Networked Governance in the Global Financial Markets

Sketch for a Foresight project; third draft, Feb. 2010 (with minor edits Apr. 2010).

Prof. D. Cliff, UK Large-Scale Complex IT Systems Initiative, University of Bristol, dc@cs.bris.ac.uk

The past 15 years has seen a surprising, but well-documented, major technology transition in most industrialised economies to the use of advanced telecoms (wireless mobile and wired broadband) for accessing information via the internet and the world-wide-web. There are at least three major technologies that look set to combine in coming decades, in such a way that major transitional disruptions (either positive or negative) look likely for western industrialised societies such as the UK in terms of how citizens interact with each other, how business is done between firms, and how both firms and people interact with institutions of government.

The three technologies likely to combine in disruptive ways are the world-wide-web (especially in its currently much-vaunted incarnation as "Web2.0"), coupled with the easy accessibility of cheap remotely-accessible ultra-high-power computing facilities (so-called "Cloud Computing"), coupled with advanced apparently-intelligent automated processing of machine-readable information (so-called "computational intelligence" and "semantic web" technologies). But these three technologies are not the only ones making rapid advances. It is already clear that advances in materials and manufacturing, in the psychopharmacology of drugs for cognitive enhancement, and in brain-machine interfaces, all also have significant disruptive potential, thereby increasing the potency of that mix and adding to the uncertainty of various possible futures, and potentially also increasing the desirability of those futures.

Some, but certainly not all, of the issues to be addressed here could be described as "Web Science" in that they involve richly interconnected webs (i.e., networks) of interacting entities with complex dynamics. Ted Nelson, the American information-technology pioneer, sociologist, and philosopher, coined the term "intertwingling" in 1974[1] to express the often-unacknowledged difficulty encountered when attempting to divide richly cross-connected topics into hierarchical or sequential structures. The factors to be explored in this proposed Foresight project are manifestly "intertwingled" in that there is not necessarily one central technology driving the suggested project. Rather, this project is proposed because there are multiple simultaneous developments in science and technology, and in the maturation of business processes and regulatory environments and societal norms and expectations, that all look set to converge in ways that justify examination by Foresight's methods.

As the Foresight Programme has now been running for 15 years, specifications of its new projects need to be ever more careful to avoid repeating work from its past projects. The project sketched here is motivated by the desire to address issues central to one major potential lead government department: Her Majesty's Treasury. HMT has yet to benefit from a Foresight project focused on issues of central concern to that department. In essence, the proposition here is to run a Foresight project that will explore, and attempt to understand, issues of stability and controllability in future global financial markets.

The project proposed here meets all of the Foresight project-selection criteria[2]. It requires looking ahead at least 10 years, in areas where the outcomes are uncertain. This is because the future rate of change is rapid, the direction of current trends is uncertain, and multiple different trends may converge or indeed diverge. Science and technology are the main drivers of change; the outcomes can be influenced to some extent, and are manifestly significant for the economy and/or society. This project is not duplicating work carried on elsewhere and yet it builds on areas of active research; it clearly requires an inter-disciplinary approach to the relevant science and technology, and would bring together groups from business, academia, and government (most notably via the involvement of regulatory authorities): the issues to be explored here do not seem capable of resolution by a single group or capable of being addressed by a set of experts from a single domain. This project has the clear potential to command the support of the groups most likely to be able to influence the future, and to be owned by a lead government department. It is also addressing a critically important issue: the factors driving stability and instability of the financial markets in the future.

In recent years, sizeable investments have been made in the UK science-base that start to better enable it to tackle the problems discussed here. The UK Engineering and Physical Sciences Research Council (EPSRC) has funded the establishment of major doctoral training centres, and at least one national research initiative, which could all make valuable contributions; several UK universities have also recently created relevant centres aimed at tackling complex problems that require collaboration across several disciplines.[3]

Stability of the global financial markets is a central theme of this proposed project, but is not the only issue: the nature of work and employment in the future financial industry is also open to question. Higher levels of automation, enabled by ongoing reductions in the cost of computing technology and increases in algorithmic speed and sophistication, allows ever more work to be done by ever fewer people. And the work that does actually need to be performed by humans may not necessarily need to be performed by people in a high-cost location such as New York's Wall St, or the City of London. Bangalore, Beijing, or Bangkok may all be future contenders.

There are two primary aims of this proposed project. One aim, as would be expected of any Foresight project, is to explore issues of how various outputs from science, engineering, and technology (such as the sciences of network mathematics, of complex adaptive systems, of socio-technical systems engineering, of economics and finance, of industrial organisation, and the psychology and sociology of individual and collective behaviours of market participants) are likely to combine in fundamentally interesting ways. The second, and more significant, aim is to deliver findings of genuine value to HMT, enabling a better predictive understanding of key issues that will affect stability in future networked global financial markets; of the extent to which negative influences on stability and resilience can be influenced, avoided, and controlled by regulatory and policy mechanisms, or by other means; and the implications for national and international economies.

Science, Technology, Business, and Governance in Future Financial Markets.

Proposed lead/owner department: HM Treasury

The global financial markets have been proactive early-adopters of new technologies for most of their history. In the past quarter of a century, since the instigation of the "big bang" switch to paperless electronic trading, the City of London has led the world in the adoption of new ICT technologies for the provision of electronic trading facilities, and the associated distribution of data and news feeds. This hunger for new technologies looks unlikely to be particularly diminished in future, even after the unprecedented turmoil of the 2007/2008 credit-crunch crisis.

Furthermore, the new regulatory framework that is being introduced in response to the events of 2007/08 seems set to greatly increase the demand for accurate and sophisticated risk-management in many (and possibly all) of the institutions that participate in the global financial markets; and that demand is likely to be met by ever more compute-hungry risk-modelling systems. Risk management has become so heavily dependent on computational power that the constant availability of high-power computing facilities is now vital to most risk-management functions. Moreover, as institutions seek to cut costs and boost profitability in the new environment, any function that can be delegated to computerised automation is likely to be so: for many roles, humans may remain in the control loop but there will be fewer such humans and their primary role will be that of monitoring and supervising the automated technology.

Several commentators have noted that the rich web of communications connectivity and massive information-processing capacity distributed across investment banks, fund-management companies, and electronic exchanges (and other trading venues) forms a large-scale and highly complex information-technology system whose full complexity is probably unknown and possibly unknowable.

With this lack of knowledge comes the potential for surprisingly unpredictable behaviour, volatility, and failure.

To take one example: there is little doubt in informed circles that the credit crunch crisis was triggered at least in part by the use of advanced computer-optimization techniques to automatically create or fine-tune structured credit instruments (such as Collateralised Debt Obligations) that no single trader, and indeed no single financial institution, fully understood the detailed rationale for, nor the full systemic risk implications thereof. However, the crisis was certainly not *solely* a consequence of advanced technology being misapplied or going wrong. Analysing the global financial markets as a purely technical system misses the point: the markets are *socio-technical* systems, in the sense that key elements in any causal explanation of their operation are people, as well as machines; and those people are not external "users" of the system, rather they are vital components within the system.[4]

The possible *global* systemic effects of a widespread (and *locally* rational) reliance on computer-optimized instruments for spreading risk were simply not well understood with respect to the potential impact to the individual trader's position, nor to the potential impact to that trader's employer's corporate/institutional position, and nor to the possible negative effects on the wider macro-level system of interconnected corporations and institutions. Not, that is, until it was too late.

But all is not lost: global financial systems nevertheless do yield to scientific analysis. In recent decades a new set of scientific tools and techniques has been developed to study *complex systems*: 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. In addition to exhibiting emergent behaviour, many complex systems of significant interest are *adaptive* (in mathematical language, they show path-dependencies) such that the response of the system to a particular change in its inputs or internal activity may be dependent on its history of previous inputs or activity-states: colloquially speaking, the system may appear to "learn from experience" or "evolve in response to changing circumstances". Since the late 1980's a growing number of scientists have been attempting to understand the financial markets as complex adaptive systems (CAS), and 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.[5]

It is fair to say that, during the run-up to and unfolding of the recent crisis, the local-vs-global issues were poorly understood not only by the market participants (the traders, their management, and the shareholders in their companies) but also by the market regulators and political authorities. The so-called "Persaud Paradox", that the observation of safety creates risk (where large numbers of market participants take very similar risk-reducing "safe" positions and thereby, via the near-homogeneity of their positions, greatly increase the overall systemic risk) was spoken about as something of an idle curiosity when the phrase was first coined in a 2005 *Financial Times* article,[6] but within a couple of years events in the real world had conspired to make its downside effects a rather stark reality.

In the last five to ten years there has been a very rapid growth in "automated execution" or "algorithmic trading" systems in the financial markets, where autonomous and adaptive software programs do the trading that was once entrusted to human traders on the dealing floors. As this has taken effect, fewer human traders have been employed – increasingly in an oversight role rather than as active market participants – and those humans are, if some plausible reports are to be believed, ever more tempted to employ "nootropic" cognitive-enhancer pharmaceuticals to sharpen their responses as one additional tactic for safeguarding their own jobs, by beating the human competition. If there are roles that human traders remain better-suited to than machine traders are, then those roles are likely to involve ever more tightly-coupled interaction with ever-more complex machine systems.

The nature of the human-machine interactions in the financial markets, and the interfaces that are best suited to those interactions, looks to be ripe for a series of innovations based on new interaction and interfacing technologies such as multi-touch screens, 3D displays, ungameable interfaces, haptic input/output devices, and direct interfacing to the nervous system. For each such technology, if it gives the trader, or the trader's employer, a significant advantage, it is likely to be seized upon by as many traders as can afford the technology and also feel they can afford the risks.

But developments in pharmaceuticals that sharpen a trader's thinking, or in interface technologies that lessen the time a trader takes to instruct a trade, are likely only to offer medium-term gains. As with many arms-races, new innovations only confer competitive advantage for as long as it takes the competitors to copy that innovation or to come up with counter-innovations of their own: as soon as all traders are using a particular new technology, the playing field is levelled again.

In fact, in the visible longer term (and well within the time-horizon of a Foresight project) it seems entirely reasonable to speculate that the number of human traders involved in the financial markets will fall very dramatically indeed. While unlikely, it is not impossible that human traders will simply no longer be required at all.

Although the last 30 years of robotic shop-floor automation in manufacturing industries has made the replacement of unskilled/blue-collar labour by computer technology a commonplace event, the near-elimination of a sector of the human workforce who are widely (self-) described as possessing uncommonly high levels of intelligence, quick-wittedness, and calmness-under-fire, to be replaced by mere machines, perhaps seems a somewhat more fanciful proposition.[7]

Nevertheless, the likely replacement of human traders by machine systems is entirely plausible when one considers two factors: first, the rapid growth in the amount of information that can need to be assimilated and integrated in order to make profitable trading decisions; second, the rapid reduction in the reaction time required between an event occurring in a market and traders issuing orders that can capitalise on that event. The simple fact is that we humans are made from hardware that is just too bandwidth-limited, and just too slow, to compete with coming waves of computer technology.

For sure, for the foreseeable future there are likely to always be certain classes of transactions that require human skill and judgement (and subtlety), but it seems plausible that the number of such transactions will decrease markedly in coming years. Algorithmic trading systems are increasingly able to replace human traders at the point of execution, and open up the possibility of machine-trading on short timescales that humans would simply find too labour-intensive (or error-prone) to realistically contemplate.

Financial-market traders are not the only intelligent, quick-witted and calm professionals whose high-pay high-status jobs are imminently replaceable by information technology. Aeroplane pilots, both civil and combat, are (within the defence and aerospace industry at least) now increasingly seen as dispensable "spam in the can". It is commonly accepted in the industry that the current generation of combat aircraft platforms will probably be the last to carry human flesh: remotely-operated[8] semi-autonomous unmanned air vehicles (UAVs, or "drones") can be built to serve much the same combat requirements, without having to worry about keeping a human pilot conscious in thin air or high-g manoeuvres, shielded from bullets, and able to eject. Recent combat UAVs have automatic take-off and landing (ATOL) systems fitted as standard, further adding to their autonomy. In a similar vein, modern civilian airliners have autopilot systems that can readily navigate transcontinental routes and safely land the vehicle in zero-visibility weather conditions: routine deployment of full ATOL systems for commercial passenger jets seem entirely plausible within the foreseeable future. At that point, the airliner's pilot will be largely redundant, but will no doubt retained for two reasons: reassuring the passengers (who might otherwise find the prospect of sitting in the belly of a flying robot for several hours on transoceanic flights a somewhat disheartening prospect); and dealing with the "black swan" extremely rare but very-high-consequence events that the autopilot and ATOL may simply not have been programmed to deal with (such as the January 2009 safe ditching of US Airways Flight 1549 in the Hudson River at Manhattan Island after loss of both engines from bird-strike during climb-out).

It seems reasonable enough to ask whether the future of traders at the points of execution in the financial markets may be going the same way as combat pilots (likely to be entirely replaced by drone technology) or as passenger-jet pilots (likely to be retained primarily for customer-assurance and crisis-management). Right now, plausible stories could be argued for either outcome.

There is a third outcome, one that is quite popular on the trading floors of major investment banks, which is that neither of the above two outcomes occurs, and the traders continue in their jobs much as before. The comforting appeal of this third story is clear: politically, it is extremely difficult for the manager of a trading floor to introduce technology that threatens the future of the traders on that floor: in most of the world's major financial centres there is a high degree of worker mobility, and a good trader feeling threatened may be tempted to leave for a competitor, the day after the annual bonus is paid.

Of course, there is no single definition of "trader" and in practice the impact of automated trading technology will vary over specific types of trading roles. For execution-only "sales traders" working in investment banks, and probably also for market-makers or specialists, the impact is likely to come sooner rather than later, and to be largely very negative when it arrives. For proprietary traders and investment managers, the negative effects may be less pronounced and take longer to occur, because automated portfolio selection and optimization management techniques are not so well developed, and require significant technical challenges to be addressed before human-level performance is achievable by machines.

Students of technology innovation in various industry sectors have repeatedly observed that when a new technology or innovation disrupts an industry, often the dominant positions of the previous major incumbent companies in that sector are damaged or displaced by the arrival of new companies that are better able to understand and exploit the innovation.[9]

With the recent changes in the regulatory frameworks of the financial markets in Europe and in the USA (most notably, MIFID in the EU), the previously widespread heavily vertically integrated business model of investment banking acting as the "sell side", selling their access to centralised market liquidity to "buy-side" fund managers, is clearly entering a disruptive de-verticalisation phase. In consequence, a new business ecosystem of small and medium-sized suppliers of component technologies ("middleware" such as market-data event aggregators/processors, trader interfaces, and so on) and new alternatives to existing structures (such as the "dark pools of liquidity" provided off-exchange by so-called Alternative Trading Systems) is rapidly growing. Fund-management companies can now pick-and choose their technology components, and their trading venues, and replicate much – possibly all – of the functionality that they previously paid investment banks to provide.

There is a joke that in the near future a typical trading floor will consist of a large room full of high-power server computers, and in the corner there will be one man and one dog. The dog's job? To make sure no-one touches the computers. The man's job? To feed the dog. It is generally not a good idea to tell this joke on the trading floor of an investment bank.

The transition from trading floors full of people, screens, and phones to highly depopulated server-rooms is already well underway.

Major exchange-operators have for several years offered "hosting" services whereby a number of "proximity servers" are placed in uniform proximity to the main server that runs the exchange's order-matching, transaction processing, and reference-data functions. Here, "uniform proximity" means essentially that each of the proximity-servers is connected to the main exchange server by a cable of the same length. Thus all the proximity servers have access to the exchange server with exactly the same degree of communications latency (i.e., the time it takes for the electronic order to leave the trader-server and arrive at the exchange-server) and transaction time (i.e. the time it takes for the exchange server to execute the order).

Latency is a key issue in automated trading: if your trading algorithm consistently beats mine on trades of a particular type, that simply won't matter to me if I can get my trader executing in the market sooner than yours – orders from your better trader will arrive at the market only to find that my orders have already been executed, denying yours the opportunity to trade.

Investment banks and fund management companies can no longer always tolerate the few milliseconds of delay that it takes the small number of bytes in an electronic order to leave the institution's intranet, traverse the telecoms lines that connect the institution to the exchange-operator perhaps only a few city-blocks away, and then once inside the exchange find their way to the main server. To get around this, institutions now increasingly develop trading algorithms at their own offices, and then upload them to the hosting proximity server at the exchange, where they are guaranteed a transmission latency that is identical to those of their competitors. Trading algorithms running on low-latency server architectures now routinely operate on millisecond timescales.

That is, they may routinely identify an opportunity to trade, execute an order (a buy, say), wait for the position that results from execution of the order to become profitable, and then execute a fresh order (a sell, say) to lock in the value, all within a few seconds, and possibly in a fraction of one second.[10]

This is known as high-frequency trading. Common current estimates for major equity (stocks-and-shares) exchanges around the world are that 50-70% of trading volume is now generated by high-frequency algorithmic trading systems.

In the coming one to two decades, it is clear that the next major shift in the arms race could be to trading in the sub-millisecond time domain, which we'll refer to here as Very High Frequency (VHF) trading.

To do this, it will be necessary to move away from using trader-servers that are built from commodity hardware (i.e., racks of PC motherboards with standard Intel CPU chips running Windows or Unix or Linux operating systems): such commodity hardware is general-purpose, and the generality of the design introduces elements that cause unnecessary delays.

Eliminating these delays is straightforward if the trading algorithms cease to run as program-code on a general-purpose computer, and instead are "cast in silicon". That is, the trading algorithms can be implemented on dedicated silicon chips, via established techniques such as re-purposing high-performance computer-graphics boards (which often have tens or hundreds of parallel computer "cores" built into each graphic chip), using well-established technologies such as field-programmable gate arrays (FPGAs), or using newly developed technologies such as the software-defined-silicon for programmable multi-core chips as pioneered by the UK company Xmos.[11]

Further down the line, yet within the realms of possibility, quantum computing devices may come commercially viable within a couple of decades or so, allowing progress from the microsecond-range of VHF to the nanosecond-range, which we'll refer to here as UHF trading.

As automated high frequency trading becomes ever more commonplace, and as VHF and UHF trading supplant it in due course, so the timescales on which small dramas can rapidly become major crises will shrink by several orders of magnitude. In a plausible future involving high proportions of automated VHF or UHF trading, the losses incurred by errors or failures in automated trading systems could grow very large very quickly and it's not impossible that such losses reach ruinous levels before they are even noticed by human operators. Of course, decent software should involve self-checks and fail-safes, but the point here is that even decent software sometimes isn't quite as infallibly decent as it was intended or hoped to be.

This is not mere speculative scaremongering. Already, for several years, automated trading systems have been engaged in autonomous high-frequency trading (primarily for risk management) in the global foreign exchange (FX) markets. Daily global FX transactions in the main financial markets dwarf those in equities: each day, well over $3trillion changes hands in FX transactions (daily global equity volumes are roughly one tenth of this).[12]

The massive liquidity, geographic dispersion, and round-the-clock continuous operation of the global FX market means that it is often cited as the real-world system closest to the idealization of perfect competition explored by economics theorists. Nevertheless, traders familiar with high-frequency automated FX trading systems are now well-accustomed to seeing occasional "glitches" in the market's exchange rates of major currencies: such events occur two or maybe three times a year. In such a case, an automated FX trading system somewhere will develop a problem and quote a wildly inaccurate price; if the price is a "giveaway", well below the true market rate, the rest of the global FX market participants (many of them running their own automated trading systems) rapidly switches to trading with the error-hit system, buying currency for cheap prices that cause significant losses for that system's parent institution. Not uncommonly, the speed and magnitude of the market's rush to exploit the mistake by one automated trading system means that the incorrect price being quoted by that system then affects, or even becomes, the *global* exchange rate for that pair of currencies, until such time as the owner of the error-hit system disables it.

Such "stuck prices" may last only for a few minutes, so they never make the evening news, but they can happen even between the currencies of G8 economies. These glitches are entirely artefacts of the increased reliance on technology: no single human trader could trade in sufficient volume at sufficient speed to go undetected. Thus far, participants in the global FX markets have merely absorbed the losses from stuck prices as part of the cost of doing business; it would seem that the potential for major catastrophe has been ruled out by the simple line of reasoning that no major catastrophe has yet occurred.

Similar reasoning, known as "normalisation of deviance" in the academic sociology literature dealing with studies of technology failure, led NASA to lose not one but two space shuttles.[13, 14]

*

For any and all of these potential developments in the financial markets, regulatory authorities need to be aware of their likelihood and their probable implications well in advance. It could be argued, with the benefit of 20:20 hindsight, that all the technology factors involved in the 2007/2008 crisis were readily identifiable by the mid-to-late 1990's: the ways in which those technology factors were deployed to support innovation in trading products, and the ways in which those innovations were a response to the existing regulatory framework, were also not so obtuse or unexpected to be reasonably described as wholly unforeseeable. People just didn't ask the right questions.

While some of the technologies or techniques were admittedly not of practicable use in 1997/1998, the strong likelihood that the "Moore's Law" ongoing exponential reduction in the cost of the necessary technologies meant that it was relatively trivial to predict that the necessary computing power, or network bandwidth, would become readily affordable within the following decade or so.

*

The proposal here then is for a Foresight project that explores the possible financial market large-scale complex socio-technical systems that seem likely on a 10-to-30-year horizon, in order that technology providers, market participants, regulators, and policy makers can all be informed and work together to avoid major problems like those experienced in the past two years.

The project proposed here would seek to combine the expertise and experience of financial-market practitioners, regulators, economists, financial mathematicians, organisational sociologists, and the computer scientists and engineers that drive the development of future banking technologies and structures. The UK is home to world-leading industry practitioners, academics, and experts in designing and enforcing financial regulatory systems. Such a project would be an investment by the British Government (of whatever political persuasion) in safeguarding the UK's position as a world-leading centre for innovation and expertise in the global financial markets, and also in safeguarding the UK economy (and indeed the wider global economy) against future unforeseen and unprecedented socio-technology-led instabilities and crises like that which arose over 2007-2008.[15]

Acknowledgements

The current version of this document has benefitted very much from comments provided by several anonymous reviewers, to all of whom I am very grateful.

Endnotes

---

1 See T. Nelson, *Computer Lib/ Dream Machines* (Microsoft Press, 1987; reprinted from a 1974 original).

2 Listed at www.foresight.gov.uk/About/Themes/Criteria_for_selection_as_a_foresight_project.asp.

3 The relevant UK academic centres and multi-site initiatives that could (or should) be involved in the Foresight project proposed here include:

- The Bristol Centre for Complexity Sciences: http://bccs.bristol.ac.uk/
- The Bristol/Bath Systems Engineering Doctoral Training Centre: http://www.bristol.ac.uk/engineering/systemscentre
- The Cambridge Centre for Financial Analysis and Policy: http://www.cfap.jbs.cam.ac.uk/
- The Centre for Financial Computing: http://www.c4fc.net/
- The UK Large-Scale Complex IT Systems (LSCITS) Research & Training Initiative: www.lscits.org
- The Oxford Centre for Collaborative Applied Mathematics: http://www.maths.ox.ac.uk/groups/occam
- The Southampton Institute for Complex Systems Simulation: http://www.icss.soton.ac.uk/
- The UCL Financial Computing Doctoral Training Centre: http://www.cs.ucl.ac.uk/financialcomputing/
- The Warwick Complexity Science Doctoral Training Centre: http://www2.warwick.ac.uk/fac/cross_fac/comcom/dtcsite/

4 A reviewer of an earlier version of this document argues that the (social) lack of understanding of systemic risk implications were in fact the primary problem. He/she writes:

> "The assumption that National Housing Indices in the USA could only fall by 2% in an extreme move (the underlying risk model of S&P and Moodys used for AAA accreditation) looks ridiculous in hindsight. It was widely believed because Nationwide House Prices had not fallen since the Great Depression. The implied loss potential of financial instruments based on such assumptions was vastly lower than the losses after the 33% fall which occurred (see: http://www2.standardandpoors.com/spf/pdf/media/rfc_rmbs_091009.pdf)."

I agree with all of that, but don't think it's the whole story, for reasons clear in the main text: I think the most plausible analysis of the recent crisis is neither that it was entirely technical in origin, nor that it was entirely social; rather, it was a consequence of interconnected/"interwingled" socio-technical issues. For discussion of a view of the future of the financial markets viewed as a component in a larger complex socio-ecological networked system, see http://www.sbs.ox.ac.uk/centres/insis/projects/Pages/financial-scenarios.aspx. And for detailed socio-technical analyses of the financial markets, see the various works of the Edinburgh Professor of Sociology Donald MacKenzie, including: *Do Economists Make Markets? On the Performativity of Economics* (Princeton University Press, 2007); *An Engine, Not a Camera: How Financial Models Shape Markets* (MIT Press, 2008); and *Material Markets: How Economic Agents are Constructed* (OUP, 2008).

5 See, for example, the following edited collections: P. Anderson, K. Arrow, & D. Pines (eds), *The Economy as an Evolving Complex System,* Addison-Wesley, 1989; B. Arthur, V. Morrison, S. Durlauf, & D. Lane (eds) *The Economy as an Evolving Complex System II,* Addison Wesley, 1997; N. Johnson, P. Jefferies, & P. Hui (eds) *Financial Market Complexity,* OUP, 2003; D. Challet, M. Marsili, & Y. Zhang (eds) *Minority Games: Interacting agents in financial markets,* OUP, 2004; and L. Blume & S. Durlaf (eds) *The Economy As an Evolving Complex System III: Current Perspectives and Future Directions,* Addison-Wesley, 2005.

6 See e.g. http://www.ft.com/cms/s/1/c84064da-1661-11da-8081-00000e2511c8.html

[7] For a longer discussion of this and related issues, see the briefing paper co-authored by IBM and The Economist Intelligence Unit, titled *The Trader Is Dead: Long Live The Trader*. Available from: http://www-935.ibm.com/services/us/index.wss/ibvstudy/imc/a1024121?cntxt=a1005266

[8] Current UAV/drone technology, where the aircraft is remotely operated by a "pilot" on the ground, has provoked a heated debate within the US armed forces: the US Air Force is insistent that weapon-carrying drones be operated by trained USAF pilots, i.e. individuals who have previously qualified to fly conventional human-piloted bomber and fighter aircraft; but US Army ground commanders argue that their troops' needs are best served by having the drones operated from command posts in theatre: see, e.g. http://articles.latimes.com/2008/mar/21/nation/na-predators21. Moreover, there are reports that significant numbers of USAF pilots complain that they are dissatisfied by drone duty: they joined up to fly in real aeroplanes, not to sit thousands of miles away from theatre and remotely operate flying robots via video links (drones currently flying in Iraq and Afghanistan are "flown" by human pilots working at video consoles in Las Vegas: see http://www.timesonline.co.uk/tol/life_and_style/men/article5944961.ece). And so, as the degree of automation and autonomy in the drones increases, the argument for having drones operated exclusively by USAF combat pilots becomes ever weaker, the demands on the trained USAF pilots becomes ever more trivial, and the ease with which their role can be replicated, and hence replaced, by ground troops becomes ever greater.

[9] An anonymous reviewer added this helpful example from the recent history of the UK/EU financial markets:

> "A UK example would be LIFFE (London Financial Futures Exchange) – they ignored electronic trading which was being developed by the Germans – as a result of which they went out of business (were in fact bought out by EuroNext.) In 1996 LIFFE was by far the biggest futures exchange in Europe having started in 1982. The French MATIF was second largest and the German DTB third (less than 25% of the market). In mid 2007 DTB had less than 25% of the DTB contract. After introducing electronic trading the DTB had more than 50% by October and by December LIFFE had less than 10% of the market. The French Matif was similarly wiped out and then absorbed into EuroNext. This forced a move from open outcry to electronic exchanges across Europe."

[10] A well-informed BBC Radio 4 programme in the *File on Four* series, broadcast in November 2009, covered this in some depth and included lengthy interviews with several senior figures involved in leading-edge high-frequency trading; the broadcast prompted press coverage for the comments made on it by HM Treasury minister, Lord Myners. A complete transcript of the 30-minute *File on Four* broadcast is available as pdf from here: http://news.bbc.co.uk/1/shared/bsp/hi/pdfs/20_10_09_fo4_casino.pdf and the BBC News story on Myners' comments is here: http://news.bbc.co.uk/1/hi/business/8338045.stm.

[11] See www.xmos.com

[12] The most recent (2007) triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity published by the Bank for International Settlements showed an average total daily dollar volume of $3,080bn in the global foreign exchange markets. For comparison, data for the New York Stock Exchange (NYSE) published by its parent company NYSE-Euronext shows that in January 2010 the total dollar volume for NYSE Group's exchanges had a daily average of $174bn (January total was $1,324.65bn, over 19 trading days); figures released by the London Stock Exchange (LSE) show average daily trade volumes on LSE Group's markets of £7.1bn ($11.1bn) for January 2010. NYSE-Euronext Group and LSE Group are the largest equity exchanges by volume in the US and Europe, respectively.

- www.nyxdata.com/nysedata/asp/factbook/viewer_edition.asp?mode=table&key=3133&category=3
- http://www.londonstockexchange.com/about-the-exchange/media-relations/press-releases/2010/monthlymarketreportjanuary2010.htm
- http://www.bis.org/publ/rpfxf07t.htm

[13] See: D. Vaughan, *The Challenger Launch Decision: Risky technology, culture, and deviance at NASA,* (University of Chicago Press, 1996); and W. Starbuck & M. Farjoun, eds, *Organisation at the Limit: Lessons from the Columbia disaster,* (Blackwell Publishing, 2005).

[14] An anonymous reviewer commented on this point at length, and it is useful to quote the comment verbatim here:

> "I would make an even stronger point: the stock market crash of 1987 (Hong Kong down 45%, UK down 26% etc.) was exacerbated by two technological issues: the first was the feedback loop of programme selling (automated selling as the market declined). In London the equity trading screens jammed due to overload (poor bandwidth as a result of poor design – each time a price was called for a single stock the system sent all stocks on a page). The result was a three-hour delay in prices which led to increased uncertainty and more selling.
>
> "As for quant funds running equity high frequency programmes – they lost an average of 20% over three days during the week of August 6th 2007. To quote Professor Andy Lo of MIT: "Based on empirical evidence from TASS hedge-fund data as well as simulated performance of a specific long-short equity strategy we hypothesize that the losses were initiated by the rapid unwinding of one or more sizable quantitative equity market-neutral portfolios. Given the speed and price impact with which this occurred it was probably the result of a sudden liquidation by a multi-strategy fund or proprietary trading desk possibly due to margin calls or a risk reduction. These initial losses then put pressure on a broader set of long/short and long-only equity portfolios causing further losses on Aug 9th by triggering stop loss and deleveraging policies." (A. Khandani & A. Lo "What happened to the Quants in August 2007?", working paper available from http://web.mit.edu/alo/www/articles.html). In fact it was one large portfolio liquidation that triggered the event.
>
> "So a blow-up of algorithmic trading has already occurred and furthermore it occurred due to the build-up of similar positions and (ironically) the risk management rules which dictated that as prices moved all other funds had to cut risk (by more selling of the same position) – see the paper by Philip Bond on crashes for more examples ("Phenomenology of Bubbles, Crashes and Extreme Events in Financial Markets", *Journal of Economic and Financial Modelling*, Spring 2008.)"

[15] [This note added April 2010:] In January 2010 the US Securities and Exchange Commission (SEC) released a 74-page public consultation document, *Concept Release on Equity Market Structure,* exploring issues in market structure and regulation. The first-page summary of that document reads as follows:

> "The Securities and Exchange Commission ("Commission") is conducting a broad review of the current equity market structure. The review includes an evaluation of equity market structure performance in recent years and an assessment of whether market structure rules have kept pace with, among other things, changes in trading technology and practices. To help further its review, the Commission is publishing this concept release to invite public comment on a wide range of market structure issues, including high frequency trading, order routing, market data linkages, and undisplayed, or "dark," liquidity. The Commission intends to use the public's comments to help determine whether regulatory initiatives to improve the current equity market structure are needed and, if so, the specific nature of such initiatives."

The document is free to download from http://www.sec.gov/rules/concept/2010/34-61358.pdf.