judd, 2006.

Given the success of Darley & Outkin’s work, which is now over a decade old, it seems entirely plausible to propose that a similar complex-adaptive-systems, evolutionary agent-based, predictive simulation model could be constructed to assess the dynamics and behavior of individual financial markets, or indeed of the entire global financial market system. Obviously, it would be a major endeavour to create such a model, requiring national-scale levels of investment and ongoing funding to provide appropriate resources of human capital and computing power.

Nevertheless, there is an obvious precedent in most advanced economies: very many countries fund, as a national utility, a meteorological agency such as the UK’s Met Office¹⁴. Combining real-time sensor data from satellites and ground-based observation stations with historical data and advanced, highly compute-intensive, predictive simulation models, the Met Office is able to give accurate near-term weather forecasts with a high spatial precision.

The famously chaotic nature of weather systems (Lorenz, 1963) means that accurate longer-term predictions remain more problematic, and the same is very likely to be true of long-term predictive models of the financial markets, but there is a well-established technique used in meteorological forecasting that should also be of use modeling the markets: so-called *ensemble forecasting*, where the same model is re-run many hundreds or thousands of times, with each fresh run having minor variations in the initial conditions, and/or a different sequence of random numbers generated in the modeling of stochastic factors (see, e.g., Smith, 1995, 2002). From a thousand runs

¹⁴ <http://www.metoffice.gov.uk>

(say) of a model aimed at predicting the weather 48 hours into the future, it may be that 243 of the simulations show snowfall on a particular area, 429 show rain, and the rest predict no precipitation; with these results, the forecast for two day's time would be a 24% chance of snow, a 43% chance of rain, and a 33% chance of it staying dry. In this sense then, the forecast is a probability function over the space of possible outcomes. Here we have only three mutually exclusive outcomes; a more sophisticated model might give a probability density function (PDF) over the space of possible precipitation levels measured to the nearest millimeter per unit of area, and also a separate PDF over the space of possible ambient temperatures, measured to the nearest degree Celsius; taken together, the two PDFs would form a prediction of whether water would fall from the sky, and whether it would fall as rain or as snow.

So, the chaotic nature of financial markets is not necessarily an impediment to the development of predictive simulation models, so long as sufficient computing resources are made available to allow for ensemble forecasting. In fact, it is likely that the real value of the ensemble forecasting work would be in running very many simulations (perhaps tens or hundreds of thousands or more) in the search for those extremely rare but devastatingly problematic combinations of circumstances that have become widely known as *Black Swan* events (Taleb, 2007). It seems reasonable to describe the May 6th Flash Crash as a Black Swan event, and maybe the likelihood of such an event could have been predicted in advance, if a suitably detailed simulation model had been available beforehand. Of course the simulation would not have predicted that the crash would occur on May 6th, and would probably not have identified the precise trigger event. But it does seem entirely reasonable to argue that an appropriate model may have identified in advance the existence of a nonzero probability that if a certain type of order is executed in sufficiently large volume with certain (lack of) constraints on its execution pattern, that order could interact with the existing population of traders (both human and machine) to cause a "hot-potato" dynamic leading to a sudden, largely irrational, mass sell-off, exposing stub-quote values as reference prices, and leading major exchange-operators to declare self-help against each other, which is the current official story (CFTC & SEC, 2010a,b).

The possibility of such a sequence of events does not seem to have been much discussed prior to May 6th; perhaps if an appropriate national-level or international-level modeling facility had been operational, people would have been aware of the latent risk. Central government treasury departments in most economies have for many years (since before the advent of electronic computers) run large-scale macro-economic models for forecasting, but as far as we are aware there are no mature models used to understand and predict issues of systemic risk in the financial markets. Such a systemic-risk market simulator system could also be used for training market practitioners and regulators in dealing with rare but extreme situations, in much the same way as civil and combat aeroplane pilots are trained to deal with various rare but serious aircraft system failures by flying many hours of simulator practice, so that in the unlikely event of such a failure occurring on a real flight, the pilot can rely on her lessons learned and experience gained in the simulator. The rescue of *Apollo 13* owed an awful lot to the availability of accurate simulation models (physical electro-mechanical ones rather than purely virtual computer simulations) at NASA Mission Control. The simulators had been developed to train the astronauts in dealing with various mid-mission failure situations, including using the Lunar Excursion Module

as a “lifeboat”, as was necessary on *Apollo 13*; after the explosion on *Apollo 13* the simulators also became the test-bed for evaluating novel procedures necessary to keep the crew safe and the crippled ship on its return course.

Simulation models used in complex systems engineering are typically not intended for training humans *within* the socio-technical system being simulated; rather, any human agents within the real system are *also* simulated in the model of that system. Nevertheless, the use of simulation models as scientific evaluation and training tools for humans dealing with unusual complex situations has a long history: see, e.g., Sloan (1981) and Dorner (1990, 1997), yet there is currently very little in the way of comparable use of personnel training/evaluation simulators in the financial markets. Trainee traders typically learn the ropes by running “dummy” accounts, keeping a record of trades that they would have made, but did not actually execute, so that any losses are merely on paper; this can be done using live market data, and trading strategies can also be back-tested on historical data. A notably more sophisticated simulator, integrating real-time price feeds, was developed in a collaboration between the University of Pennsylvania and Lehman Brothers, the Penn-Lehman Automated Trading project, described by Kearns & Ortiz (2003).

While techniques such as these work well as training for situations where the trader’s activity has no immediate effect on the prices of the securities being traded, they cannot readily model *market impact*, where the mere act of revealing the intent to buy or sell a large quantity of a security means that other traders in that security (potential counterparties to the trade) alter their prices before the transaction occurs, in anticipation of the change in price that would otherwise result after the transaction has executed. Furthermore, simulators based on regurgitating historical data offer essentially nothing toward understanding the current or future overall system-level dynamics of the system: they can tell you what happened, but not what might happen next, nor what might have happened instead. Simulators for evaluating trading strategies on historical data are sometimes referred to as financial-market “wind-tunnels” (e.g. Galas *et al.*, 2010). A financial-market wind-tunnel is certainly useful in refining the dynamics of an individual trading strategy, in much the same way as a traditional engineer’s wind tunnel is useful in refining the aerodynamics of a new aeroplane or car. But financial-market wind-tunnel simulators are of zero help in understanding systemic issues such as financial stability, for much the same reason that an aerodynamicist’s wind tunnel can tell you nothing about system-level phenomena such as traffic congestion in a city’s street, nor air safety in a nation’s skies.

More fancifully, it may also be worth exploring the use of advanced simulation facilities to allow regulatory bodies to act as “certification authorities”, running new trading algorithms in the system-simulator to assess their likely impact on overall systemic behavior before allowing the owner/developer of the algorithm to run it “live” in the real-world markets. Certification by regulatory authorities is routine in certain industries, such as nuclear power or aeronautical engineering. We currently have certification processes for aircraft in an attempt to prevent air-crashes, and for automobiles in an attempt to ensure that road-safety standards and air-pollution constraints are met, but we have no trading-technology certification processes aimed at preventing financial crashes. In the future, this may come to seem curious.

We're not arguing here that predictive simulation models are a "silver bullet", an easily achievable panacea to the problem of assessing systemic risk and identifying black-swan failure modes: developing and maintaining such models would be difficult, and would require a major research investment. It seems very likely that quantitative analytical techniques such as probabilistic risk assessment (see e.g. Stamatelatos *et al.*, 2002a, 2002b; Dezfuli *et al.*, 2009; Hubbard, 2009) and probabilistic model-checking (e.g. Calinescu & Kwiatkowska, 2010; Calinescu, Kikuchi, & Kwiatkowska, 2010) would also need to be involved, in sufficiently extended forms, to help constrain the (otherwise impossibly vast) space of possible situations and interactions that would need to be explored by the simulations.

While there is no shortage of challenges in simulating the technical entities in socio-technical systems, simulating the social entities is almost always even more problematic, and this is something that doesn't have to be addressed by meteorological forecasting systems. Whether individual human agents, or groups of humans operating and interacting as teams or large organizations, the social entities in a socio-technical system are frequently present in virtue of the fact that they are needed to perform roles and discharge responsibilities with levels of flexibility, adaptability, and subtlety that are beyond the capability of automated systems. Modelling those kind of issues certainly presents a large number of deep technical challenges, and it is fair to say that the representations of social entities in many HISE models are often quite primitive: simple probabilistic models of humans switching from "safe" to "error" status are not uncommon. More sophisticated nondeterministic behavioural models such as those based on Markov chains (e.g. Haccou & Meels, 1994; Benveniste *et al.*, 2003), and computational implementations of models of behaviour and motivation from the ethology literature (such as Lorenz's well-known hydraulic model explained in his 1966 book *On Aggression*) have all been explored in the research field that studies mechanisms for the generation or simulation of adaptive behaviour in animals (including humans) and synthetic agents, including those that are needed to model human ingenuity and adaptivity in predictive simulation models. One of the biggest drivers for this research is the need for creating believable synthetic agents in virtual environments such as computer games, yet the work presents deep challenges and is also directly relevant to simulations of real-world scenarios for training and evaluation purposes (so-called "serious games")¹⁵: see, e.g., Blumberg, 1996; Ivanov, 2002; Tomlinson & Blumberg, 2002; Horswill 2009. In some limited domains, for instance the modeling of emergency egress by crowds of humans from stricken structures (such as burning buildings or sinking ships), where there is reasonable data for how humans do behave in such circumstances, such

¹⁵ See, for example, the Serious Games Institute at <http://www.seriousgamesinstitute.co.uk>, the Serious Games Initiative at <http://www.seriousgames.org/>, and the various research outputs from FutureLab on Games and Learning, Serious Games in Education, Game-Based Experience in Learning, and Teaching with Games, all available at <http://www.futurelab.org.uk/projects/>. An extensive report on the use of serious games in military education and training was produced by Caspian Learning for the UK Ministry of Defence: http://www.caspianlearning.co.uk/Mod_Defence_Academy_Serious_games_Report_04.11.08.pdf.

models have proven to be genuinely insightful (see, e.g., Johnson, 2005, 2006, 2008; Johnson & Nilsen-Nygaard, 2008)¹⁶.

The significant cost of constructing and operating such a simulation facility could possibly be met from the public purse via general taxation, or could perhaps be funded by direct contributions from the major financial corporations (banks, fund-management companies, exchange operators, insurers, etc.) operating in a particular country or group of countries. If funded as a public-sector project, it would of course be necessary to recognize that in addition to the significant technical challenges, the establishment of such a simulator facility also present significant budgetary challenges and the entire endeavour would need to stand up to a thorough cost-benefit analysis: this is an issue expanded upon by Bullock (2011).

However, it is not the case that the only way of building or running such a simulation facility is via public-sector financing. It is possible that a group of financial institutions could collaborate on, and co-fund, the necessary capital expenditure at start-up and ongoing operational costs. A UK precedent for this, albeit in a different industry sector, is the independent non-profit company CFMS Ltd¹⁷ that is jointly owned and operated by founding partners Airbus, BAE Systems, Frazer-Nash Consultancy, MBDA UK, Rolls-Royce, and Williams Formula 1 Engineering. CFMS exists to advance the theory and practice of simulation-based design processes, and has invested in its own high-performance computing facilities available in its Advanced Simulation Research Centre (ASRC). Given the importance of aerodynamics to many of the founding partners, there is a focus on computational fluid dynamics modeling in CFMS/ASRC, which is of no direct relevance to the world of finance. Nevertheless, the success of CFMS and ASRC shows that independent companies can indeed come together to co-found and co-run shared facilities as an investment in pre-competitive research and development capability.

If a major simulation facility was constructed, revenue could be generated from levying charges for anyone wanting access to it, and also possibly from using it as a training or certification facility. The potentially massive cost involved is not necessarily a disincentive: if the simulator was constructed on a minimal budget of (say) several hundred thousand pounds, it would be reasonably easy for financial corporations such as a hedge funds or investment banks to fund their own rival internal projects, probably much better-resourced, which would then detract from the public-good shared-utility nature of what is proposed here.

But, if the national-level simulator was funded by tens or hundreds of millions of pounds (and assuming that these pounds were spent wisely) then it is plausible that it would be so well resourced, and hence so much more detailed and/or accurate, that no private corporation could reasonably hope to compete with it, then all private corporations reliant on its results would have an incentive to contribute to the running costs, and the intellectual content, of the simulator facility as a common good. The facility would then be a pre-competitive shared resource: all contributing corporations would have access to details of its design and construction, and all would have access to its facilities for running experiments. Corporations would nevertheless be free to

¹⁶ See also <http://www.massivesoftware.com/real-world-simulation-gallery/>.

¹⁷ See www.cfms.org.uk

compete on the basis of what questions they ask of the simulator (details of each corporation's specific experiments could be kept confidential), and in how they then use the results from their experiments.

Of course the counterargument to developing a single utility facility is that this would concentrate risk: if the one national simulator is wrong, and everyone is using results from that simulator, then everyone's expectations or predictions are wrong at the same time. This is also manifestly true of national weather-system simulator facilities, and there is no shortage of examples of entire nations being taken by surprise when their state-funded monopoly weather-forecasting services got it wrong.¹⁸

One approach to mitigating this risk may be to enforce so-called "*n*-plex redundancy", as is common in the design of controllers for aerospace and defence systems, where the same control-system functionality is implemented by *n* multiple parallel systems, each designed and implemented by different independent suppliers, often constrained to not use the same core technologies (such as particular processor chips, programming languages and compilers, third-party suppliers, etc). The rationale for such an approach is that, while each of the *n* redundant systems may have one or more failure modes, the likelihood of all *n* systems having the same (or overlapping) vulnerabilities is greatly reduced by the active prevention of them sharing common components and/or development paths. Thus, so the argument goes, while one or more of the individual systems may fail from time to time, the remaining parallel redundant systems will most probably remain operational, and thereby coherent control will be maintained. So, maybe the best approach is for a national agency to commission some small number *n* of competing predictive simulation models, adopting or guided by the principle of *n*-plex redundancy, in the hope that the collective indications from the suite of *n* independent simulations can be trusted more than the lone voice of a single model.

A more thorny issue is the effect of the feedback loop from the model(s) back to the market systems being modeled. Results from a predictive simulation model of the weather do not actually alter the weather, but results from a market simulation may have a significant effect on the subsequent behavior of agents within the real-world markets that the simulator is a model of. There is prior evidence of self-fulfilling prophecies driving market dynamics, such as the theory that market activity is somehow affected by the number of sunspots. There is no *a priori* causal mechanistic explanation for why sunspots might affect market activity, but someone once proposed that there was at least a correlation between sunspot numbers and markets rising or falling; all that was then required was for enough people to believe in the correlation and to allow that belief to alter their trading activity in the markets. This

¹⁸ On October 15th, 1987, a UK Met Office forecaster reassured viewers on the BBC prime-time evening weather broadcast that there was not a hurricane coming, in an attempt to quell earlier speculation. Later that night the south of England was hit by the worst hurricane-force windstorm for over 250 years, with speeds gusting to 120mph for several hours, causing huge amounts of damage and unprecedented levels of disruption for days afterwards. Other nations' meteorological forecasting services on mainland Europe, using different monitoring and prediction models, had given more accurate forecasts of the windy weather that night.

shared belief then *became* the causal link: if enough people are counting sunspots and using that to drive their market behaviour, then an increase in the number of sunspots will indeed affect the market in the manner that was “predicted” by their belief, thereby reinforcing the conviction of those who already hold the belief and helping to convert non-believers. The causal feedback loop from predictive simulations back to the real-world markets is something that will need to be handled well, but it is not necessarily a problem: the feedback could have a positive effect, dampening unwelcome dynamics.

To conclude, we observe that there is an old saying: “if it ain’t broke, don’t fix it”. This is certainly wise guidance in very many situations. But it is important to remember that for some systems, when they do actually break, they go so catastrophically wrong so superhumanly fast that the safest option for such a system really is to fix it while it ain’t broke, because that is the only decent chance you’ll get. This is the case for many large-scale complex IT systems (LSCITS). Ensemble forecasting via *n*-plex redundant predictive simulation models is not cheap, is not easy, and is certainly far from perfect, but it may just be the best option currently available.¹⁹

The novelty of this proposal can perhaps be judged by the fact that the most recent comprehensive UK industry-focused review examining mechanisms for achieving supervisory control of systemic risk (Bonisch & Di Giammarino, 2010) mentions predictive simulation modeling only obliquely, in passing; but that same report also mentions the Flash Crash only once, in passing, too, as if such a manifestly deviant event was already normalized.

Nevertheless, we are certainly not the only people to be making such proposals: see, e.g. (Farmer & Foley 2009; Economist, 2010; Harford, 2011; Salmon, 2011), and the UK Government Office for Science’s recent *Foresight* project exploring the future of computer trading in the financial markets has commissioned two excellent reviews that discuss aspects of the idea in more detail: see Bullock (2011) and Farmer & Skouras (2011). The UK already has significant investments in university research centres that could make valuable contributions to this approach.²⁰

In his April 2009 speech *Rethinking the Financial Sector*, Andy Haldane, Executive Director for Financial Stability at the Bank of England, argued that three steps were necessary to safeguard against another series of events like the 2007/08 subprime crisis: all three steps deal with the global network of interacting financial institutions. Haldane’s argument was that we should work first to map that network; then take steps to better manage and regulate the existing network; and then explore useful ways in which the network could be restructured or otherwise modified. We contend that all three of these steps (map, manage, & modify) could, and in fact should, be performed

¹⁹ In the interests of balance, for recent counterarguments to the use of simulation models, see Turkle (2009).

²⁰ Major UK academic research centres that could be involved include: the Bristol Centre for Complexity Science (<http://bccs.bristol.ac.uk>); the Bristol/Bath Systems Engineering Centre (www.bristol.ac.uk/eng-systems-centre/); the Southampton Institute for Complex Systems Simulation (www.icss.soton.ac.uk); the UCL PhD Centre for Financial Computing (<http://fc.cs.ucl.ac.uk/phd-centre>); the York Centre for Complex Systems Analysis (www.yccsa.org); and the UK Large-Scale Complex IT Systems Initiative (www.lscits.org).

via an appropriate simulation-model-based engineering approach: creating and maintaining the model would be Haldane's mapping exercise; once operational, the effects of different regulatory actions, and any potential restructuring of the financial network could be explored and evaluated in the model too.

4 Summary

The Flash Crash of May 6th 2010 was a sudden and dramatic failure in a ultra-large-scale software-intensive socio-technical system (the US financial markets) with prices running wild at a speed and magnitude of volatility that were without historical precedent. The fact that there was not major lasting damage to the global financial markets is perhaps more due to luck than judgement: if the down-spike in the Flash Crash had occurred five minutes before market close in New York, it's plausible that could have triggered a contagious global sell-off that then went on to wrap around the world.

Yet from a broader perspective it is clear that the Flash Crash was just one more in a sequence of failures of risky technology, and quite plausibly such an event was made more likely via a prior process of financial-market practitioners becoming increasingly tolerant of unexpected events, previously thought to be unacceptable, not resulting in disasters: that is, via a process of normalization of deviance.

The problems posed by attempting to engineer and manage reliable ultra-large-scale complex adaptive socio-technical systems of systems are becoming ever more clear, but further research is needed to develop appropriate tools and techniques. System-of-systems issues of scaling, normal failure, heterogeneity via organic growth, and emergent behavior all have to be addressed. Parallel running of multiple redundant predictive simulation models is one approach that may now be applicable for assessing and controlling systemic risk in the financial markets.

The engineering of LSCITS and ULS socio-technical ecosystem system-of-systems is in its infancy: it has significant differences from traditional engineering of smaller-scale systems, and developing rigorous trusted approaches may turn out to be a long haul. The UK's LSCITS Initiative and the USA's Ultra-Large-Scale (ULS) Systems Initiative are each articulations of national strategic concerns. Both represent a sizeable step toward developing a new community of practitioners and researchers who are conversant with all the necessary subfields that can contribute to addressing issues in the science and engineering of such systems, forming those communities of practice will take several years of sustained investment. Without doubt this is not merely responding to a national need but an international one. We, the authors of this report, welcome any researchers, practitioners, regulators, policy-makers or sponsors who would like to become involved in the LSCITS and/or the ULS Systems initiatives. The intellectual challenges are significant, but not insurmountable; the potential societal savings are massive, and the scale is truly global.

Acknowledgements. This chapter is slightly revised and extended/updated from a report that we co-authored in early 2011 for the UK Government Office for Science's *Foresight* project on the future of computer trading in the financial markets. We thank the following people for valuable conversations and/or for their comments on previous versions of this document: Prof. Philip Bond, University of Bristol and University of

Oxford; Prof. Seth Bullock, University of Southampton; Andy Haldane, Bank of England; Kevin Houston, FIX Protocol Ltd; Prof. David Parkes, Harvard University; Lucas Pedace, UK Government Office for Science; Dr. John Rooksby, University of St Andrews; Tim Rowe and his colleagues at the UK Financial Services Authority; Prof. Ian Sommerville, University of St Andrews; Dr. Gillian Tett, *The Financial Times*; and Nigel Walker, UK Financial Services Knowledge Transfer Network.

Appendix: High-Integrity Large-Scale Complex Ecosystems

In this Appendix we take a quick tour through the concepts and approaches from current systems engineering that are relevant to the discussion just presented, but for which going into detailed explanation or definition would have been a distracting diversion from the flow of our argument. In sequence, here we briefly review high-integrity approaches to systems engineering (Appendix A.1); the definitions of Systems-of-Systems (A.2) and Complex Adaptive Systems (A.3); and then selected current leading-edge approaches to the high-integrity engineering of complex adaptive systems-of-systems (A.4).

A.1 High-Integrity Systems Engineering

High-integrity engineering techniques for safety-critical systems have a long heritage, and it's simply beyond the scope of this document to provide a comprehensive review of all the relevant background literature; for detailed surveys, see the review chapters in the recent PhD theses by Alexander (2007, pp.29-55), Despotou (2007, pp.41-76), and Hall-May (2007, pp.33-72).

It is commonplace in real-world engineering situations to be dealing with systems that simply cannot be guaranteed to be *absolutely* safe because key components in the system are known not to be *absolutely* reliable. If one of the key components is known to be 99.99999% reliable, that is an admission that there is a 0.00001% chance of failure; if failure of that component compromises the safety of the overall system, then there is a risk (small, but nonzero) that the system will become unsafe. Safety engineering has developed techniques for estimating the causal chains of events leading to failure, the attendant risks of failure, the effects of failure, and for reducing those risks and limiting their effects; in this sense then, risk and reliability are two sides of the same coin.

One of the earliest forms of risk and reliability assessment method, developed in the 1960's US aerospace and missile programs, is fault-tree analysis (FTA). FTA operates by the engineer first identifying "basic events" such as a fuse blowing or a relay-switch failing to open. Significant combinations of these basic events are then aggregated into a graph structure much like a family tree: compound events are formed via "gate" nodes that link basic events. It may be that basic events E1 and E2 and E3 *all* have to occur for a particular output fault F1 to occur: on the graph the event nodes E1, E2, and E3 would be shown as "daughters" of F1, with F1 denoted as an "and" gate. Other types of gate include: "or" (any one or more of the daughters triggers the compound fault); "combination" (the compound fault is triggered by any n

or more of the daughters occurring, for $n > 1$); “exclusive or” (exactly one daughter will act as the trigger); “priority and” (the daughter events have to all occur in a specific sequence); and “inhibit” (the daughter event occurs at the same time as some enabling condition). The daughter nodes of a compound event are not required to be basic events: they can be other compound events, and so it is possible to construct deep trees showing how basic events, combinations of basic events, and combinations of those combinations, can each combine to contribute to particular faults or failures in the system under analysis. Fault-tree analysis distinguishes between failure *effects* (such as a switch failing to make contact), failure *modes* (such as the switch’s contacts being broken, or the contacts having a very high resistance), and failure *mechanisms* by which those modes may come about (such as high resistance on the switch contacts being caused by corrosion of the contact surfaces, or by an insulating coating having been spilled onto them); this well-known safety-critical engineering practice is known as Failure Modes and Effects Analysis (FMEA). For further details, see e.g. Stamatelatos *et al.* (2002b).

FMEA and FTA, as just described, are essentially qualitative, deterministic, approaches. In recent years, there has been a concerted move toward developing quantitative approaches where numeric values represent measures of risk. An obvious, intuitive, risk metric is the probability of failure, and so the field is widely known as probabilistic risk assessment (PRA).²¹ Over much the same period, the field of mathematical statistics has undergone something of a revolution in the rapid adoption of the so-called *Bayesian* approach as an alternative to the long-established, traditional, *frequentist* approach, and this has been reflected in the PRA literature. For instance, in 2002 NASA published a 323-page guide to PRA procedures for its managers and practitioners (Stamatelatos *et al.*, 2002a) based on traditional frequentist statistics, but then in 2009 it published a new 275-page guide to PRA using Bayesian methods (Dezfuli *et al.*, 2009). Some authors, most notably Hubbard (2009), have argued forcefully that PRA should be the only game in town, but PRA is not without its critics and detractors: see, for example: Parry (1996); Slovik (1999); and Apostolakis (2004).

The opening page of NASA’s 2002 guide to PRA neatly summarises the history of its adoption in that organization:

“Legend has it that early in the Apollo project the question was asked about the probability of successfully sending astronauts to the moon and returning them safely to Earth. A risk, or reliability, calculation of some sort was performed and the result was a very low success probability value. So disappointing was this result that NASA became discouraged from further performing quantitative analyses of risk or reliability until after the Challenger mishap in 1986. Instead, NASA decided to rely on the Failure Modes and Effects Analysis (FMEA) method for system safety assessments. To date, FMEA continues to be required by NASA in all its safety-related projects.

²¹ Some authors (e.g. Apostolakis, 2004) instead refer to Quantitative Risk Assessment, to cover the possibility that the numerical values being manipulated are not strictly interpretable as probabilities.

“In the meantime, the nuclear industry picked up PRA to assess safety almost as a last resort in defense of its very existence. This analytical method was gradually improved and expanded by experts in the field and has gained momentum and credibility over the past two decades, not only in the nuclear industry, but also in other industries like petrochemical, offshore platforms, and defense. By the time the Challenger accident occurred, PRA had become a useful and respected tool for safety assessment. Because of its logical, systematic, and comprehensive approach, PRA has repeatedly proven capable of uncovering design and operation weaknesses that had escaped even some of the best deterministic safety and engineering experts. This methodology showed that it was very important to examine not only low-probability and high-consequence individual mishap events, but also high-consequence scenarios which can emerge as a result of occurrence of multiple high-probability and nearly benign events. Contrary to common perception, the latter is oftentimes more detrimental to safety than the former.” (Stamatelatos *et al.*, 2002a, p.1)

NASA’s series of public-domain guides on FTA, frequentist PRA, and Bayesian PRA (Stamatelatos *et al.*, 2002a; Stamatelatos *et al.*, 2002b; Dezfuli *et al.*, 2009, respectively) talk in terms of estimating and assuring *system* safety/reliability: they do not involve themselves in the distinction between systems, and systems-of-systems (SoS), which was informally introduced earlier. However, for the discussion that follows, we need to take a brief diversion into a more precise definition of what precisely we mean here by “SoS”.

A.2 Systems-of-Systems: Directed, Collaborative, Coalition, and Ecosystem

Probably the most-cited paper in the SoS literature is Maier’s “Architecting Principles for Systems of Systems” (1998), and we will use Maier’s careful definition of a SoS here. Maier proposed two primary characteristics that distinguish a SoS: a system that did not exhibit these two characteristics was, in his terms, not to be considered as a SoS “...*regardless* of the complexity or geographic distribution of its components.” (Maier 1998, p.271, original emphasis). Maier’s definition reads as follows:

“A system-of-systems is an assemblage of components which individually may be regarded as systems, and which possess two additional properties:

“Operational Independence of the Components: If the system-of-systems is disassembled into its component systems the component systems must be able to usefully operate independently. That is, the components fulfill customer-operator purposes on their own.

“Managerial Independence of the Components: The component systems not only *can* operate independently, they *do* operate independently. The component systems are separately acquired and integrated but maintain a continuing operational existence independent of the system-of-systems.”

(Maier, 1998, p.271, original emphasis)

A strict interpretation of Maier's definition of SoS would argue that the US Space Shuttle, even at one second before launch, is not a system of systems. The Orbiter, its external fuel tank, its left and right SRBs, and the launch-pad and support-tower that they all lift off from, do not have immediate *operational independence*: that is, they were all intimately designed to work with each other. It might perhaps be argued that with a little tinkering the SRBs could be re-engineered to usefully operate independently (as warhead-carrying long-range missiles, perhaps), but that would be clutching at straws: even if that were true, there is no real sense in which any of the Shuttle's component systems exhibit Maier's second property, of *managerial independence*, and on that basis the Shuttle at launch is simply not an SoS. At launch, each of the shuttle's component systems is under the collective, coordinated, combined command of NASA (the precise nexus of that command is something that is constructed by the interaction of, and shifts dynamically between, Mission Control on the ground, and the astronauts onboard the Shuttle).

Precisely because of Maier's definition, earlier in Section 2 of this paper we were careful not to describe the Shuttle as a SoS. Nevertheless, it is clear that the global financial markets network, or even "just" the financial markets operational in one of the major global hubs such as London or New York, satisfy both the operational independence and managerial independence criteria. Maier goes on to note that SoS can be classified as *Directed* (built and managed to fulfill specific purposes), or *Collaborative*, or *Virtual*. His definition of collaborative SoS reads as follows:

"Collaborative systems-of-systems are distinct from directed systems in that the central management organization does not have coercive power to run the system. The component systems must, more or less, voluntarily collaborate to fulfill the agreed upon central purposes."
(Maier, 1998, p.278).

In Maier's terms, a virtual SoS is then a SoS that is neither directed nor collaborative, i.e. it is one for which there is no central management authority, and also no agreed upon central purposes. Maier is explicit that he considers national economies to be virtual SoS; and it seems obvious that in Maier's terms the global financial markets are also virtual SoS. But classifying the markets as a virtual SoS simply because of their *absence* of central management and centrally agreed purpose glosses over some important richness in the network of interacting institutions within the financial markets. The markets involve varying numbers of various types of institution (e.g., investment banks, hedge funds, exchange operators, insurers, technology providers). The organizations that participate in the markets (and those that regulate them too) serve different purposes; some of them are in direct competition with other institutions (sometimes in zero-sum terms), others are in collaborative relationships with one or more other institutions; and such institutions come and go over time. Sommerville *et al.* (2012) have recently coined the term "Coalition of Systems" to describe this class of SoS; before that, Valerdi *et al.* (2008) referred to "No Single Owner SoS", and Northrop *et al.* (2006) coined the term *socio-technical ecosystems*, to capture the same notion that these SoS can be represented as a web of interacting constituents: in some cases the interactions are collaborative, in others they are competitive, all within the one SoS. It seems unarguable that the technology-enabled global financial markets of today, and in the future, are ecosystem-SoS.

The development of techniques for maintaining and managing high-integrity large-scale ecosystem-SoS is a new and significantly under-researched field. Fewer than five years ago, eight authors from industry and academia co-authored a paper (De Laurentis *et al.*, 2007) calling for an international consortium on SoS engineering to be established, to better understand the problems and solution strategies associated with SoS, yet their conception of a SoS was phrased in terms of "...heterogeneous independently operable systems to achieve a unique purpose" (p.68) – that is, they concentrated on a conception of SoS that is better suited to Maier's directed/collaborative SoS than the ecosystem-SoS of Northrop *et al.* Books and research papers exploring how to engineer robustly scalable socio-technical systems are currently few and far between (but see Abbot & Fisher, 2009; Rooksby, Rouncefield, & Sommerville, 2009; Baxter & Sommerville 2010).

The primary reason for that is because the development of reliable practices, and engineering teaching, for ensuring or assuring the integrity or safety of a SoS is a current research challenge; one that is being actively pursued by the world's leading research groups in high-integrity systems engineering, and even those leading researchers would admit that it is not yet a solved problem. In contrast to traditional engineering teaching, with its emphasis on designing "from scratch", starting (metaphorically at least) with a clean sheet of paper, most SoS instead arise from organic processes of aggregation and accretion, where pre-existing systems are integrated as constituents into the SoS. In almost all large-scale SoS, there is significant heterogeneity (which itself changes over time) because different constituents in the SoS were added at different stages in the development of the SoS and arrived via differing design and implementation paths. In their 2008 book *Eating the IT Elephant: Moving from Greenfield Development to Brownfield*, senior IBM staff Richard Hopkins and Kevin Jenkins made the analogy between the greenfield/brownfield distinction in civil engineering, and modern-day large-scale complex IT projects. A greenfield engineering project is one in which construction takes place on a previously undeveloped site, allowing a "clean-sheet" approach at the design stage, with relatively little preparatory work required on-site before construction, and with relatively few constraints on the construction process. A brownfield project is one in which the site has previously been built on and hence may require significant clearing operation before construction, with the possibility of the added complexity from the requirement that existing structures must be retained and their viability maintained during the construction phase (Hopkins & Jenkins, 2008).

Even if a large-scale SoS *was* the product of a clean-sheet engineering design process and was initially constructed from homogeneous constituents, sheer largeness-of-scale implies that at any one time it is almost definite that some of those constituents will have failed and be needing replacement (so-called *normal failure*). Those replacement constituents may not be exactly identical to the originals, and so the SoS becomes a heterogeneous, brownfield engineering problem.

The challenge of determining the safety of a SoS is neatly summarized by Alexander, Kazakov, & Kelly (2006):

"In a conventional system, ...the system boundary is well defined and the components within that boundary can be enumerated. When a safety analyst postulates some failure of a component, the effect of that failure can be propagated through the system to reveal whether or not the failure results in a

hazard. This is not always easy, because of the complexity of possible interactions and variability of system state, hence the need for systematic analysis techniques, automated analysis tools and system designs that minimize possible interactions. To make the task more tractable, most existing hazard analysis techniques... deal with only a single failure at a time; coincident failures are rarely considered.

“In an SoS, this problem is considerably worse. The system boundary is not well defined, and the set of entities within that boundary can vary over time, either as part of normal operations... or as part of evolutionary development... Conventional tactics to minimize interactions may be ineffective, because the system consists of component entities that are individually mobile. In some cases... the entities may be designed to form ad-hoc groupings amongst themselves. Conventional techniques may be inadequate for determining whether or not some failure in some entity is hazardous in the context of the SoS as a whole.”

The prospect of component entities being “individually mobile” was relevant to Alexander *et al.* because their work concentrated on SoS in defence applications, where the constituent entities in the SoS are often individual battlefield units (e.g., troops, tanks, unmanned vehicles, etc). While there is no direct *physical* correlate of spatial mobility in the computerized global financial markets, there is a reasonable equivalent in the *virtual* space defined by the network of current interactions between agents in the markets: just as a tank might physically move from one location to another on a battlefield in order to engage with the enemy or withdraw to a position of safety, so a trading agent (human or machine) might establish a connection with a potential counterparty, or terminate an existing connection. In both the tank battle and the trading scenario, the key factor that is altered is the network of links from the node in question (the tank, the trader), to other nodes in the network (enemy units, other traders) with which that node might have meaningful interactions (exchange of fire, exchange of bids/offers).

But this “mobility” issue of the network of meaningful interactions changing dynamically is not the only issue that confuses the task of understanding or managing an ecosystem SoS. Each of the nodes in the network, i.e. each of the constituent entities, is likely to be both *nonlinear* and *adaptive*. For the sake of the argument here, we’ll simply define “nonlinearity” as a meaning that the entity’s “outputs” (i.e., its responses or behavior) are not a simple linear function of its “inputs” (i.e., readings from its sensors, say); and we’ll adopt a similarly simple definition of “adaptive”: the entity is adaptive if its “outputs” may change over time, in consequence of the particular time-sequence of “inputs” that the entity is exposed to. Readers familiar with the mathematical economics literature will recognize this notion of adaptation as similar to “path-dependency”; colloquially we can think of the entity “learning from experience” or “evolving its response over time”. In recent decades, a new set of scientific tools and techniques has been developed to study systems composed of networks of interacting nonlinear adaptive entities. That field is known as Complexity Science, and the networked nonlinear adaptive systems are known as Complex Adaptive Systems.

A.3 Complex Adaptive Systems

In complexity science, *complex systems* are commonly defined as systems that are composed from large numbers of components, where each component interacts with some number of other components, and where there are nonlinearities in the nature of the component interactions and/or in the responses of the components themselves, which compound across the entire system in such a way that the overall system-level behaviour is difficult or perhaps impossible to predict accurately, even when one is given complete or near-complete information about the individual components and their interactions. The system-level behaviour is said to *emerge* from the network of interacting components and their constituent behaviours, forming a whole that is in some reasonable sense more than the sum of its parts. Substituting the word “constituent” for “component” in that description and it is clear that for very many SoS of practical importance, the SoS is manifestly a complex system. In addition to exhibiting emergent behaviour, many complex systems of significant interest are *adaptive* (in the sense informally introduced in the previous paragraph), and this also is surely true of many constituents in SoS, hence many SoS are instances of Complex Adaptive Systems (CAS). Since the late 1980’s a growing number of scientists have been attempting to understand the financial markets as CAS, and have been exploring the links between the financial markets and other CAS, both naturally-occurring and engineered artefacts. There is growing evidence that the emergent behaviour, phase changes, instabilities, and hysteresis seen in many other complex systems are also to be found in the financial markets: see, for example: Anderson, Arrow, & Pines (1989); Arthur, Morrison, *et al.* (1997); Johnson, Jefferies, & Hui (2003); Challet, Marsili, & Zhang (2004); and Blume & Durlaf (2005).

A small but growing number of researchers in the (systems-of-) systems engineering community have, in recent years, turned their attention to whether tools and techniques from complexity science can help in the brownfield engineering of robust, scalable, large-scale, systems: that is, they are exploring the consequences of taking a CAS approach to the creation and management of such large-scale systems and SoS: see, for example, Bar-Yam (2005); Braha *et al.* (2006); Sheard, & Mostashari (2008); Polacek, & Verma, (2009); and Sillitto (2010). Thus far, only a small amount of this work has addressed issues directly relevant to the financial markets but some notable work has been produced; see, e.g.: Harman & Bar-Yam, 2008; and the Nasdaq study by Darley & Oatkin (1997), which was discussed in more detail in Section 4.

Very often, such approaches involve exploring the system using so-called Multi-Agent Simulation (MAS) models, where a computer simultaneously models each of the constituents (or “agents”) in the network of interacting adaptive nonlinear entities, resolving the consequence of each entity’s interaction with its environment (which in most cases will include one or more other such entities), often using fine time-slicing or discrete-event simulation techniques. The agents in the simulation may adapt their responses over time either by implementing machine-learning techniques (for learning “within the lifetime” of the agent) and/or by implementing a process inspired by Darwinian evolution, a so-called *genetic algorithm* (a simulated population of agents, adapting to its niche over successive generations via a process of random variation and “survival of the fittest” directed selection: each agent’s behaviour or performance at the task at hand being determined at least in part by “genes” that can be passed on to successor agents: see e.g. Goldberg, 1987). Very often, the reliance on computer simulation models is a consequence of the mathematical nonlinearities in the system

being analytically intractable: that is, they are sufficiently complicated and complex that the tools for expressing them as a set of equations and then deriving formal proofs of certain statements about the system, via manipulation of the equations, is simply not possible.

For introductions to the use of CAS/MAS models in understanding social, economic, and socio-technical systems, see the texts by Epstein & Axtell (1996) and Axelrod & Cohen (2000). For examples of early machine-learning adaptive trading agents, see Cliff (1997) & Gjerstad & Dickhaut (1998), for the story of how those agents beat human traders, see Das *et al.* (2001). With regard to the application of evolutionary approaches, there has been heavy use of “replicator dynamics” (a technique pioneered in the theoretical study of evolution in biological systems) for exploring the interactions between different types of trading strategies, and identifying stable equilibria in the interaction dynamics (e.g., Walsh *et al.*, 2002; Vytelingum, Cliff, & Jennings 2008); and also various researchers have used genetic algorithms to create trading agents, and the market-mechanisms they operate in, co-adapted to each other by evolution (e.g., Phelps *et al.*, 2002; Cliff, 2003; Byde, 2003; Cliff, 2009; Phelps *et al.*, 2010). Evolutionary adaptation and co-adaptation in biological systems has served as a productive metaphor for economic dynamics at various levels for several decades (see, e.g., Nelson & Winter, 1982; Hodgson, 1993; Ormerod, 2006; Stephens & Waelbroeck, 2009); and there are other aspects of biological systems, such as the interconnected web of dependencies in natural ecosystems, that can offer fruitful insights into the functioning of financial systems (see, e.g., May *et al.*, 2008; Haldane & May, 2011; also Johnson, 2011). Sources of inspiration are not limited to biological systems: studies of the complex dynamics and size-vs-frequency distributions of earthquakes also offer insights for students of markets crashes: see Sornette (2002).

CAS and MAS approaches are not limited to the exploration of economic and financial systems: the approach is now pretty-much a standard item in the toolboxes of biologists, urban planners, military strategists, movie animators, safety architects, and practitioners of many more application areas in science and engineering. Several research teams have worked on developing general-purpose simulators (with associated visualization and analysis tools) for exploring CAS and MAS: for details of an example generic simulator and reviews of related work see Polack, Andrews, & Sampson (2009); and Polack *et al.* (2010).

In the course of this section’s discussion thus far, we’ve briefly surveyed high integrity systems engineering, and the definitions of systems of systems (SoS) and of complex adaptive system. Now we draw those three strands together and explore the current state, and future prospects for, high-integrity safety-critical engineering of complex adaptive ecosystem SoS.²²

²² We recognize that this is a long and cumbersome phrase. A shorter alternative might be “wicked systems”, first coined as a technical term in information systems engineering by Metcalf (2005) in direct reference to Rittel & Webber’s (1973) notion of “wicked problems”. But, given the current widespread disaffection in the media and general public with the banking sector, it seems prudent to avoid the potential confusion between the technical sense of “wicked” and the morally judgmental one, confusion that might arise in talking about trying to develop new engineering approaches for dealing with the “wicked systems of the financial markets”.

A.4 Engineering High-Integrity Complex Adaptive Ecosystem System-of-Systems

All approaches to risk assessment and safety-critical engineering involve the notion of a *model*. Rather than attempting to observe and manipulate the real physical system in its real operating environment, the model is instead an abstract representation of those aspects of the system that the engineers believe to be necessary and sufficient to reason about in order to achieve the task at hand. So, in this sense, a fault-tree diagram for some system is a model of that system. The fault-tree can be reasoned about, argued over, and altered to make it a better or worse representation of the real system, and the fault-tree can be manipulated to arrive at specific answers to specific questions, without having to touch the real system. The fault-tree is an explicit, diagrammatic, model of the system, suitable for risk assessment. But, as we have seen, the same system's risk assessment could instead be approached via Bayesian PRA, in which case the model will be a set of coupled equations and the associated prior probabilities.

In high integrity systems engineering, it is recognized that all models are developed iteratively, that they pass through a lifecycle: after an initial model is proposed, experience with the real system may reveal that the model needs refinement and improvement, the model is altered appropriately, but subsequent experience may again reveal the need for additional alterations. Eventually, it is hoped, the model will stabilize as more is known of the system. Of course, if the system itself is changing over time (as is almost definite in a socio-technical ecosystem SoS), the safety-engineer's model is forever playing catch-up; there will always be a strong likelihood that some aspect of the SoS is not yet known, not yet captured in the safety model.

Recognizing this, in recent years many researchers and practitioners involved in the engineering of high-integrity systems of systems have turned to predictive computer simulation models as a way of exploring "what if" scenarios. Such simulations are typically highly compute-intensive, and it is only with the ongoing Moore's-Law reductions in the real costs of computer power that such approaches have become practicable. In a predictive simulation, the model is expressed as interacting processes within the computer: such simulations may involve manipulating numeric values according to given equations (as in PRA); and they may also represent the model, or its outputs, via explicit diagrammatic visualizations (as in fault-tree analysis). Computer simulations offer the advantage of taking exhaustive "brute force" approaches to exploring system safety: for some systems, it is feasible to simulate the system in every possible combination of values for all variables of interest – the entire "state-space" of the system (that is, the space of all possible states it could ever find itself in) can be explored by the computer, given enough time. If the entire state-space is explored, and no unanticipated failures are discovered in the model, then (so long as the model is an accurate representation of the real system) the system's reliability is known completely. This technique of brute-force simulation has been particularly successful in the microelectronics industry, where the responses of new designs for silicon chips are explored exhaustively in simulation before the chip is fabricated for real: mistakes discovered at the simulation stage are *much* cheaper to fix than if the error is discovered only after the chip has been manufactured.

However, for many real-world systems, the state-space is sufficiently large that brute-force exhaustive searching is simply not possible. The combinatorics of state-spaces often involve exponentials-of-exponentials: equations of the form $v=w$ -to-the-power-(x -to-the-power-(y -to-the-power- z))), and numbers such as v can grow astronomically huge, much larger than the number of atoms in the known universe, for only moderate values of w , x , y , and z . Attempting exhaustive search of such vast state-spaces is possible in theory, but the sun will burn out long before the search is over. So, for many real systems, sophisticated techniques are required to cleverly sample only selected points or areas in the system's state-space. Developing such techniques is a current research issue, even in microelectronics where the state-spaces of current chips have now grown to routinely be beyond the size where exhaustive search is practicable (see, e.g. Hsueh & Eder, 2006).

Researchers concerned with risk assessment and safety assurance in SoS have developed increasingly sophisticated simulation modelling techniques (see, e.g., De Laurentis & Han, 2006; Parisi *et al.*, 2008; Clymer, 2009; Kewley & Tolk, 2009), and researchers interested in developing generic simulation tools for the study of complex adaptive systems have learnt from the methods developed in high-integrity systems engineering (Polack, Andrews, & Sampson, 2009). Some recent work has explored the possibility of feeding the outputs of simulation models directly into machine learning (ML) algorithms, so that the ML system can discover or learn rules and regularities that can neatly summarise the behavior of the system (see, e.g., Eder, Flach, & Hsueh, 2006; Alexander, 2007). Nevertheless, researchers remain cautiously aware that the model is only that: only a model, an abstraction. The models are used to explore possible circumstances and situations that may be very rare, and/or disastrous, in the real system. Alexander *et al.* (2006) comment that this approach is one that Dewar *et al.* (1996) refers to as "weak prediction":

"[Dewar *et al.*, 1996] note that "*subjective judgement is unavoidable in assessing credibility*" and that when such a simulation produces an unexpected result "it has created an interesting hypothesis that can (and must) be tested by other means". In other words, when a simulation reveals a plausible system hazard, other, more conventional analyses must be carried out to determine whether it is credible in the real system. Therefore, the role of the simulation analysis is to narrow down a huge analysis space into one that is manually tractable." (Alexander *et al.*, 2006)

One of the biggest challenges at present concerns modeling the *social* elements in socio-technical SoS: people and groups of people can be surprisingly sophisticated (and surprisingly stupid), and representing their relevant nonlinear, adaptive, nondeterministic behavior in a simulation model is certainly not easy.

Although it is undoubtedly difficult to capture human ingenuity and adaptivity, there are well-developed techniques in the CAS literature that can serve as good proxies: most notable of these is the use of co-evolution as a process for driving stochastic search through a space of possible designs or strategies, giving rise to what can appear to be a form of "artificial creativity".

The seminal example of this approach was described in a paper by Hillis (1990): Hillis used simulated evolution, a genetic algorithm (GA), to automatically design algorithms for sorting lists of numbers into numeric order; each "individual" in his

GA's population was a particular algorithm, and the sequence of steps in each individual's algorithm were specified by its "genes" (each step involved comparing a pair of numbers, and if necessary swapping their places in the list to make them be in the right numeric order); each individual's probability of reproduction (i.e., its *fitness*) was determined by how many test-lists it sorted successfully.

Initially, Hillis worked with a set-up where the test-lists were fixed in advance: when he did this, his GA could reliably evolve individual algorithms that did well at sorting the specific lists in the test set, but did poorly when presented with a novel list, one that was not in the test set. To counteract this, Hillis re-worked his system so that the test-lists were *also* an evolving population: the test-set was a population of lists, the particular numbers in each list were specified via its "genes" and the "fitness" of each list was determined by how "difficult" it was, i.e., by how many of the sorting algorithms failed to sort it. Thus the population of sorting algorithms, and the population of test-lists, made up a competitive co-evolutionary system, much like a predator-prey or parasite-host dynamic: the fitness of each sorter-algorithm depended on how many lists it could sort; the fitness of each list depended on how many sorter-algorithms it could defeat; and the two populations co-evolved over time. The co-evolutionary system was much more productive, and readily discovered sorting algorithms that rivalled the best-known human-designed ones.

Since Hillis' paper, several CAS researchers have demonstrated the power of co-evolution as a force for generating novel solutions and designs (see, e.g. Sims, 1994; Funes & Pollack 1999; Cartledge & Bullock, 2004; Cliff & Miller 2006; Stuermer *et al.* 2009), it seems entirely plausible that co-evolutionary processes could be used to approximate the effects of human ingenuity and creativity in socio-technical systems. Perhaps more importantly, co-evolutionary processes could also be used to explore the state-space of simulated ecosystems SoS, in the search for conditions that reveal unanticipated failure modes, in much the same way as Hillis's population of test-lists searched for methods of "failing" his population of sorting algorithms. This would allow semi-automated generation of hypotheses about how the real system might fail.

References

1. Abbot, M., Fisher, M.: *The Art of Scalability: Scalable Web Architecture, Processes, and Organizations for the Modern Enterprise*. Addison-Wesley (2009)
2. Alexander, R., Kazakov, D., Kelly, T.: *System of Systems Hazard Analysis Using Simulation and Machine Learning*. In: Górski, J. (ed.) *SAFECOMP 2006*. LNCS, vol. 4166, pp. 1–14. Springer, Heidelberg (2006)
3. Alexander, R.: *Using Simulation for Systems of Systems Hazard Analysis*. PhD Thesis, Department of Computer Science, University of York, UK (2007)
4. Anderson, P., Arrow, K., Pines, D. (eds.): *The Economy as an Evolving Complex System*. Addison-Wesley (1989)
5. Angel, J.: *Opening Remarks at SEC Roundtable on Shortselling (May 5, 2009a)*, <http://www.sec.gov/comments/4-581/4581-2.pdf>
6. Angel, J.: *Letter to the Securities and Exchange Commission (June 19, 2009b)*, <http://www.sec.gov/comments/s7-08-09/s70809-3758.pdf>
7. Angel, J.: *Letter to the Securities and Exchange Commission (September 21, 2009c)*, <http://www.sec.gov/comments/s7-08-09/s70809-4658.pdf>
8. Angel, J., Harris, L., Spratt, C.: *Trading in the 21st Century (2010)* (unpublished manuscript), <http://www.sec.gov/comments/s7-02-10/s70210-54.pdf>

9. Angel, J.: Letter to the Securities and Exchange Commission (April 30, 2010a), <http://www.sec.gov/comments/s7-02-10/s70210-172.pdf>
10. Angel, J.: Testimony to the US Senate (December 8, 2010b), <http://msb.georgetown.edu/story/1242666871500.html>
11. Apostolakis, G.: How Useful is Quantitative Risk Analysis? *Risk Analysis* 24(3), 515–520 (2004)
12. Arthur, B.: *The Nature of Technology: What it is and how it evolves*. Allen Lane (2009)
13. Arthur, B., Morrison, V., Durlauf, S., Lane, D. (eds.): *The Economy as an Evolving Complex System II*. Addison Wesley (1997)
14. Axelrod, R., Cohen, M.: *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. Free Press (2000)
15. Bar-Yam, Y.: *Making Things Work: Solving Complex Problems in a Complex World*. Knowledge Press (2005)
16. Baxter, G., Sommerville, I.: Socio-technical Systems: From design methods to systems engineering. *Interacting with Computers* 23(1), 4–17 (2010)
17. Benveniste, A., Fabre, E., Haar, S.: Markov Nets: Probabilistic Models for Distributed and Concurrent Systems. *IEEE Transactions on Automatic Control* 48(11), 1936–1950 (2003)
18. Beinhocker, E.: *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics*. Harvard Business School Press (2007)
19. Blas, J.: High-speed trading blamed for sugar rises. *The Financial Times* (February 8, 2011), <http://www.ft.com/cms/s/0/05ba0b60-33d8-11e0-b1ed-00144feabdc0.html#axzz1JIX0tWXX>
20. Blumberg, B.: *Old Tricks, New Dogs: Ethology and Interactive Creatures*. Ph.D. Thesis, MIT Media Lab. (1996)
21. Blume, L., Durlaf, S.: *The Economy as an Evolving Complex System, III*. Addison-Wesley (2005)
22. Bonen, Z.: *Evolutionary Behavior of Complex Socio-Technical Systems*. Working Paper #1056-79. Alfred P. Sloan School of Management, MIT (1979)
23. Bonisch, P., Di Giammarino, P.J.: *Achieving Supervisory Control of Systemic Risk*. Report jointly produced by UK Financial Services Knowledge Transfer Network, JWG, and Paradigm Risk (October 2010), <http://www.jwg-it.eu/library.php?typeId=11>
24. Bootle, R.: *The Trouble with Markets: Saving Capitalism from Itself*. Nicholas Brealey Publishing (2009)
25. Braha, D., Minai, A., Bar-Yam, Y.: *Complex Engineered Systems: Science Meets Technology*. Springer (2006)
26. Bullock, S.: *Prospects for Large-Scale Financial Systems Simulation*. Driver Review DR14, Foresight Project on the Future of Computer Trading in the Financial Markets. UK Government Office for Science (2011), <http://www.bis.gov.uk/assets/foresight/docs/computer-trading/11-1233-dr14-prospects-for-large-scale-financial-systems-simulation.pdf>
27. Bye, A.: *Applying Evolutionary Game Theory to Auction Mechanism Design*. In: *Proceedings of the 2003 ACM Conference on E-Commerce*, pp. 192–193 (2003)
28. Calinescu, R., Kwiatkowska, M.: *Software Engineering Techniques for the Development of Systems of Systems*. In: *Choppy, C., Sokolsky, O. (eds.) Monterey Workshop 2008*. LNCS, vol. 6028, pp. 59–82. Springer, Heidelberg (2010)
29. Calinescu, R., Kikuchi, S., Kwiatkowska, M.: *Formal Methods for the Development and Verification of Autonomic IT Systems*. In: *Cong-Vinh, P. (ed.) Formal and Practical Aspects of Autonomic Computing and Networking: Specification, Development and Verification*. IGI Global (2010) (to appear)
30. Cartlidge, J., Bullock, S.: *Combating Coevolutionary Disengagement by Reducing Parasite Virulence*. *Evolutionary Computation* 12(2), 193–222 (2004)

31. CFTC & SEC. 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, 2010a), <http://www.sec.gov/sec-cftc-prelimreport.pdf>
32. CFTC & SEC. 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 (September 30, 2010b), <http://www.sec.gov/news/studies/2010/marketevents-report.pdf>
33. CFTC & SEC. Recommendations Regarding Regulatory Responses to the Market Events of May 6, 2010. Summary Report of the Joint CFTC and SEC Advisory Committee on Emerging Regulatory issues (February 18, 2011)
34. Challet, D., Marsili, M., Zhang, Y. (eds.): *Minority Games: Interacting agents in financial markets*. OUP (2004)
35. Chilton, B.: Stopping Stammering: Overcoming Obstacles in Financial Regulatory Reform. In: Speech of Commissioner Bart Chilton to the Goldman Sachs Global Commodity Conference, London (2011), <http://www.cftc.gov/pressroom/speechestestimony/opachilton-43.html>
36. Cliff, D., Keen, J., Kwiatkowska, M., McDermid, J., Sommerville, I.: Large Scale Complex IT Systems (LSCITS) Research Programme. Research proposal to the UK Engineering and Physical Sciences Research Council (2006), <http://lscits.cs.bris.ac.uk/docs/LSCITSproposalRP1.pdf> (submitted December 2006, commenced April 2007)
37. Cliff, D.: Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, Hewlett Packard Labs (1997)
38. Cliff, D., Bruten, J.: Animat Market-Trading Interactions as Collective Social Adaptive Behavior. *Adaptive Behavior* 7(3&4), 385–414 (1999)
39. Cliff, D.: Explorations in evolutionary design of online auction market mechanisms. *Journal of Electronic Commerce Research and Applications* 2(2), 162–175 (2003)
40. Cliff, D., Miller, G.: Visualising Coevolution with CIAO plots. *Artificial Life* 12(2), 199–202 (2006)
41. Cliff, D.: ZIP60: Further Explorations in the Evolutionary Design of Trader Agents and Online Auction-Market Mechanisms. *IEEE Transactions on Evolutionary Computation* 13(1), 3–18 (2009)
42. Cliff, D.: Networked Governance in the Financial Markets. Foresight strategic briefing paper, for UK Government Office of Science & Technology, Department of Business, Innovation, and Skills (2010), http://www.cs.bris.ac.uk/home/dc/Foresight_NetGov_v2a.pdf
43. Cliff, D., Brown, D., Treleaven, P.: Technology Trends in the Financial Markets: A 2020 Vision. Driver Review DR3, Foresight Project on the Future of Computer Trading in the Financial Markets. UK Government Office for Science (2011), <http://www.bis.gov.uk/assets/foresight/docs/computer-trading/11-1222-dr3-technology-trends-in-financial-markets.pdf>
44. Clymer, J.: *Simulation-Based Engineering of Complex Systems*, 2nd edn. Wiley-Blackwell (2009)
45. Colander, D., Föllmer, H., Haas, A., Goldberg, M., Juselius, K., Kirman, A., Lux, T., Sloth, B.: The Financial Crisis and the Systemic Failure of Academic Economics. Kiel Working Paper 1489, Kiel Institute for the World Economy (2009)
46. Collingridge, D.: *The Management of Scale: Big Organizations, Big Decisions, Big Mistakes*. Routledge (1992)
47. Darley, V., Outkin, A.: *A NASDAQ Market Simulation: Insights on a Major Market from the Science of Complex Adaptive Systems*. World Scientific (2007)

48. Das, R., Hanson, J., Kephart, J., Tesauro, G.: Agent-Human Interactions in the Continuous Double Auction. In: Proceedings IJCAI 2001 (2001)
49. De Laurentis, D., Dickerson, C., DiMario, M., Gartz, P., Jamshidi, M., Nahavandi, S., Sage, A., Sloane, E., Walker, D.: A Case for an International Consortium on System-of-Systems Engineering. *IEEE Systems Journal* 1(1), 68–73 (2007)
50. De Laurentis, D., Han, E.: System-of-Systems Simulation for Analyzing the Evolution of Air Transportation. In: Proceedings of the 25th International Congress of the Aeronautical Sciences, pp. 1–10 (2006)
51. Demos, T.: US panel on flash crash urges rule changes. *The Financial Times* (February 18, 2011a), <http://www.ft.com/cms/s/0/417134ea-3b84-11e0-9970-00144feabdc0.html#axzz1EOx4E4Gg>
52. Demos, T.: Quick View: Blown away by the flash crash report. *The Financial Times* (February 19, 2011b), <http://www.ft.com/cms/s/0/bf6017b0-3baa-11e0-a96d-00144feabdc0.html#axzz1EOx4E4Gg>
53. Demos, T.: Plans to avert “flash crash” draw opposition. *The Financial Times* (March 22, 2011c), <http://www.ft.com/cms/s/0/3a3e52a0-54a9-11e0-b1ed-00144feab49a.html#axzz1Ht6fUrUu>
54. Despotou, G.: Managing the Evolution of Dependability Cases for Systems of Systems. PhD Thesis, Department of Computer Science, University of York, UK (2007)
55. Dewar, J., Bankes, S., Hodges, J., Lucas, T., Saunders-Newton, D., Vye, P.: Credible Uses of the Distributed Interactive Simulation (DIS) System. Technical Report MR-607-A, RAND (1996)
56. Dezfuli, H., et al.: Bayesian Inference for NASA Probabilistic Risk and Reliability Analysis. NASA SP-2009-569 (2009), <http://www.hq.nasa.gov/office/codeq/doctree/SP2009569.pdf>
57. Dorner, D.: The logic of failure. *Philosophical Transactions of the Royal Society of London, Series B* 327(1241), 463–473 (1990)
58. Dorner, D.: *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*. Perseus (1997)
59. Easley, D., Lopez de Prado, M., O’Hara, M.: The Microstructure of the Flash Crash: Flow Toxicity, Liquidity Crashes and the Probability of Informed Trading. *The Journal of Portfolio Management* 37(2), 118–128 (2011)
60. Economist. Agents of Change. *The Economist* 396(8692), 76 (2010) (Note that The Economist has a standard policy of not showing author bylines for articles written by regular staff journalists)
61. Eder, K., Flach, P.A., Hsueh, H.-W.: Towards Automating Simulation-Based Design Verification Using ILP. In: Muggleton, S., Otero, R., Tamaddoni-Nezhad, A. (eds.) *ILP 2006*. LNCS (LNAI), vol. 4455, pp. 154–168. Springer, Heidelberg (2007)
62. Epstein, J., Axtell, R.: *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press (1996)
63. Epstein, J.: *Generative Social Science: Studies in Agent-Based Computational Modelling*. Princeton University Press (2007)
64. Farmer, D., Foley, D.: The economy needs agent-based modeling. *Nature* 460, 685–686 (2009)
65. Farmer, D., Skouras, S.: An ecological perspective on the future of computer trading. Driver Review DR6, Foresight Project on the Future of Computer Trading in the Financial Markets. UK Government Office for Science (2011), <http://www.bis.gov.uk/assets/foresight/docs/computer-trading/11-1225-dr6-ecological-perspective-on-future-of-computer-trading.pdf>
66. Flinders, K.: The Evolution of Stock Market Technology. *Computer Weekly* (November 2, 2007), <http://www.computerweekly.com/Articles/2007/11/02/227883/The-evolution-of-stock-market-technology.htm>

67. Funes, P., Pollack, J.: Computer Evolution of Buildable Objects. In: Bentley, P. (ed.) *Evolutionary Design by Computers*, ch. 17. Morgan Kaufman (1999)
68. Galas, M., Brown, D., Treleaven, P.: ATRADE Platform: Algorithmic Trading & Risk Analytics Development Environment. Unpublished manuscript, Department of Computer Science, University College London (2010), <http://fc.cs.ucl.ac.uk/mscfc/virtual-trading-floor>
69. Gjerstad, S., Dickhaut, J.: Price Formation in Double Auctions. *Games and Economic Behavior* 22, 1–29 (1998)
70. Grant, J.: Quick View: US looks at European-style circuit breakers. *The Financial Times* (May 19, 2010), <http://cache.ft.com/cms/s/0/139ddd44-6325-11df-99a5-00144feab49a.html#axzz1Ij5j9zds>
71. Goldberg, D.: *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley (1987)
72. Gorton, G.: *Slapped by the Invisible Hand: The Panic of 2007*. OUP (2010)
73. Goth, G.: Ultralarge Systems: Redefining Software Engineering? *IEEE Software* 25(3), 91 (2008)
74. Gray, J.: On eScience: A Transformed Scientific Method. In: Hey, T., Tansley, S., Tolle, K. (eds.) *The Fourth Paradigm: Data-Intensive Scientific Discovery*, pp. xvii–xxxi. Microsoft Press (2009)
75. Haccou, P., Meelis, E.: *Statistical Analysis of Behavioral Data: An Approach Based on Time-Structured Models*. Oxford University Press (1994)
76. Haldane, A.: Rethinking the Financial Network. Text of a speech given at the Financial Student Association, Amsterdam (April 2009), <http://www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf>
77. Haldane, A., May, R.: Systemic risk in banking ecosystems. *Nature* 469, 351–355 (2011)
78. Hall-May, M.: *Ensuring Safety of Systems of Systems*. PhD Thesis, Department of Computer Science, University of York, UK (2007)
79. Harford, T.: What we can learn from a nuclear reactor? *The Financial Times* (January 14, 2011), <http://www.ft.com/cms/s/2/cea7b256-1def-11e0-badd-00144feab49a.html#axzz1DN62IXnB>
80. Harman, D., Bar-Yam, Y.: Technical Report on SEC Uptick Repeal Pilot. NECSI Technical Report 2008-11, New England Complex Systems Initiative (2008)
81. Hillis, D.: Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D* 42, 228–234 (1990)
82. Hodgson, G.: *Economics and Evolution: Bringing life back into economics*. Polity Press (1993)
83. Hollnagel, E., Woods, D., Leveson, N. (eds.): *Resilience Engineering: Concepts and Precepts*. Ashcroft (2006)
84. Hopkins, R., Jenkins, K.: *Eating the IT Elephant: Moving from Greenfield Development to Brownfield*. IBM Press (2008)
85. Horswill, I.: Very Fast Action Selection for Parameterized Behaviors. In: *Proceedings of the Fifth International Conference on Foundations of Digital Games (FDG 2009)*, Orlando (2009)
86. Hsueh, H., Eder, K.: Test Directive Generation for Functional Coverage Closure Using Inductive Logic Programming. In: *Proc. IEEE International High Level Design Validation and Test Workshop (HLDVT)*, pp. 11–18 (2006)
87. Hubbard, D.: *The Failure of Risk Management. Why It's Broken and How to Fix It*. John Wiley (2009)
88. Institute for International Finance. *Interim Report of the IIF Committee on Market Best Practices* (April 2008), <http://www.iif.com/download.php?id=SDzcEc8juCI=:/>

89. Ivanov, Y.: State Discovery for Autonomous Learning. Ph.D. Thesis, MIT Media Lab. (2002)
90. Johnson, C.: Lessons from the Evacuation of the World Trade Center, September 11, 2001, for the Future Development of Computer Simulations. *Cognition, Technology, & Work* 7, 214–240 (2005)
91. Johnson, C.: Using Evacuation Simulations to Ensure the Safety and Security of the 2012 Olympic Venues. *Safety Science* 46(2), 302–322 (2008)
92. Johnson, C., Nilsen-Nygaard, L.: Extending the Use of Evacuation Simulators to Support Counter-Terrorism: Using Models of Human Behaviour to Coordinate Emergency Responses to Improvised Explosive Devices. In: Simmons, R., Mohan, D., Mullane, M. (eds.) *Proceedings of the 26th International Conference on Systems Safety* (2008)
93. Johnson, C.: The Application of Computational Models for the Simulation of Large-Scale Evacuations Following Infrastructure Failures and Terrorist Incidents. In: *Proceedings of NATO Research Workshop on Computational Models of Risk to Infrastructure*, May 9–13. NATO (2006)
94. Johnson, N., Jefferies, P., Hui, P. (eds.): *Financial Market Complexity*. OUP (2003)
95. Johnson, N.: Proposing Policy by Analogy is Risky. *Nature* 469, 302 (2011)
96. Kearns, M., Ortiz, L.: The Penn-Lehman Automated Trading Project. *IEEE Intelligent Systems*, 22–31 (November/December 2003)
97. Kewley, R., Tolk, A.: A Systems Engineering Process for Development of Federated Simulations. In: *SpringSim 2009: Proceedings of the 2009 Spring Simulation Multiconference*. Society for Computer Simulation International (2009)
98. Kindleberger, C.: *Manias, Panics, and Crises: A History of Financial Crises*. John Wiley (2001)
99. LeBaron, B.: Agent Based Computational Finance: Suggested Readings and Early Research. *Journal of Economic Dynamics and Control* 24, 679–702 (2000)
100. LeBaron, B., Arthur, B., Palmer, R.: The Time Series Properties of an Artificial Stock Market. *Journal of Economic Dynamics and Control* 23, 1487–1516 (1999)
101. Levy, M., Levy, H., Solomon, S.: *Microscopic Simulation of the Financial Markets: From Investor Behavior to Market Phenomena*. Academic Press (2000)
102. Lewis, M.: *The Big Short: Inside the Doomsday Machine*. Allen Lane (2010)
103. Lorenz, E.: Deterministic Nonperiodic Flow. *Journal of Atmospheric Science* 20, 130–141 (1963)
104. Lorenz, K.: *On Aggression*. Routledge Classics (1966/2002)
105. MacKenzie, D.: *An Engine, Not a Camera: How Financial Models Shape Markets*. MIT Press (2008a)
106. MacKenzie, D.: *Material Markets: How Economic Agents are Constructed*. Oxford University Press (2008b)
107. MacKenzie, D., et al. (eds.): *Do Economists Make Markets? On the Performativity of Economics*. Princeton University Press (2008)
108. Maier, M.: Architecting Principles for Systems of Systems. *Systems Engineering* 1(4), 267–284 (1998)
109. May, R., Levin, S., Sugihara, G.: Ecology for Bankers. *Nature* 451, 893–895 (2008)
110. Meerman, M., et al.: *Money and Speed: Inside the Black Box*. Documentary produced by VPRO (Dutch public broadcaster) (2011), available as an iPad application <http://itunes.apple.com/us/app/money-speed-inside-black-box/id424796908?mt=8&ls=1#>
111. Metcalfe, M.: Strategic knowledge sharing: a small-worlds perspective. In: Hart, D., Gregor, S. (eds.) *Information System Foundations: Constructing and Criticizing*, Australian National University Press (2005)
112. Mitchell, M.: *Complexity: A Guided Tour*. OUP (2009)
113. Mullane, R.: *Riding Rockets: The Outrageous Tales of a Space-Shuttle Astronaut*. Simon & Schuster (2006)

114. Nelson, R., Winter, S.: *An Evolutionary Theory of Economic Change*. Harvard University Press (1982)
115. Northrop, L., et al.: *Ultra-Large-Scale Systems: The Software Challenge of the Future*. Technical Report. Carnegie Mellon University Software Engineering Institute (2006)
116. Ormerod, P.: *Butterfly Economics: A New General Theory of Economic and Social Behaviour*. Faber (1998)
117. Ormerod, P.: *Why Most Things Fail: Evolution, Extinction, & Economics*. Faber (2006)
118. Palmer, R., Arthur, B., Holland, J., LeBaron, B., Tayler, P.: Artificial economic life: a simple model of a stockmarket. *Physica D: Nonlinear Phenomena* 75(1-3), 264–274 (1994)
119. Parry, G.: The Characterization of Uncertainty in Probabilistic Risk Assessments of Complex Systems. *Reliability Engineering and System Safety* 54, 119–126 (1996)
120. Parisi, C., Sahin, F., Jamshidi, M.: A discrete event XML based system of systems simulation for robust threat detection and integration. In: *Proc. 2008 IEEE International Conference on System of Systems Engineering* (2008)
121. Perrow, C.: *Normal Accidents: Living with High-Risk Technologies*. Basic Books, New York (1984)
122. Phelps, S., McBurney, P., Parsons, S., Sklar, E.: Co-evolutionary Auction Mechanism Design: A Preliminary Report. In: Padget, J., Shehory, O., Parkes, D., Sadeh, N., Walsh, W.E. (eds.) *AMEC 2002*. LNCS (LNAI), vol. 2531, pp. 123–142. Springer, Heidelberg (2002)
123. Phelps, S., McBurney, P., Parsons, S.: Evolutionary Mechanism Design: A Review. *Autonomous Agents and Multi-Agent Systems* 21(2), 237–264 (2010)
124. Polack, F., Andrews, P., Sampson, A.: The Engineering of Concurrent Simulations of Complex Systems. In: *Proc. 2009 IEEE Congress on Evolutionary Computation*, pp. 217–224 (2009)
125. Polack, F., Andrews, P., Ghetiu, T., Read, M., Stepney, S., Timmis, J., Sampson, A.: Reflections on the Simulation of Complex Systems for Science. In: *Proc. International Conference on Engineering of Complex Computer Systems (ICECCS 2010)*, pp. 276–285. IEEE Press (2010)
126. Polacek, G., Verma, D.: Requirements Engineering for Complex Systems: Principles vs. Rules. In: *Proceedings of the Seventh Annual Conference on Systems Engineering Research, CSER 2009* (2009)
127. Reason, J.: *The Human Contribution: Unsafe Acts, Accidents, and Heroic Recoveries*. Ashgate (2008)
128. Rittel, H., Webber, M.: Dilemmas in a General Theory of Planning. *Policy Sciences* 4, 155–169 (1973)
129. Roberts, K.: Some Characteristics of One Type of High Reliability Organization. *Organization Science* 1(2), 160–176 (1990)
130. Rooksby, J., Rouncefield, M., Sommerville, I.: Testing in the Wild: The Social and Organisational Dimensions of Real-World Practice. *Journal of Computer Supported Cooperative Work* 18(5-6), 559–580 (2009)
131. Salmon, F.: Algorithmic trading and market-structure tail risks. Reuters (January 13, 2011), <http://blogs.reuters.com/felix-salmon/2011/01/13/algorithmic-trading-and-market-structure-tail-risks/>
132. Schelling, T.: Dynamic Models of Segregation. *Journal of Mathematical Sociology* 1, 143–186 (1971)
133. Sheard, S., Mostashari, A.: Principles of Complex Systems for Systems Engineering. *Systems Engineering* 12(4), 295–311 (2008)
134. Sillitto, H.: *Design Principles for Ultra-Large-Scale Systems* (2010) (unpublished draft manuscript)
135. Sims, K.: Evolving 3D Morphology and Behavior by Competition. In: Brooks, R., Maes, P. (eds.) *Artificial Life IV Proceedings*, pp. 28–39. MIT Press (1994)
136. Sloan, S.: *Simulating Terrorism*. University of Oklahoma Press (1981)

137. Slovik, P.: Trust, Emotion, Sex, Politics, and Science: Surveying the Risk-Assessment Battlefield. *Risk Analysis* 19(4), 689–701 (1999)
138. Smith, L.: Accountability and Error in Ensemble Forecasting (1995) (manuscript), <http://people.maths.ox.ac.uk/lenny/ecmwf96.pdf>
139. Smith, L.: What might we learn from climate forecasts? *Proceedings of the National Academy of Sciences of the United States of America* 99, 2487–2492 (2002)
140. Stamatelatos, M., et al.: Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners. V1.1 (2002a), <http://www.hq.nasa.gov/office/codeq/doctree/praguide.pdf>
141. Stamatelatos, M., et al.: Fault Tree Handbook with Aerospace Applications. Version 1.1. (2002b), <http://www.hq.nasa.gov/office/codeq/doctree/fthb.pdf>
142. Sommerville, I., Cliff, D., Calinescu, R., Keen, J., Kelly, T., Kwiatkowska, M., McDermid, J., Paige, R.: Large-Scale Complex IT Systems. *Communications of the ACM* (July 2012)
143. Sornette, D.: *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton University Press (2002)
144. Steiner, C.: Did We Dodge Another Flash Crash on September 1? (2010) Forbes blog, <http://blogs.forbes.com/christophersteiner/2010/09/02/did-we-just-dodge-another-flash-crash-yesterday/>
145. Stephens, C., Waelbroeck, H.: Algorithm Switching: Co-Adaptation in the Market Ecology. *Journal of Trading*, 1–15 (Summer 2009)
146. Stuermer, P., Bucci, A., Branke, J., Funes, P., Popovici, E.: Analysis of coevolution for worst-case optimization. In: *Proceedings GECCO 2009, the 11th Annual Conference on Genetic and Evolutionary Computation* (2009)
147. Taleb, N.: *The Black Swan: The Impact of the Highly Improbable*. Allen Lane (2007)
148. Tesfatsion, L., Judd, K. (eds.): *The Handbook of Computational Economics. Agent-Based Computational Economics, vol. 2*. North-Holland (2006)
149. Tett, G.: Fool's Gold: How Unrestrained Greed Corrupted a Dream, Shattered Global Markets, and Unleashed a Catastrophe. Little, Brown (2009)
150. Tomlinson, B., Blumberg, B.: Social Synthetic Characters. *Computer Graphics* 26(2) (2002)
151. Tuinenga, P.: *SPICE: A Guide to Circuit Simulation and Analysis Using PSpice*. Prentice Hall (1988)
152. Turtle, S.: *Simulation and its Discontents*. MIT Press (2009)
153. Valerdi, R., Axelband, E., Baehren, T., Boehm, B., Dorenbos, D., Jackson, S., Madni, A., Nadler, G., Robitaille, P., Settles, S.: A research agenda for systems of systems architecting. *International Journal of System of Systems Engineering* 1(1&2), 171–188 (2008)
154. Vaughan, D.: *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. University of Chicago Press (1997)
155. Vaughan, D.: On Slippery Slopes, Repeating Negative Patterns, and Learning from Mistakes. In: Starbuck, W., Farjoun, M. (eds.) *Organization at the Limit: Lessons from the Columbia Disaster*, pp. 262–275. Wiley Blackwell (2005)
156. Vaughan, D.: NASA Revisited: Theory, Analogy and Public Sociology. *American Journal of Sociology* 112(2), 353–393 (2006)
157. Vytelingum, K., Cliff, D., Jennings, N.: Strategic bidding in continuous double auctions. *Artificial Intelligence* 172(13), 1700–1729 (2008)
158. Waldrop, M.: *Complexity: The Emerging Science at the Edge of Order and Chaos*. Simon & Schuster (1992)
159. Walsh, W., Das, R., Tesauro, G., Kephart, J.: Analyzing Complex Strategic Interactions in Multi-Agent Games. In: *Proceedings of AAAI 2002 Game Theoretic and Decision Theoretic Agents Workshop*, Edmonton, Canada (2002)
160. Weick, K., Sutcliffe, K.: *Managing the Unexpected*. Jossey Bass (2007)