

The Flash Crash of May 6th 2010: WTF?

Dave Cliff

Director, UK Large-Scale Complex IT Systems Research & Training Initiative
www.lscits.org

Department of Computer Science, University of Bristol, Bristol BS8 1UB, UK.
+44 79 77 55 22 50; dc@cs.bris.ac.uk

Copyright © Dave Cliff, November, 2010.

(This version with minor edits and an extended Section 4, December 2010)

Abstract

In most advanced economies, large-scale socio-technical systems are increasing in complexity, and in socio-economic criticality, while our ability to engineer and manage such systems is probably not increasing at the same pace. The “Flash Crash” of May 6th 2010 is evidence of what may become more commonplace in future, and is evidence to support the concern that in advanced economies we may already be reliant on large-scale complex IT systems that support critical social and economic functions, and yet for which we cannot predict their failures until it is too late. Nevertheless, there is an established literature devoted to the study of technology failures, and that literature may be of some help: in the financial markets, practitioners and regulators should now be asking the question “Why Technology Failures?” (WTF?). I also argue here 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 potentially identifying “black swan” failure modes in the simulations, before they occur in real life, by which time it is typically too late.

Provenance

This whitepaper is a close approximation to a transcript of the words I’ve spoken in a presentation first given at a workshop at the Royal Society of London, organized by the *Foresight* unit of the UK Government’s Office of Science & Technology, in July 2010. An extended version of the same presentation was given at the TTI Vanguard *Matters of Scale* conference later the same month, and at the GII Doctoral Summer School on Ultra-Large-Scale Systems in Bari, Italy, in September 2010. After that, I gave a revised and extended version at the opening of the inaugural annual conference of the UK Financial Services Knowledge Transfer Network (FS-Net) in October 2010. A revision of the *Matters of Scale* version then formed the basis for the opening presentation at the first public symposium of the UK Large-Scale Complex IT Systems (LSCITS) National Research & Training Initiative, in November 2010 (for further details of how this relates to the LSCITS Initiative, see Cliff *et al.*, 2006; Calinescu *et al.*, 2010). I am grateful to the various audience members of these presentations for their feedback, and for their requests that I write it all down. Thanks in particular to John Rooksby and Linda Northrop, for their comments on parts of earlier written versions of this text.

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 in 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 approx 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 660 points, its biggest ever one-day loss of points, representing the disappearance of around one trillion 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 haywire, 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.

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 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 almost random.

Two weeks later, the US Securities and Exchange Commission (SEC) and the Commodity Futures Trading Commission (CFTC) jointly released an interim report into the events of May 6th (CFTC&SEC, 2010a) that established 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), the consequent lack of a single unifying “consolidated tape” showing unique timestamps for all events in all the markets, and 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 a only few seconds or less. When the SEC/CFTC final report on the Flash Crash was finally published on September 30th, nearly five months after the event (CFTC&SEC, 2010b), it stated 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.

I argue here that the Flash Crash is best understood as a failure in a large-scale complex socio-technical system of systems (SoS), and that there are other such socio-economically significant SoS in which similar problems or failures seem likely, or at least plausible, in future. Unpacking that assertion requires some care, so we’ll start first with a discussion of notable technology failures, then bring the conversation back to discussion of failures of the financial markets, and then extrapolate out to other large-scale complex socio-technical SoS. Such SoS are very often the result of organic growth and unplanned accretion rather than clean-sheet engineering design, thereby involving or acquiring significant degrees of variability in components and heterogeneity of constituent systems. For this reason traditional engineering techniques cannot necessarily be trusted to deliver acceptable solutions. Therefore, new approaches are required.

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 was responsible for failures, but rather that the failures were inherent, latent, in the system; this seems especially so in the case of *socio-technical systems*, i.e. systems whose dynamics and behaviour can only be properly understood by including human agents (such as operators and/or users) within the system boundary.

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 recovery procedures could be rehearsed and refined before being relayed to the astronauts. The value of simulation models is something that we will return to 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 the number of deaths resulting from that event is widely held to be many thousands.

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 gave a detailed study 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. 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 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 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, immediately following the loss of *Columbia* in 2003, because of her work on *Challenger*, Vaughan was 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, 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 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. Prior to loss of *Columbia*, sixty-four previous missions had suffered strikes from insulating material breaking away during launch and hitting the orbiter, each such strike was technically a violation of design requirements (most notably, in 1988 on mission STS-27, insulation broke away from an SRB and damaged 700 of the heat-insulating tiles on shuttle *Atlantis*) but repairing the damage from insulation strikes became increasingly seen as 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 so-called High-Reliability Organizations (HROs) 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).

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 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 which 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 the events of 2007-08 have been recounted from a number of journalistic perspectives, of which Lewis's (2009) and Tett's (2009) are notably thorough and well written. Tett's perspective is particularly insightful: she is a senior 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, or at least very close relatives 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 MacKenzie and his colleagues (MacKenzie 2008a, 2008b; MacKenzie *et al.*, 2008), although all of those pre-date the turmoil of the past three years. In a strategic briefing paper written for the UK Government's Office of Science, I discussed the danger that normalization of deviance posed in high-frequency trading systems in the global financial markets, and the possibility of major catastrophe happening within very short time-scales; the final version of that paper (Cliff, 2010) was submitted to the government nine days before the Flash Crash.

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, through a twenty-year stage of being phone-and-screen-based rather than face-to-face, but still largely requiring a human at each end of the phone or screen. Within the past decade, however, 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. Many automated trading systems have limiters and circuit-breakers built in to prevent major losses, but as was seen on May 6th, the system-wide interaction between multiple independently-engineered 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.

The global financial markets, considered as a system, is made up of components, of constituent systems such as the various electronic exchanges and the automated and the manual trading systems operated by the various investment banks and hedge funds, all of which have been developed, procured, operated and managed independently. That is, the current global financial markets are, from a technology perspective, *systems of systems* (SoS). 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. 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 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.

So, the concerns expressed in 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 summarize schematically in a single three-line graph, which we show in Figure 1.

The UK Research and Training Initiative in the Science and Engineering of LSCITS was started in 2007 as a national strategic investment with a primary aim being to foster the formation of a new community of researchers and practitioners with training and experience appropriate for dealing with future engineering dominated by LSCITS issues. In the next section there follows a very brief overview of the UK LSCITS Initiative, for more details, see (Cliff *et al.*, 2006; Calinescu *et al.*, 2010); and for an overview of a very similar American initiative, the CMU/SEI Ultra-Large-Scale Systems Project, see (Northrop *et al.*, 2006).




FIGURE 1: The LSCITS 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. I am deliberately non-committal on the precise timing of this crossover point: it could be a decade or more away, it could have happened already.

3. The UK LSCITS Initiative

The range of tools and techniques required to fully address issues in the science and engineering of LSCITS spans a spectrum from the mathematics of complex adaptive nonlinear systems, to the study of organizational processes that enable the engineering of complex socio-technical systems. The need for an interdisciplinary approach is manifest. The UK LSCITS Initiative has been structured to involve all the constituent research fields that are directly relevant; our primary interest is what happens at the intersections of these fields. The constituent fields are:

- *Complexity Science*: the study of interacting coupled nonlinear dynamical systems that exhibit system-wide emergent behavior; and especially complex adaptive systems (CAS), where there is some path-dependency or adaptive processes within the components that make up the system, such that it exhibits ‘evolution’ or ‘learning from experience’ over time. See, for example, Waldrop (1992), Axelrod & Cohen (2000), Bar-Yam (2005), Braha *et al.*, (2006), and Mitchell (2009).
- *Predictable Software Systems (PSS)*: the study of formal methods for computer systems engineering, involving the creation of models of existent or proposed systems, expressed in a formal language that can then be subject to automated theorem-proving and/or simulations for comprehensive search of the state-space. When the model has been proven or demonstrated to satisfy its requirements, it can serve as input to automated code-generators. PSS techniques show great potential, but their applicability in (ultra-)large-scale systems is limited by the poor scaling properties of current methods; improving their scalability is a current research issue.
- *High Integrity Systems Engineering (HISE)*: recognizing that PSS approaches cannot currently fulfill all the needs of real-world applications where software is responsible for maintaining the integrity of safety-critical systems (such as aerospace systems, nuclear power stations, defence command & control), there is a distinguished body of engineering research and teaching that has been developed for engineering such high-integrity systems. The intellectual heritage of much of this work is traceable back to classical engineering control theory, with its separation of *plant* (the thing to be controlled) and *controller* (the thing that does the controlling); in consequence, historically HISE work addressed systems that have relatively simple control interfaces (buttons, dials, throttles) and considers the human users as essentially external to the plant-controller system. In many systems of significant interest, that approach is entirely appropriate. However, for many systems of systems, it is necessary to consider the human agents as components within the system, i.e. to adopt a socio-technical perspective. HISE work now embraces such issues, linking to work on social-technical systems and embracing the system of systems paradigm.
- *Socio-Technical Systems Engineering (STSE)*: historically, the majority of research in socio-technical systems has been largely *a posteriori* descriptive, using techniques from human-computer-interaction, organizational and social psychology, and workplace ethnography, to formulate insightful analyses of existing systems. There is a less well-developed tradition of developing tools and techniques that are predictive and prescriptive; i.e., of developing *a priori* methods for engineering

socio-technical systems. Extending socio-technical analysis techniques into software-systems engineering is a current research challenge.

- *Complexity in Organizations (CiO)*: it is manifest that many LSCITS have come into existence to support large-scale networks of interacting groups or organizations. The interacting groups or organizations may be separate independent enterprises, or they may be divisions within some larger single enterprise. Here the word “enterprise” I intended to include public sector organizations as well as private firms; and the groups or organizations can be any size. Very often, the complexity in the LSCITS is present as a direct reflection of the complexity in the (network of) organizations, and understanding that organizational complexity is a deeply non-trivial problem. Thus, the LSCITS Initiative has a commitment to work with organizational theorists, economists, and management and political scientists to understand these social aspects of socio-technical LSCITS.
- *Novel Computation Approaches (NCA)*: finally, as the size of socio-technical LSCITS increases, many traditional engineering techniques do not scale well. Much engineering practice, that has served us so well for decades, is based on divide-and-conquer: a process of recursive hierarchical decomposition is applied, breaking a system into subsystems, which are in turn decomposed into sub-sub systems, continuing until the system is decomposed to a sufficiently low level that the engineering of each component is tractable. Such hierarchical decompositions often result in engineered systems having hierarchical tree-structured control architectures, which scale badly because of the need for communication from the leaf-nodes up to the root node, this is a source of delays and noise, and the root node is a significant vulnerability, as a single point of failure. In the past two decades a variety of novel, decentralized, approaches to management and control of large-scale and distributed systems have been developed; many of them stemming from Complexity Science studies of CAS. The Initiative’s research is intended to include explorations of the application of these novel decentralized approaches in LSCITS contexts. They are perhaps most relevant in the design, management, and control of ultra-large-scale IT facilities required for “cloud computing”, discussed further in (Calinescu *et al.*, 2010)

From its outset, the LSCITS Initiative has acknowledged that there are unlikely to be productive lines of inquiry in trying to forge a single unifying intersection between *all* of these constituent fields, nor even in trying to force intersections between every possible pairing of these subfields; nevertheless there is a coherent path through the fields and their intersections, linking from one end of the spectrum to another, in a configuration that (Calinescu *et al.*, 2010) refer to as “The LSCITS Stack”, as illustrated in Fig. 2.

In case it is not already obvious, relating the LSCITS Stack to the large-scale complex IT-centric socio-technical system that we see in today’s flash-crashing global financial markets is relatively straightforward. For the purposes of this paper, it’s probably sufficient to merely ask some questions: if some authority or agency, or a lone investment fund or hedge fund even, had created a PSS-style model, could that have been used to identify the possibility of events such as the Flash Crash, before it happened? Would the Flash Crash have been possible at all if the people and organizations who develop and manage the IT systems that enable automated trading had used HISE approaches and been subject to the same stringent regulations (and, indeed, legislation) as are commonplace in sectors such as defence, aerospace, and nuclear power? Could the tools of Complexity Science (typically a mix of mathematical

analysis and empirical computer simulation models) have been used to identify the likelihood of the Flash Crash before the event, and given guidance on how to identify if and when the system as a whole was entering the danger-zone (the “basin of attraction in its phase-space”, in the technical language) where flash-crash type events were much more likely, or perhaps inevitable? Could decentralized NCA techniques be developed to re-engineer the financial markets to make such failures less likely?

None of these questions are intended as rhetorical. In truth, all of them are open research issues. In the next section, I’ll explore one particular approach, at the intersection of several of the fields in the LSCITS stack, that seems likely enough to provide a major payoff that it could be worth the (considerable) investment required.




FIGURE 2: The “Stack” of overlapping research fields which constitute the LSCITS Initiative. The acronyms are defined and explained in the text.

4. 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 prescriptive (saying what should be done differently in future, to predict or prevent such failures from re-occurring).

One possible approach is to accept that technology failures are simply to be expected 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, “shit 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 project-management, for all LSCITS including those underlying the global financial markets, but such changes are likely to take time, and in the 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 paragraphs, I 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 *Arianne* 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, 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 four 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 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).

Of course, computational simulations are 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 it is much less commonplace to simulate entire markets at a fine-grained level to study issues in overall system behaviour. Nevertheless, this is now quite certainly within the realms of possibility. In a recent book, Darley & Outkin (1997) relate their use of complex adaptive systems (CAS) 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 makers 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; 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 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 investigate various what-if scenarios by using user-defined strategies.”

(Darley & Outkin, 1997, p.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).

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 Office¹. 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. From a thousand runs (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

¹ www.metoffice.gov.uk

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 I am 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 ones rather than computer simulations) at NASA Mission Control. The simulators had been developed to train the astronauts in dealing with failure situations, but after the explosion on *Apollo 13* they also became the test-bed for evaluating potential rescue procedures. The use of simulation models as scientific evaluation and training tools for humans dealing with complex situations has a long history: see, e.g., Sloan (1981) and Dorner (1990, 1997).

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 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, but have no trading-technology certification processes aimed at preventing financial crashes, in the future, this may come to seem curious.

I’m not arguing here that predictive simulation models are a “magic 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.

The significant cost of constructing and operating such a simulation facility could 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, etc.) operating in a particular country or group of countries. Revenue could be generated from levying charges for anyone wanting access to the simulator, and also possibly from using the simulator as a training or certification facility (discussed further below). 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. However, 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. 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 compete on the basis of what questions they ask of the simulator (details of each corporation's specific experiments could be kept confidential), and how they 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.² 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.

² 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. Weather forecasting services on mainland Europe, using different monitoring and prediction models, had given more accurate forecasts of the weather that night.

There is an old saying, “if it ain’t broke, don’t fix it”. This is 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 really is to fix them while they ain’t broke, because that’s the only decent chance you’ll get. This is the case for most large-scale complex IT systems (LSCITS). Ensemble forecasting via n -plex redundant predictive simulation models is not cheap, and is certainly far from perfect, but it may be the best option currently available.³

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 mentions predictive simulation modeling only in passing (Bonisch & Di Giammarino, 2010). Nevertheless, I am certainly not the only person to be making such proposals: see, e.g. (Economist, 2010).

5. Summary

The Flash Crash of May 6th 2010 was a sudden and dramatic failure in a 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 wrapped 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 large-scale complex socio-technical 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. Running multiple redundant predictive simulation models is one approach that may now be applicable for assessing and controlling systemic risk in the financial markets.

The UK LSCITS Initiative and the USA’s Ultra-Large-Scale Systems Initiative are each national-level strategic initiatives. 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. The engineering of LSCITS is in its infancy, it has some significant differences from engineering smaller-scale systems, and developing rigorous trusted approaches may be a long haul; any researchers, practitioners, regulators, policy-makers or sponsors who would like to become involved in the LSCITS initiative are very welcome to get in touch.

³ For recent counterarguments to the use of simulation models, see Turkle (2009).

References

W. B. Arthur (2009). *The Nature of Technology: What it is and How it Evolves*. Allen Lane.

R. Axelrod & M. Cohen (2000). *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. Free Press.

Y. Bar-Yam (2005). *Making Things Work: Solving Complex Problems in a Complex World*. Knowledge Press.

E. Beinhocker (2007). *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics*. Harvard Business School Press.

L. Blume & S. Durlaf (2005). *The Economy as an Evolving Complex System, III*. Addison-Wesley.

P. Bonisch & P.J. Di Giammarino (2010). *Achieving Supervisory Control of Systemic Risk*. Report jointly produced by FS KTN, JWG, and Paradigm Risk.

D. Braha, A. Minai, & Y. Bar-Yam (2006). *Complex Engineered Systems: Science Meets Technology*. Springer

R. Calinescu, & M. Kwiatkowska (2010). Software Engineering Techniques for the Development of Systems of Systems. In: Choppy, S. and Sokolsky, O. (editors), *Foundations of Computer Software: Future Trends and Techniques for Development*, vol. 6028 of LNCS, pp. 59-82, Springer.

R. Calinescu, S. Kikuchi, & M. Kwiatkowska (2010). Formal Methods for the Development and Verification of Autonomic IT Systems. To appear in: Cong-Vinh, P. (editor), *Formal and Practical Aspects of Autonomic Computing and Networking: Specification, Development and Verification*, IGI Global.

R. Calinescu, D. Cliff, J. Keen, T. Kelly, M. Kwiatkowska, J. McDermid, R. Paige, & I. Sommerville (2010). *The UK Large-Scale Complex IT Systems (LSCITS) Initiative*. Manuscript available from <http://lscits.cs.bris.ac.uk/docs/LSCITSOverview2010.pdf>

CFTC & SEC (2010a). *Preliminary Findings Regarding the Market Events of May 6th, 2010*. Report of the staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory issues. May 18th 2010: <http://www.sec.gov/sec-cftc-prelimreport.pdf>

CFTC & SEC (2010b). *Findings Regarding the Market Events of May 6th, 2010*. Report of the staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory issues. September 30th, 2010. <http://www.sec.gov/news/studies/2010/marketevents-report.pdf>

D. Cliff, J. Keen, M. Kwiatkowska, J. McDermid, & I. Sommerville (2006). *Large Scale Complex IT Systems (LSCITS) Research Programme*. Research proposal to the UK Engineering and Physical Sciences Research Council; submitted December 2006, commenced April 2007. <http://lscits.cs.bris.ac.uk/docs/LSCITSProposalRP1.pdf>

D. Cliff (2010). *Networked Governance in the Financial Markets*. Foresight strategic briefing paper, for UK Government Office of Science & Technology, Department of Business, Innovation, and Skills. Available from: http://www.cs.bris.ac.uk/home/dc/Foresight_NetGov_v2a.pdf

- D. Colingridge (1992). *The Management of Scale: Big Organizations, Big Decisions, Big Mistakes*. Routledge.
- V. Darley & A. Outkin (2007). *A NASDAQ Market Simulation: Insights on a Major Market from the Science of Complex Adaptive Systems*. World Scientific.
- H. Dezfuli, et al. (2009). *Bayesian Inference for NASA Probabilistic Risk and Reliability Analysis*. NASA SP-2009-569:
<http://www.hq.nasa.gov/office/codeq/doctree/SP2009569.pdf>
- D. Dorner (1990). The logic of failure. *Philosophical Transactions of the Royal Society of London, Series B*. **327**(1241): 463-473.
- D. Dorner (1997). *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*. Perseus.
- Economist (2010). Agents of Change. *The Economist*, **396**(8692):76. [Note that *The Economist* has a standard policy of not showing author bylines for articles written by regular staff journalists].
- M. Gardner (1971). On Cellular Automata, Self-Reproduction, the Garden of Eden, and the Game 'Life', *Scientific American*, 224(2):112-118.
- J. Gray (2009). "On eScience: A Transformed Scientific Method" in T. Hey, S. Tansley, & K. Tolle (editors), *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Microsoft Press. Pp. xvii—xxxii.
- A. Haldane (2009). *Rethinking the Financial Network*. Text of a speech given at the Financial Student Association, Amsterdam, April 2009. Available from:
<http://www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf>
- E. Hollnagel, D. Woods, & N. Leveson, editors, (2006). *Resilience Engineering: Concepts and Precepts*. Ashcroft.
- D. Hubbard (2009). *The Failure of Risk Management. Why It's Broken and How to Fix It*. John Wiley.
- C. Kindleberger (2001). *Manias, Panics, and Crises: A History of Financial Crises*. John Wiley.
- M. Lewis (2010). *The Big Short: Inside the Doomsday Machine*. Allen Lane.
- E. Lorenz (1963). Deterministic Nonperiodic Flow. *Journal of Atmospheric Science*, **20**:130–141.
- D. MacKenzie (2008a). *An Engine, Not a Camera: How Financial Models Shape Markets*. MIT Press.
- D. MacKenzie et al., editors (2008). *Do Economists Make Markets? On the Performativity of Economics*. Princeton University Press.
- D. MacKenzie (2008b). *Material Markets: How Economic Agents are Constructed*. Oxford University Press.

- M. Mitchell (2009). *Complexity: A Guided Tour*. OUP.
- R. Mullane (2006). *Riding Rockets: The Outrageous Tales of a Space-Shuttle Astronaut*. Simon & Schuster.
- L. Northrop *et al.* (2006). *Ultra-Large-Scale Systems: The Software Challenge of the Future*. Technical Report. Carnegie Mellon University Software Engineering Institute.
- P. Ormerod (1998). *Butterfly Economics: A New General Theory of Economic and Social Behaviour*. Faber.
- P. Ormerod (2006). *Why Most Things Fail: Evolution, Extinction, & Economics*. Faber.
- C. Perrow (1984). *Normal Accidents: Living with High-Risk Technologies*. New York: Basic Books.
- J. Reason (2008). *The Human Contribution: Unsafe Acts, Accidents, and Heroic Recoveries*. Ashgate.
- K. Roberts (1990). Some Characteristics of One Type of High reliability Organization. *Organization Science* **1**(2):160-176.
- S. Sloan (1981). *Simulating Terrorism*. University of Oklahoma Press.
- M. Stamatelatos *et al.* (2002a). *Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners*. Version 1.1. www.hq.nasa.gov/office/codeq/doctree/praguide.pdf
- M. Stamatelatos *et al.* (2002b). *Fault Tree Handbook with Aerospace Applications*. Version 1.1. www.hq.nasa.gov/office/codeq/doctree/fthb.pdf
- N. Taleb (2007). *The Black Swan: The Impact of the Highly Improbable*. Allen Lane.
- G. Tett (2009). *Fool's Gold: How Unrestrained Greed Corrupted a Dream, Shattered Global Markets, and Unleashed a Catastrophe*. Little, Brown.
- P. Tuinenga (1988). *SPICE: A Guide to Circuit Simulation and Analysis Using PSpice*. Prentice Hall.
- S. Turkle (2009). *Simulation and its Discontents*. MIT Press.
- D. Vaughan (1997). *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. University of Chicago Press.
- D. Vaughan (2005). On Slippery Slopes, Repeating Negative Patterns, and Learning from Mistakes. In Starbuck, W. & Farjoun, M., editors (2005) *Organization at the Limit: Lessons from the Columbia Disaster*. Wiley Blackwell. Pp. 262-275. <http://www.sociology.columbia.edu/pdf-files/vaughan01.pdf>
- D. Vaughan (2006). NASA Revisited: Theory, Analogy and Public Sociology. *American Journal of Sociology*, **112**(2):353-393. <http://www.sociology.columbia.edu/pdf-files/nasa.pdf>

M. Waldrop (1992). *Complexity: The Emerging Science at the Edge of Order and Chaos*. Simon & Schuster.

K. Weick & K Sutcliffe (2007). *Managing the Unexpected*. Jossey Bass.