



Government
Office for

Science

 Foresight

An ecological perspective on the future of computer trading

**The Future of Computer Trading in Financial
Markets - Foresight Driver Review – DR 6**

Contents

2. The role of computers in modern markets	7
2.1 Computerization and trading	7
Execution.....	7
Investment research	7
Clearing and settlement	7
2.2 A coarse preliminary taxonomy of computer trading	7
Execution algos	8
Algorithmic trading	9
High Frequency Trading / Latency arbitrage	9
Algorithmic market making	9
3. Ecology and evolution as key principles for understanding financial markets and computer trading	12
3.1 Inherent limitations of equilibrium, efficiency and rationality.	12
3.2 The ecological view of markets.....	14
3.2.1 Framework	14
3.2.2 Example: Portfolio balancing and pairs trading	15
3.2.3 Markets do not always tend towards the ideal of increasing efficiency	18
3.2.4. Evolution towards efficiency can be slow	19
3.2.5 The path to efficiency may involve instabilities.....	20
3.4 Need for a large-scale market simulation of market ecology	21
4. Regulation for the microstructure ecology	21
4.1 Towards an evolutionary ecology view of competition policy.....	22
Computer trading poses unique challenges for regulators.....	23
4.2 Possible regulation to discourage market instabilities	25
4.2.1 Leverage and systemic risk in the ecology	25
4.2.3 Correlated strategies / crowded trades and systemic risk	26
4.2.4 Encouraging diversity for the sake of stability	26
4.2.5 Inability of the AT ecology to cope with unexpected events.....	27
4.2.6 Controlling financial innovation: new financial products as invasive species	28
4.2.7 Multiple equilibria in the ecology	28
5. On measuring the current ecology of computer trading.....	29
5.1 The need for measurements	29
5.2 On the type of data that is necessary	29

5.4 Early warning diagnostics for market crashes based on patterns in the market ecology	30
6. A long term view: How might the ecology evolve over the next decade?	30
7. Conclusion: How to prepare for the future	35
References	37

J. Doyne Farmer

Santa Fe Institute
jdf@santafe.edu

Spyros Skouras

Athens University of Economics and Business
skouras@aueb.gr

August 21, 2011

This review has been commissioned as part of the UK Government's Foresight Project, The Future of Computer Trading in Financial Markets. The views expressed do not represent the policy of any Government or organisation.

An ecological perspective on the future of computer trading

We adopt an ecological microstructure perspective of financial markets and use it to consider the impact of computer trading on the fairness, competitiveness and stability of markets now and in the future. The ecological perspective is particularly appropriate for computer trading because it has increased specialization and has made the behaviour and interactions of agents in the market much more systematic and measurable than when trading based on human judgment was dominant. We view regulatory policy as a key element of the ecology itself and a key driver shaping markets of the future. After reviewing a number of important trends in market ecology, we conclude that immediate regulatory initiative is necessary for careful measurement of market ecology with the related goals of (1) developing real-time warning signals for systemic risk; (2) informing an evidence-based approach to systemic stability policy and to competition policy that properly accounts for the multi-platform nature of modern financial markets; and (3) building a deeper theoretical understanding of markets based on large-scale simulations.

1. Introduction

Computer trading in financial markets is a natural and inevitable consequence of technological progress, and is almost as old as computers themselves. Computers facilitate basic market activities such as trade, liquidity provision, information dissemination, accounting, risk reduction and regulatory monitoring. They also facilitate the gathering, processing and analysis of historical data, provide the ability to rapidly and systemically monitor many data streams at once and execute trades automatically. These abilities enable the construction of algorithmic trading (AT) strategies that take advantage of market inefficiencies. Despite ever-increasing computer power, computer algorithms nonetheless remain far from fully intelligent, and from an economic point of view provide a clear example of what is called “bounded rationality”. Bounded rationality means that individual computer algorithms are necessarily highly specialized, which in turn means that to perform all the functions demanded of them, the universe of computer algorithms is necessarily highly diverse. Consequently the universe of computer algorithms is best understood as a complex ecology of highly specialized, highly diverse, and strongly interacting agents.

Algorithmic strategies have a natural life cycle – a given strategy’s performance tends to degrade with time as it becomes more widespread, reducing the inefficiencies that allowed it to exist in the first place. New strategies are continually being created in response to changing market conditions, while the propagation of successful strategies changes the market. Strategies, markets and regulations co-evolve in competitive, symbiotic or predator-prey relationships as technology and the economy change in the background. While there may be overall trends toward efficiency, new inefficiencies are created all the time. This dynamic process creates instabilities that drive regulatory changes, at the same time that regulatory changes create new inefficiencies and drive new strategies.

The role of computers in markets is thus an inherently dynamic process that can only be properly understood from an ecological and evolutionary perspective. The particular market ecology of computers that exists today is the outcome of many historical accidents, not just in the path of technological progress, but also specific low-level details in market design, regulation, habits, conventions and more broadly, patterns in human trading.

In this paper we synthesize evidence from disparate strands of research and also present new results, with the goal of articulating the ecological view of markets and presenting some of the insights it already offers. We present a vision of what might be possible in the future with a more vigorous program of data gathering, simulation, and basic research, and argue that such a program is essential if we are to understand the future of computer trading.

This paper is structured as follows. In Section 2 we discuss what computer trading is and its role in modern markets, leading to the discussion of our ecological perspective of financial markets in Section 3. In Section 4 we discuss insights from this approach that are relevant for regulation, in particular with the goals of promoting fair and competitive markets as well as ensuring systemic stability. In Section 5 we discuss the importance of properly measuring this ecology, some features known about the current ecology of computer trading strategies and how real-time measurements could help provide a warning system for incipient problems. In Section 6 we make some extrapolations of what we see as the broad trends in the market ecology before concluding with the implications of our work in Section 7.

2. The role of computers in modern markets

Computers are now completely integrated into modern markets. This section reviews how they have affected basic market functions, such as execution and clearing and settlement, how they have reduced personnel, and give an example of how regulatory changes affect participant behaviour. We also present a very tentative taxonomy of market activities using computers.

2.1 Computerization and trading

Execution

First, computers affect the trading process through their adoption by exchanges to automate the process of receiving orders and generating trades for investors. Studying this transition to computerized execution which started with the Toronto exchange in 1977, Jain (2005) finds huge benefits globally (in 101 countries) when equity exchanges made these changes, primarily because stock prices went up (by almost 30%!) as did liquidity, while the rate of return required by investors went down, making investment projects in these countries more attractive. Note that without the adoption of automated exchanges, algorithmic trading could never have grown to its current importance, which underlines the importance of symbiosis of various technologies in financial markets. Similarly, Hendershott and Moulton (2010) find that the adoption of a new technology by NYSE in 2006 that decreased execution times from 10 seconds to less than 1 second made prices more efficient (though it also raised bid-ask spreads).

Investment research

Second, they have had a huge impact on trading through their role in investment research and in the development of market monitoring and communication tools for traders while they have also played a role in research that led to the development of new financial instruments (including complex instruments associated with the subprime crisis of 2007).

Clearing and settlement

Third, automation in clearing and settlement now plays an important role in the trading process. For example, the NYSE was closed for 45 days between 1928 and 1933 to allow back offices to catch up on work. The fact that clearing and settlement still takes days on many markets adds a layer of counterparty risk which from discussions with market participants we understand can become an important consideration in times of extreme market turbulence. There are other consequences of computer trading that are equally important yet often overlooked. For example IBM (2006) estimates a 90% reduction of employment in trading related positions. The location of financial centres may also be affected by speed issues and the economies of market agglomeration. Long term trends in computer trading are discussed in Section 6.

2.2 A coarse preliminary taxonomy of computer trading

There are many reasonable schemes for classifying trading styles and market participants often use terminology in different, overlapping or conflicting ways. This widespread pressure towards classification in itself shows that thinking in terms of a taxonomy for trading styles is natural and useful. It also indicates the commercial pressure on money managers to pursue strategies with well-defined characteristics that conform to patterns that have become accepted as standard ("species" in the ecology). For example, the Lipper TASS hedge fund database classifies thousands of hedge funds into eleven categories: convertible arbitrage, dedicated short-bias, emerging markets, equity market neutral, event driven, fixed income arbitrage, global macro, long/short equity, managed future, multi-strategy and fund of funds.

Relatedly, the CS/Tremont indexes provide average returns of funds in each of these categories and three additional categories: event driven multi-strategy, distressed and risk arbitrage. These represent very coarse classifications of money management styles but much finer classifications are also possible.

The very fact that hedge funds are classifiable indicates that money managers cannot be pursuing fully optimized strategies since conforming to an investment style means restricted investment options which cannot add to performance in terms of risk-adjusted returns on invested capital. Rather their investment goals are skewed by marketing and other commercial considerations as well as their own abilities and expertise.

It is interesting to note that several dimensions of computer trading can be analyzed as taxonomies of diverse groups, including exchange venues (Domowitz, 1993) though we will not pursue this here, where we instead focus on a taxonomy of AT.

The considerations we discussed for hedge funds apply a fortiori to AT. AT systems usually exploit relatively narrowly defined market patterns and often require infrastructure that is specific to certain types of asset classes so they are even more well defined than hedge fund classifications. Below we discuss a very coarse and preliminary classification or taxonomy of AT that is relevant for the rest of our paper. This taxonomy is essential to begin measuring the ecology which we later argue is an essential and overdue first step. It is also necessary as any ecology driven simulation model would need to calibrate the population of these styles of investments as well as the qualitative features of the styles themselves. A more refined high resolution taxonomy would add further accuracy to a simulation model, but below we list the main algorithms that we consider essential in any simulation model for today's markets.

Execution algos

Execution algos have both a beneficial and a harmful role in price discovery. To the extent that they seek liquidity that is dispersed across markets and time they may be helpful for price discovery as they minimize temporary price fluctuations due to random fluctuations in liquidity. On the other hand, they attempt to hide the information of the order being executed by trading in a pattern that exploits any weaknesses in the market's price discovery process.

Furthermore, several extremely popular execution algos will most likely exacerbate short-term market trends and therefore may lead to greater volatility and more noisy prices. Widely used simple execution algos are usually relatively agnostic about what the 'correct' price for a trade is and instead place orders based on some other criterion with the implicit assumption that the market is performing price discovery reasonably. If a market becomes dominated by execution algos there is little reason to believe the market will perform price discovery effectively.

To illustrate, consider the three most popular styles of execution algos - volume weighted average price (VWAP), participation and passive algos. VWAP algos are designed to execute at prices close to the volume weighted price over some time interval and therefore when a large trade occurs at some price they will cause additional trading at that price reinforcing price trends (a large order of this type played a key role in the 'flash crash'). Participation algos attempt to trade in line with random variations in volume over some interval and therefore will reinforce volume trends; and passive algos attempt to trade at the best bid and ask prices available at any point in time with little weight on what the 'correct' price is.

Dark aggregators may also be viewed as types of execution algos which are market mechanisms the sole purpose of which is to hide information that is relevant for price discovery from other market participants.

Algorithmic trading

There exists a very broad range of investment strategies that are often implemented algorithmically. As examples, we will discuss pairs trading and statistical arbitrage at length in Section 3.2.2 (an up-to-date overview of such strategies can be found for example in Narang 2009). It is well-known in the algorithmic trading community that as various investment strategies have become automated and widely adopted, the time horizon over which they operate has decreased, their performance has declined and their leverage has increased. This is also visible from trending patterns of performance of algorithmic strategies over time - see Figure 1 and also Lim and Brooks (2011). It is natural therefore that Khandani and Lo found (2007) found that these kinds of investment strategies have become increasingly crowded.

Chaboud et al (2009) find that AT is beneficial in that it decreases volatility on the foreign exchange market for three currency pairs in 2006 and 2007 and Hendershott et al (2011) find that AT improves liquidity while Hendershott and Riordan (2009) find AT smooths fluctuations in liquidity over time.

It is worth noting that algorithmic investment strategies also involve a layer of execution algos (which however remain distinct usually even in their implementation as code). Below we discuss two types of AT which we consider particularly important for the market ecology and refer to elsewhere in the text, namely High Frequency Trading (HFT) and Algorithmic Market Making.

High Frequency Trading / Latency arbitrage

As discussed also earlier, HFT is a term used for strategies the performance of which is based on their speed. The ability to process speed fast is a major advantage because it means that effectively some market participants can trade on information that is not yet available to others.

The effect of HFT on markets is controversial. BMO Capital Markets (2009) argues that the effect of HFT on the Canadian Market has been detrimental and Zhang (2010) finds that HFT tends to increase volatility. Zhang and Powell (2011) suggest that HFT was responsible for the flash crash of 2010 and therefore of an adverse impact this has had on market confidence which they believe has been very large. Jarrow and Protter (2011) and Cartea and Penalva (2010) develop models in which HFT is generally harmful. On the other hand, Jovanovic and Menkveld (2010) find that a HFT market maker in Dutch stocks leads to a reduction of spreads by 29% while Brogaard (2010) finds that in the US equity markets HFT contributes to price discovery and is generally beneficial.

Tradeworx (2010) describes one situation in which HFT can - and in their experience does - exploit market microstructure imperfections to the detriment of other traders, suggesting a predator-prey relationship with other strategies.

In our view this mixed evidence is due to the fact that HFT is a very rich subset of strategies, some of which are beneficial for markets while others are not, with these effects being potentially dependent on other market conditions. It is misleading to lump all of HFT together and a detailed ecological approach is necessary. Skouras and Farmer (2011) discuss HFT in greater detail.

Algorithmic market making

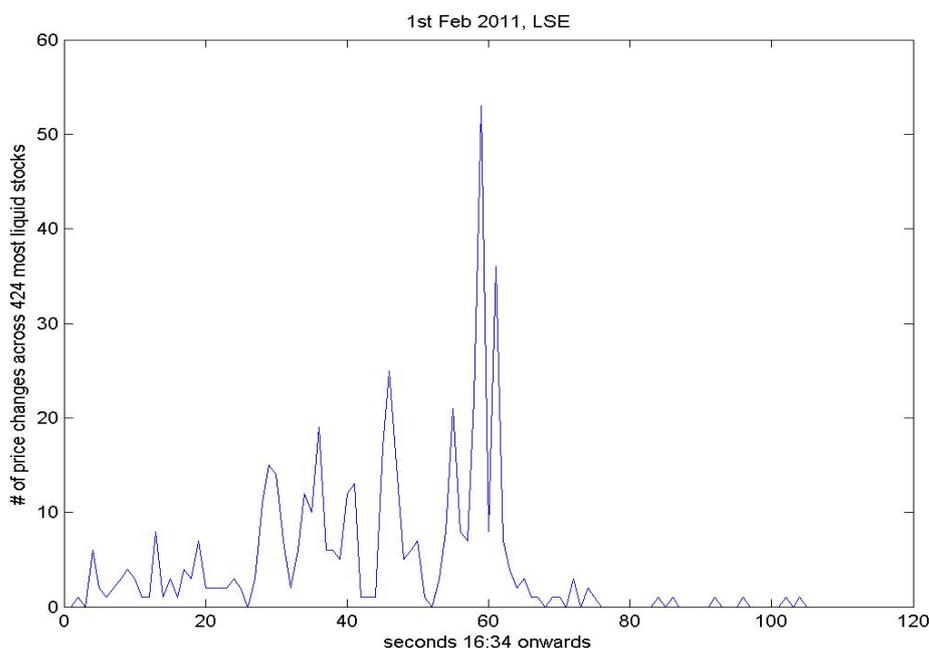
Market making has been around for centuries and indeed is practiced even outside of financial markets. The vanilla version of market making is illustrated nicely by banks which buy and sell currencies to travelers and other retail customers with the intention of selling, say, Euros to a UK customer travelling to Europe at a favourable rate relative to what they will charge when

selling Sterling to a French customer travelling to the UK. The difference is the spread and can be thought of as a fee for the liquidity the market maker provides.

We briefly discuss the history of market making because it is intimately tied with the history of execution venues and because it has played a key role in how the market ecology has evolved. It is a key part of the long term trends we also discuss in Section 6.

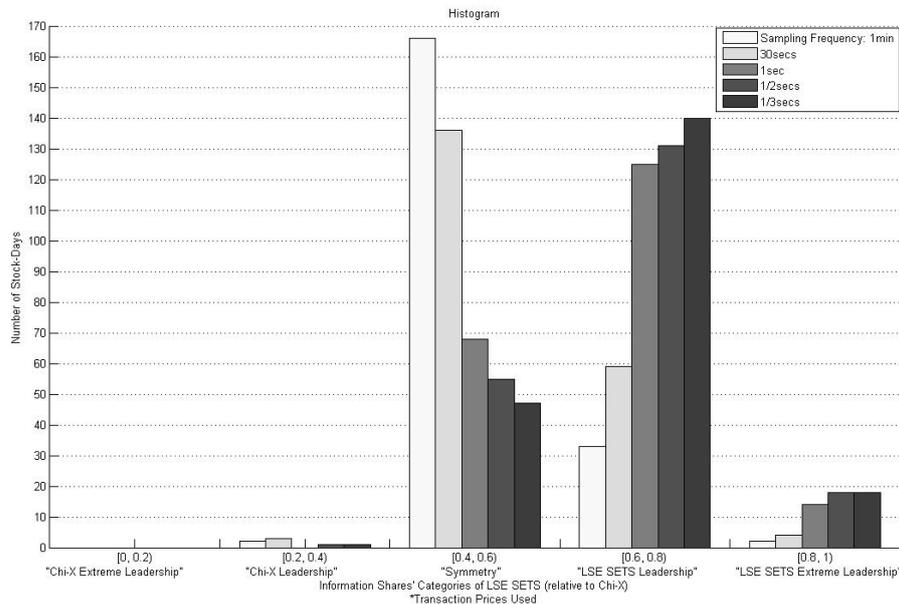
In the early days of market making, market makers were designated by the exchanges and assigned privileges and responsibilities with the intention that they would help make markets function better by providing more liquidity. As recently as a decade ago, the standard execution mechanism on NASDAQ and NYSE involved exchange appointed market makers as a counterparty to any transaction who specialized in making markets for particular instruments which they became experts in. As markets evolved and most exchanges switched to open electronic order book execution mechanisms, market making became a voluntary activity that could be pursued by any trader and also that could also be automated. Indeed, in Figure 4 below we see the characteristic market making pattern in the inventory of a particular brokerage id code for the London Stock Exchange's Stock Exchange Trading System (SETS).

Figure 4: This figure illustrates the fact that algos are carefully crafted to exploit obscure details of market design which often give a significant advantage to participants with lower latencies. On the London Stock Exchange (LSE) continuous trading stops at 16:30:00 after which new orders are accepted to participate in a closing auction that occurs at some random point in time between 16:35:00 and 16:35:30. At any point in time before that random moment, new orders are accepted and pre-existing orders can be modified or cancelled. After 16:30:00 and until the closing auction for each stock, the LSE disseminates in real-time (e.g. via Bloomberg terminals which was the source of the data in this graph) the "indicative prices" at which the auction would occur if no further orders, modifications or cancellations happened before the auction. The figure shows the number of stocks which experience a change in indicative price at each second between 16:34:00 and 16:35:30. Notice the spike at 16:34:59 (over 50 of the 424 most liquid LSE stocks experience a price change during this second) which is the latest time at which an order/modification/cancellation can occur and be certain to be in time for the auction. This suggests that algos are carefully designed to try to have a "last mover advantage" which however will ultimately belong to the fastest algo as this will be able to make decisions before the auction but with a more accurate indicative price than any slower algo. There is another interesting spike at 16:35:01 which presumably belongs to algos for which the last mover advantage is so great they are willing to risk not executing at all for the benefit of seeing all the price changes that occur up to 16:35:00.



While we cannot be sure that this market making pattern was driven by AT or of the size of the profits for this market maker who presumably was active in several stocks, we can present the audited track record of an algorithmic market making strategy that was employed by a company we co-founded (Figure 5).

Figure 5: This figure illustrates the fact that the sampling frequency at which empirical analyses are conducted can lead to very different conclusions. The large literature empirically analyzing inter-market price discovery has typically been conducted on a minute-by-minute time interval while to get satisfactory results in today's markets we need to use sub-second measurements which are not widely available. In the figure below, we plot Hasbrouck's (1995) "Information Share" measure for LSE relative to Chi-X for each stock and day covering the 13 LSE stocks that were members of the STOXX50 index during the 21 trading days of May 2009. The Information Share measure is a widely accepted measure of which market leads the other, or of the extent to which one of the markets has a prevailing influence over the other. If we try to measure which markets leads the other using data at the one minute interval as has typically been the convention in the literature, we find that on most days, most stocks have a "symmetric" information share, i.e. neither market prevails over the other. However, if we do the same analysis using data sampled every 1/3rd of a second we find that the LSE prevails with a measure in the range 60-80% in favour of LSE. In other words, LSE is the execution venue which dominates price discovery but its leadership over Chi-X is too fleeting to be measurable with the data traditionally used in these studies. Note that until very recently the LSE did not make data available time stamped at a resolution below the 1 second level and the data used here was provided to us by a market participant who collected and time stamped the data using proprietary infrastructure. For details see Axioglou (2010) who uses the same data and extensively analyzes Information Shares across European markets.



As single-instrument market making became more competitive, market makers began to search for more profit opportunities by trading the same or similar assets across execution venues and their activities. Indeed, Menkveld (2011) detects a high frequency market maker who earns close to Euros 1.5k per day with an annualized Sharpe Ratio of close to 10 making markets for around 14 stocks on Euronext and Chi-X and who seems to have played a key role in the success of Chi-X. While Menkveld (2011) calculates returns to capital based on some assumptions about the capital investment necessary to support such an operation, one interesting aspect of this kind of AT strategy is that because its inventories are very small, it has a close to zero marginal cost of deployment for a large institution that is already active in trading and has infrastructure adapted to computer trading. This is relevant for regulatory policy as discussed in Section 4.1.

3. Ecology and evolution as key principles for understanding financial markets and computer trading

3.1 Inherent limitations of equilibrium, efficiency and rationality.

The most effective approach to understanding financial markets has long been debated. For the last fifty years the dominant approach has been based on equilibrium and efficiency. This approach has the advantage of making predictions without having to impose very much structure, and has proved very useful for many purposes (see Farmer and Geanakoplos, 2009).

Nonetheless, this approach is limited. In this section we explain why it is essential for the goals of this study to go beyond theories based on equilibrium and efficiency, and in the next section we argue that to understand the future of computer trading it is absolutely essential to take a broader approach that focuses on market ecology and evolution. The terms equilibrium and efficiency are ambiguous in economics and indeed, some economists try to avoid the use of the terms altogether. For the purposes of our discussion, we use the term equilibrium models to describe workhouse economics models which neglect the study of dynamics under the perturbation of core model assumptions (such as market clearing or rationality); and the term efficiency will refer to *informational* efficiency or the property of many workhorse finance models that prices are largely (not necessarily entirely) unpredictable and that consequently investors - including those using computer trading - cannot differ in investment ability.

Under a common interpretation of the theory of efficient markets profitable trading strategies should not exist at all. There is an intrinsic paradox built into the assumption of efficient markets, originally articulated by Milton Friedman. On one hand, for markets to be efficient, there must be financial agents who detect inefficiencies and exploit them by trading on them, altering prices to make them more efficient. On the other hand, if the market is perfectly efficient there should be no profit-making opportunities to motivate such agents. Thus the market can never be perfectly efficient: There must be persistent residual inefficiencies to keep financial traders motivated to do their job of making the market more efficient.

The efficient markets hypothesis assumes that the residual inefficiencies that are left after this process takes place are small enough to be neglected. This is good enough for some purposes, but not for others. In particular, for the goals of this study, the theory of efficient markets is inadequate for understanding how markets change in time, and how instabilities can arise from the dynamic processes that drive the market away from equilibrium.

A physicist would restate the efficient markets hypothesis by saying that “markets are efficient at first order but can never be efficient at second order”. For problems where a first order analysis is adequate, it is useful to assume efficiency. For example, the theory predicts that we can’t all get rich by pursuing simple strategies, which is more or less true. For problems where a second order analysis is required, in contrast, the theory of market efficiency is by definition inadequate. It cannot describe the nature of the inefficiencies that actually exist, and it cannot describe how they change in time in a market that is not perfectly efficient. We need instead to focus our effort on understanding the nature of the inefficiencies that exist in real markets and how such inefficiencies evolve in time. We must understand why particular financial strategies exist and how diverse market activities, strategies and regulations interact to generate stable or unstable outcomes.

As we have already argued, computer trading is inherently specialized and the limitations of rationality are particularly clear. Present day computer trading strategies implement fixed algorithms, hard-wired rules of thumb that take specific actions on the receipt of given information. Because these are hard-wired, they cannot change their behaviour and therefore they cannot behave “rationally” when they encounter a situation that their programmer did not anticipate. Such rules of thumb tend to be highly specialized. They evolve slowly, on timescales from years to decades, as researchers refine them and improve them, with occasional jumps when genuinely new things are discovered. As discussed below, specialization makes the ecological view essential, and the fact that the underlying strategies change through a process of trial and error naturally suggests an evolutionary view. In the future one imagines that computer trading will employ more sophisticated “artificial intelligence” algorithms, which might make rationality a better approximation. This is still a long way off. Even humans, who are more intelligent than computer algorithms will be for a long time, often exhibit irrational behaviour. It seems unlikely that the artificially intelligent algorithms of the future will do even as well as we do. Rationality will be slow in coming to the world of computer trading.

The efficient markets view is essentially a black box theory. It does not analyze the precise mechanisms through which efficiency or other core principles such as the lack of arbitrage opportunities are attained; instead, their existence is simply postulated. To address the kind of questions raised by this project, which are inherently questions about the deviations from fully efficient and stable markets, one needs to go outside of this framework and analyze these problems in evolutionary terms. The ecological view attempts to look inside the black box and study the elements that may or may not lead to an efficient market.

To understand the present effect of computer trading strategies on the market, not to mention the future evolution of such strategies, it is necessary to understand why computer trading strategies exist in the first place, and furthermore, why the particular set of computer trading strategies that we observe in the market exists. In addition, one must understand how these strategies interact with one another, and how these interactions are changing through time. The efficient markets perspective is ill-suited for this purpose, and despite having dominated finance for several decades, has not demonstrated much success in handling this kind of phenomena. Instead, a proper understanding of market ecology and evolution is required.

The ecological and evolutionary approaches are formulated from the outset in terms of deviations from efficiency, and while they must always be consistent with efficient markets at first order, they can also explain more precisely the nature of the second order deviations that are critical for understanding real markets. The ecological approach is well-developed in biology, indeed ecology is an entire field into itself. Furthermore, the field of ecology has major regulatory impact; for example, major public works projects in the United States require

environmental impact statements that must include an ecological analysis of long-range effects. In finance, in contrast, this field is hardly in its infancy.

3.2 The ecological view of markets

3.2.1 Framework

The persistent inefficiency of markets motivates the ecological view, as developed by Farmer (2002). Under this view inefficiencies support an ecology of interacting financial traders. Because of bounded rationality, financial traders can't do everything at once – they tend to specialize. These specialized traders interact with one another as they perform basic functions, such as getting liquidity, offsetting risks and making profits. A given activity can produce profits or losses for another activity. Inefficiencies play the role of food in biological systems, providing profit-making possibilities that support the ecology.

In ecology a central notion is that of a food web, which provides a method of understanding the interactions of different species. The food web describes who eats whom. Similarly, in markets a financial food web describes who makes profits from whom. A precise definition of financial food webs was given by Farmer (2002). This is formulated in terms of the gain matrix g , which describes how the size for a given trading activity affects the returns of another trading activity. Let r_A be the average return of activity A and x_B be the size of activity B. Size refers to the capital associated with the activity, and can be defined in any of several ways; for example, it can be the associated trading volume. The gain of activity A due to activity B is

$$g_{AB} = \frac{\partial r_A}{\partial x_B}$$

Two trading activities A and B have a competitive relationship if $g_{AB} < 0$ and $g_{BA} < 0$, i.e. if increasing the size of A decreases the returns of B and increasing the size of B decreases the returns of A. They have a predator-prey relationship, with A preying on B, if $g_{AB} > 0$ and $g_{BA} < 0$, i.e. if increasing the size of B increases the returns of A and increasing the size of A decreases the returns of B. And they have a symbiotic relationship if $g_{AB} > 0$ and $g_{BA} > 0$, i.e. if increasing A increases the returns of B and increasing B increases the returns of A. Of course these relationships may be context dependent – in a market ecology with a third activity C, the size of C may effect the relationship of A and B.

In biology food webs depend on basal species that gather energy either from the sun or from thermal vents in the deep sea. Other species eat the basal species, and still other species eat them in turn, in what can be a highly complicated set of relationships involving hundreds of species. Similarly, financial food webs depend on economic activities that create inefficiencies. These inefficiencies support trading strategies that make profits, which in turn may create other inefficiencies and support other strategies. In biology, predators play valuable roles in regulating and maintaining balance among the populations of the species they prey on. The introduction of new species can disrupt this balance, destabilizing an ecology and causing irreversible changes (such as the introduction of rabbits into Australia). Similarly in markets many profit making strategies play an essential role by incorporating information into prices. There is no rule, however, that say that all trading strategies are beneficial; some strategies may make substantial profits while having little or no social value. Even strategies that normally add social value can be destabilizing, particularly when combined with other factors, such as leverage (Thurner, 2009).

As traders become more skilful, over the long term the market should tend to become more efficient. There may be circumstances, however, in which structural changes such as the

introduction of new technology or regulatory changes make the market more inefficient. Of course, as we know from the Friedman paradox, over the long-term there must always be persistent inefficiencies, but the size and the nature of these inefficiencies may vary considerably through time as the market responds to outside inputs and as the market endogenously evolves.

By using idealized representations of simple strategies, such as trend following or value investing, it is possible to compute the gain matrix and map out simple financial food webs by hand. For example, if one considers a population of trend following strategies, it is easy to see that short term strategies make profits at the expense of longer term strategies, while the profits of longer term strategies are lowered by the presence of a short term strategies. Thus trend strategies with different time scales have a predator-prey relationship, with short term strategies preying on long term strategies. Trend strategies with similar time scales, in contrast, have competitive relationships. Similarly fundamentalist strategies that respond to small deviations from fundamentals have a predator-prey relationship to those that respond to large deviations from fundamentals.

By making several approximations it is also possible to derive an analog (Farmer, 2002) of the famous Lotka-Volterra equations of population biology, which describe the circumstances under which populations are stable vs. when they oscillate. A famous example of oscillation involves rabbits and foxes: If there is an overabundance of rabbits the fox population increases, but it tends to overshoot, causing the population of rabbits to plummet until foxes begin to starve, which in turn allows the population of rabbits to rise again, completing the cycle.

In markets the population of a species is replaced by the capital deployed in a trading activity. The theory of efficient markets assumes that if any inefficiency emerges, exactly the right amount of capital will be deployed in strategies that exploit the inefficiency as needed to effectively remove it. However, if investors overshoot, and continue to invest in a strategy even after it ceases to be profitable, or pull too much money out when it becomes unprofitable, there can be persistent oscillations. Such oscillations are unlikely to be strictly periodic, but may be quite sensitive to random inputs. Thus, as long as conditions are favourable, a strategy may persist and even thrive, but when the market comes under stress there may be massive die-outs. We give an example in a moment.

3.2.2 Example: Portfolio balancing and pairs trading

A good example of a market activity that generates inefficiencies is portfolio balancing. Mutual funds, for example, account for more than four trillion dollars of stock market assets in the US alone. To diversify risks and to conform to their mandate, mutual funds strive to maintain target portfolio balance, i.e. they decide on a given set of weights w_i for each stock in their portfolio. These weights are measured in monetary terms. As prices fluctuate the portfolio goes out of balance. To rebalance the portfolio the manager must sell the stocks that went up in price and buy the stocks that went down in price. The resulting market impact or price pressure causes a price reversal – the stocks that went up go back down and the stocks that went down go back up. Thus, portfolio balancing creates an inefficiency in the form of predictable price reversals (Cont 2009, Marsili XXX) and we can therefore think of it as a kind of basal species off which other strategies will feed.

The classic strategy that exploits reversals is known as pairs trading, and is one of a family of strategies that go under the heading of statistical arbitrage. Pairs traders monitor stocks that tend to move together under normal circumstances, and buy stocks that have recently fallen in price while shorting stocks that have recently risen in price. They thus front-run the market impact of the portfolio rebalancing and exacerbate this impact: knowing that the rebalancers will soon need to buy in order to rebalance and will drive up the price, they buy before them

and sell back their inventories before the market impact dissipates. They thus make the market even less efficient by amplifying arbitrary temporary price deviations. Pairs trading clearly benefits from portfolio balancing, but not vice versa: More portfolio balancing activity implies bigger returns for pairs trading, but more pairs trading degrades the performance of mutual funds. This is thus an example of a predator-prey relationship. Generalized pairs trading strategies do not simply look at individual pairs of stocks, but group stocks into clusters based on industrial classification, response to interest rates or oil prices, etc., and then build sophisticated regression models to predict the deviations of stocks from the clusters they are members of.

Of course, it is also possible that when two similar stocks diverge in price, this is just a random effect and a deviation from fundamental valuation. In this case pairs trading is also preying on those who randomly cause the prices to diverge, but may have been beneficial in pushing prices back to their proper values. Thus, pairs trading may be good or bad for market efficiency depending on the situation.

To execute a generalized pairs trading strategy and make acceptable returns it is usually necessary to use leverage. The reason is that for the average stock about 50% of the return variance is explained by the movement of the market as a whole and about 25% is explained by the movement of the industry, leaving only 25% of the variance as idiosyncratic to the individual stock. Pairs strategies hedge out the market and the industry (and often other risk factors as well), betting only on the idiosyncratic component. The returns and volatility are correspondingly smaller, so to make profits that are comparable or better than those of the stock market as a whole, most investors use leverage. The use of leverage is typical of many computer strategies, and creates instabilities that we will discuss in more detail later.

The pairs trading strategy has been known at least since the 1920's, but it is difficult to implement manually on large scale due to the vast quantity of information that needs to be processed to find the best trades. It was not able to fully realize its potential until the 1980s when computer trading began to be widespread, allowing automatic backtesting, data monitoring, and execution of generalized pairs strategies. Because of the vast amount of information that must be processed, such models tend to be highly automated and make heavy use of computers.

From the theory of market efficiency one would expect that pairs traders should rapidly mop up all the profit potential generated by portfolio balancing, quickly making pairs trading a marginally profitable business. But while the market does indeed become more efficient, it is very slow to do so, as illustrated in Figure 1.

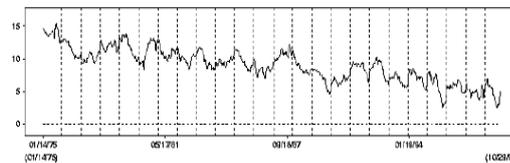
Figure 1: The strength of predictive signals and associated strategies evolves over time. This is illustrated with two signals developed by Prediction Company which are trending in opposite directions in the figure below. The first signal becomes progressively weaker as markets become more efficient while the second signal appears because of a regulatory change and becomes stronger as participants respond to this change.



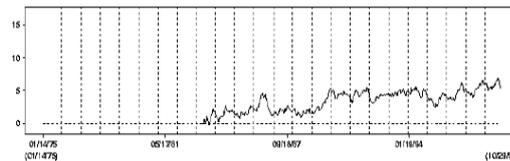
Market efficiency?

Strength of two proprietary predictive signals (1975 - 1998), (measured as smoothed average % correlation between signal and future weekly return)

Signal 1:



Signal 2:



■ Prediction Company ■

This figure measures the performance of a generalized pairs trading strategy (labeled “signal 1”) employed by Prediction Company, where one of the authors was Chief Scientist, as it was instantiated in 1998.¹ The model tracks the behaviour of about 1000 stocks, and makes daily predictions for the residual price deviations of each stock from its industry group. These predictions are tested by correlating against the future price deviation two weeks ahead. This is repeated every day for a period of 23 years, making a total of roughly 6 million predictions. The performance is measured by the correlation of the predictions with the future price movements, which is smoothed in the figure to make its behaviour more apparent.

The behaviour does not conform to the strong version of the efficient market theory. The correlation to future price residuals declines in a fairly steady way, starting at roughly 15% in 1975, and dropping to roughly 5% by 1998. The market does indeed become more efficient, but it does so very slowly. The strategy was still quite profitable in 1998, and extensions of this strategy remained profitable throughout most of the remaining decade.

However, pairs trading took a big hit in August of 2007, as explained by Khandani and Lo (2007). The underlying reason was understood by several funds (such as Prediction Company) in real-time as this problem occurred. It was driven by the economic crisis: As many institutions

¹ The algorithm used to implement the strategy was developed based on data from 1991 – 1998; the remaining data for 1975 – 1990 was acquired after the algorithm was already developed. The simulation proceeds by fitting the parameters of the algorithm in a given year and testing on data for the subsequent year, and repeating this process.

came under stress due to excessive exposure to subprime mortgages, they needed to raise cash, and began liquidating assets. Stat-arb strategies typically get in and out of positions on timescales ranging from a few days to a month, and are thus inherently liquid in normal circumstances: They can normally unwind their positions within a few days to a few weeks. Several large firms ordered their stat arb groups to liquidate. This generated unusual relative movements in the market; if stat-arb funds tended to be long on a particular stock, as these funds liquidated their position the price of this stock would be driven down, and if they tended to be short, its price would be driven up. This caused consistent losses by other stat-arb funds, many or most of whom had correlated positions, because they tended to be long on the stocks whose prices were going down and short on the stocks whose prices were going up. These losses in turn caused many hedge funds to reduce or liquidate their positions entirely, creating a snowball effect. The result was a massive die out of stat-arb funds. Based on anecdotal evidence, the funds that managed to hang on through the crisis sprang back and made enormous profits as soon as the die-out was over.

This story illustrates several things. The fact that a subprime mortgage crisis caused a mass extinction in an entirely unrelated asset class and family of trading strategies illustrates how interconnected the trading world has become. The fact that a trigger by a subset of firms to liquidate their stat-arb funds could cause such a massive die-out across this entire sub-industry was a wake-up call for the lack of diversity in the stat-arb universe, which accounts for a large fraction of computer trading. There were many funds (probably thousands) employing very similar trading strategies that learnt very quickly and very painfully how close their strategies were to those of their competitors.

In a sense the slow decline in performance for signal 1 observed in the top panel of figure 1 supports the arguments underlying market efficiency. While the path is somewhat bumpy, there is a clear overall trend for the market to become more efficient. However, the timescale is extremely slow – the approach to efficiency took decades, not years, and enormous profits were made in the meantime. In 1990 a handful of funds were pursuing stat-arb strategies, but by 2007 there were thousands of such funds. The mass extinction of 2007 eliminated a large number of funds, but hundreds to thousands of such funds still remain in business. As we will discuss later, a simple theoretical argument based on evolutionary reasoning suggests that such a slow approach to efficiency is to be expected.

3.2.3 Markets do not always tend towards the ideal of increasing efficiency

The second panel of figure 1 directly contradicts the idea that markets always become increasingly efficient over time. In this case we see a strategy (signal 2) that begins to be traded in 1983. The performance is initially rather poor – while the correlations to future price movements are positive, they are small, but during the course of the next decade they grow to roughly 5%, comparable to the generalized pairs strategy of signal 1. Thus the market does not become more efficient with time - quite the opposite - it becomes less efficient. (In fact in the following decade the profits of this strategy did eventually begin to drop as it became more widely known).

We do not have as good an understanding of the ecological relationships that support this strategy as we do for pairs trading, but what we do know is that it was initially enabled by a regulatory change regarding the disclosure of information (previously the strategy could not be implemented because the information it requires was not publicly available). We hypothesize that after this change, investors began processing this information and trading on it, but at the outset (for the subsequent two decades) they did so imperfectly, creating a strong signal for more sophisticated investors to exploit. The graph makes it clear that the regulatory change alone was not sufficient since the strategy also required time for the mass response to the

change to develop. It also illustrates how regulation can create changes that fundamentally alter the market ecology and its financial food web.

3.2.4. Evolution towards efficiency can be slow

On what timescale should we expect computer trading strategies to evolve? This can be estimated by a simple argument originally presented by Farmer (2002). The basic idea is that in order to discover a new trading strategy one must first detect an inefficiency. If the strategy is based on induction from past data,² enough data must be accumulated for the result to be statistically significant. If the inefficiency is really strong, the pattern may stand out from the noise so clearly that it will be apparent with little data, but if the inefficiency is weak, it will be less apparent and will take more data to discover, which means it will take more time.

One can easily make a crude estimate of the timescale to approach efficiency from a back of the envelope argument. Assume that, for whatever reason, an inefficiency suddenly appears. How long will it take to detect? We will measure the strength of an inefficiency based on the Sharpe ratio S , of an optimal strategy for exploiting it. The Sharpe ratio is the ratio of the expected return to the standard deviation of the returns of the strategy; for convenience we will measure it at an annual time scale. Under the idealized assumption that returns from the strategy are log-normally distributed, the average statistical confidence Σ that one expects to have in the strategy after accumulating data for t years is $\Sigma = S\sqrt{t}$, where Σ is measured in standard deviations. To achieve a 95% confidence level $\Sigma = 2$, so the time required to detect the inefficiency is $t = 4S^{-2}$. This estimate should be doubled, since one must first detect the inefficiency, then one needs to convince investors that the inefficiency is real through a live trading test (which takes the same length of time to achieve significance as it does to detect the strategy to begin with). Thus the minimum length of time required for an inefficiency to be found and attract sizeable investment capital to strategies that exploit it is:

$$t = \frac{8}{S^2}.$$

An inefficiency with an associated Sharpe ratio of 1 (a typical figure for successful investment strategies) should be expected to persist for about eight years before investors exploit it sufficiently to begin removing it. We should stress that this estimate is based on idealized assumptions, such as log-normal returns – heavy tails, autocorrelations, and other effects will tend to make the timescale even longer.

As a given inefficiency is exploited, it will become weaker and the Sharpe ratio of investment strategies associated with it drops. As the Sharpe ratio becomes smaller the fluctuations in its returns become bigger, which can generate uncertainty about whether or not the strategy is still viable. This slows down the approach to inefficiency even more. Within the framework of a highly idealized model developed by Cherkashin, Farmer and Lloyd (2009), it is possible to show that the approach to efficiency goes as $t^{-\alpha}$, where $0 \leq \alpha \leq 1$. This implies very slow convergence indeed: The trajectory is not just a power law, it is a power law with a very low exponent. The power laws comes from exactly the effect discussed above: The more the inefficiency is exploited, the weaker it becomes, the smaller the profits that can be made from it, the slower the capital deploying it grows, and the slower it evolves toward perfect efficiency.

² If an inefficiency can be predicted based on deduction, e.g. by a fundamental analysis of market conditions that does not depend on accumulating data, then of course the limits derived here do not apply. In general, though, even if one has an hypothesis based on deductive arguments, such hypotheses need to be tested, so a similar argument applies.

Interestingly, within this model the parameter α is related to the properties of a “reality map”, which relates how much real payoffs are influenced by the players’ predictions of the payoffs.³ When there is little influence of players’ perceptions on payoffs, convergence is fastest, i.e. $\alpha \approx 1$, but when real payoffs are determined entirely by perceptions, $\alpha = 0$, i.e. there is no convergence at all. This is interesting because it suggests that insofar as markets are subject to Keynesian beauty contests, the approach to efficiency is even slower. Thus assets whose value is tangible and where prices are more obviously connected to fundamentals should have less persistent inefficiencies than those whose value is more speculative (think utility companies vs. internet stocks in the 1990’s).

Because computer strategies tend to have high Sharpe ratios, they tend to evolve more quickly than other strategies. Thus we expect that in general evolution due to computer strategies will proceed more rapidly than the evolution of other strategies. Most computer trading strategies are driven by induction rather than fundamental analysis, so the analysis above applies very cleanly.

3.2.5 The path to efficiency may involve instabilities

The path toward efficiency can be indirect, with substantial instabilities generated along the way. Achieving efficiency often requires interactions between several different specialized strategies. For example, consider the strategy of exploiting an imbalance in the order book. If the order book has a preponderance of buy limit orders as opposed to sell limit orders, then without other compensating effects the price is more likely to go up than down. Such a strategy is used by many high frequency traders, who typically use this as one of many signals. According to efficient market theory one should expect that the deployment of such strategies should remove the inefficiency. But such strategies on their own cannot remove the inefficiency. For example, suppose one detects an excess of buy vs. sell orders, and places a buy market order. This removes yet another order on the sell side, making the order book imbalance even worse. While it is true that the profits from such strategies are self-limiting, due to the resulting widening of the spread, the presence of exploiting strategies does not fix the imbalance in the market – one is left with wide spreads and an even more imbalanced book. The solution is that the widened spreads create opportunities for market makers, who place orders on both sides of the book, restoring balance.

The point of the story is that maintaining an efficient, liquid, stable market often requires coordinated action between several different types of market strategies. One needs to think about the whole ecology, and cannot simply consider one component in isolation. It is possible that orderbook imbalance strategies of the type described above were involved in the flash crash, and market making strategies are now known to have been involved. This illustrates how the approach to efficiency requires coordinated changes in an ecosystem of specialized strategies, and how when there is a lack of strategy diversity instabilities can develop. One wants an ecosystem with several different types of strategies in their proper ecological balance.

³ The model is one of betting on a biased coin. The reality map is a function mapping the aggregate of the players’ bets onto the bias of the coin. A self-reinforcing reality map makes heads more likely if the players bet on heads, and tails more likely if they bet on tails; an objective reality map is one in which the bias is independent of the players’ bets. As the successful players accumulate wealth the game becomes more efficient, in the sense that a rational player who knows the reality map and the strategies of the players can extract smaller returns from this knowledge.

3.4 Need for a large-scale market simulation of market ecology

Several authors have argued for a simulation approach for studying financial markets, especially in the presence of computer trading. As early as 1989, Kim and Markowitz (1989) argued that a simulation approach could be used to explain the role portfolio insurance seems to have played in the 1987 crash and to help diagnose sources of instabilities in advance of them producing market crashes. Another driver review (Cliff et al, 2011) articulates a vision for large scale market simulations, so we will not repeat the case for simulations.

However, we do wish to emphasize the separate point that in our opinion an evolutionary ecology approach is key to the success of any simulation effort. There are many ways to build simulation models, but based on the discussion in this paper, we believe the most fruitful approach would involve an evolutionary ecology perspective.

In particular, simulation models should contain diverse types of traders (see our taxonomy in Section 2.2), the behaviors, populations, leverage and interactions of which would be calibrated using real data. With such models it becomes possible to inject new trading strategies into the ecology and study their effect. In particular one can identify the procyclical feedback loops that come from various activities, such as leverage or derivatives. One can also hope to use this approach to answer important questions that are difficult or impossible to address using traditional tools such as:

1. Might the physical location of markets change?
2. Do volume based exchange fees harm market competition?
3. Would taxing transactions or banning short-selling be beneficial or harmful?
4. Are there issues of tacit collusion between exchanges and market participants?
5. How can we prevent competition in terms of trading speed from degenerating into a wasteful and possibly harmful arms race?
6. Which strategies should be encouraged (because they provide valuable market functions, such as price discovery), and which should be discouraged (e.g. because they are destabilizing)? How can regulatory tools be used to do this?

The work that has been done so far, and the arguments developed here, barely begin to develop the concepts necessary to build the theoretical framework necessary to support such a simulation model. We believe more theoretical work needs to be done along these lines if we want to understand computer trading and that theory and simulations will inform each other.

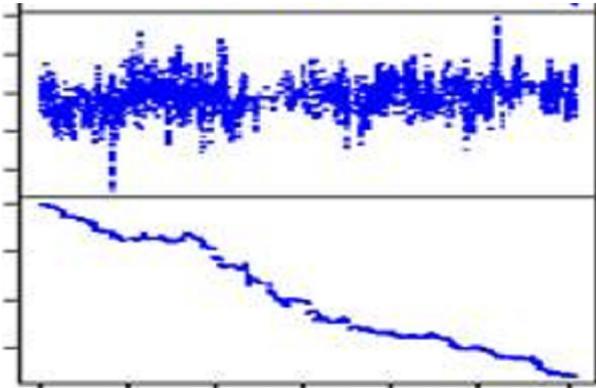
4. Regulation for the microstructure ecology

In this section we turn our attention to the role of regulation in the microstructure ecology. We examine first the role competition policy might play to ensure more competitive and fair markets before turning our attention to the role of policy in containing systemic risk. From the ecology perspective, regulators are key agents driving change and the future of markets will be influenced greatly by current policy decisions (or lack thereof).

An interesting example of how market microstructure can affect participant behaviour is illustrated in Figure 2. There we see algorithms competing for the "last-mover advantage" - they compete in terms of latency to send their orders as close as possible to the last trade (the auction close) on stocks traded on the London Stock Exchange. This illustrates how competition among market participants evolves to exploit any opportunities created by details of market design, and show how important the details of the execution protocols can be in influencing the market ecology.

Figure 2: Inventories of two of the ten most active members of the London Stock Exchange in October 2002 for a single actively traded stock (Astrazeneca). This illustrates one participant selling at a more or less constant daily rate over the entire month

(bottom graph) juxtaposed against the behaviour of a market maker who flips the sign of his inventory many times each day (top graph).



4.1 Towards an evolutionary ecology view of competition policy

The extreme heterogeneity and the well-defined nature of AT strategies interacting in the microstructure ecology suggests it is valuable to aim to create "multi-sided platforms". A multi-sided platform coordinates participants in some way that is beneficial to them because of their heterogeneity, a classic example being dating sites where men meet women who are charged different prices for access to the site to ensure a balanced population -- paying more is advantageous for men because women tend to use such sites less and it allows them to meet more women. Similarly, exchanges have customers that include listed firms, brokers and investors. Investors - and especially those employing computer trading - deploy trading strategies which themselves have multi-sided features as mentioned earlier when discussing symbiosis: liquidity providers need liquidity demanders, execution algos that ignore fundamental information need algos which ensure prices do not drift far from fundamentals, and of course, more generally buyers need sellers. Unfortunately the theory of competition policy for two-sided markets is in its infancy while that for multi-sided platform markets is at the very cutting edge of research (e.g. Evans, 2003). It is now being appreciated that this view can overhaul how industrial organization economists think about even very narrowly-defined aspects of market microstructure such as competition between exchanges (Cantillon and Yin, 2010).

Computer trading has fundamentally changed proprietary trading at major banks. This change means that the way economists have traditionally thought about it as a commercial activity is now often outdated. Traditionally, economists have analyzed prop trading in terms of a managers' skill in outperforming a benchmark asset. However, this is an inappropriate way to think of AT and especially HFT trading operations because these often provide much more steady returns than typical investments. At the same time they are not very scalable and require huge technological investments. In this sense, such operations have business models much closer to technology firms than financial institutions.

While developing new tools as the foundation of an appropriate competition policy for the microstructure ecology involves significant effort, there is clear evidence that enlightened policy in this area can be very beneficial. There is evidence in the academic literature that systematically anticompetitive behaviors can arise to exploit fine details of market microstructure and trading (Christie and Schultz, 1994). There is also evidence that policy can have very positive effects, e.g. Easley, Hendershott and Ramadorai (2009) study a natural experiment in which market access was changed to significantly equalize speed access of on and off-floor traders on NYSE and found that it led to an increase in prices of 3%, suggesting

that this kind of issue can have a huge effect on the behaviour of markets. Easley, O'Hara and Yang (2011) use a theoretical model to argue against practices which allow differential access to information (including speed advantages to HFT, an issue which seems related to the Net Neutrality debate relating to whether homogenous access to the Internet should be enforced by regulation - see e.g. Economides, 2008). That equal access to markets is viewed as a very important consideration is underlined by the fact that co-location providers make every effort to ensure that an identical service is provided to all co-location customers (e.g. the length of wires from market servers to all customers' co-located servers is the same regardless of where the servers happen to be located in the market data centre).

Computer trading poses unique challenges for regulators

Below we list some market features that should give the reader a flavour of the huge scope for anticompetitive and unfair practices in existence today and the challenges faced by regulators in disentangling complex effects.

- Conglomerate effects in large financial institutions: Conglomerate effects arise when a firm benefits from interactions of units that are neither horizontally nor vertically integrated and there exist EU guidelines, e.g. as to what conglomerate mergers should not be allowed. As an example of conglomerate effects in existing financial institutions, consider the fact that the prop trading desks of large investment banks have certain opaque advantages over independent trading desks. One obvious example is that a prop trading desk will be charged exchange fees that benefit from a volume rebate that is calculated at the level of the entire bank so that prop trading benefits from customer order flow generated at the brokerage unit. Indeed, it is also well known that execution algos that are provided by brokerage desks to customers are often the old generation of algos initially developed by prop trading desks and then abandoned as improved versions were developed. Inevitably, prop trading desks at investment banks will have a good understanding of how client flow is being executed and of any deficiencies in these algos that might be exploited. Finally, consider that the huge investment necessary to provide brokerage services by investment banks can usually be leveraged by the prop trading desks of the same banks to create AT operations at very little additional cost, giving them huge advantages relative to independent prop trading operations. Note that the Volcker rule may lead to a separation of prop trading and investment banking in the US but its logic is based on a different rationale - that of systemic risk (discussed in the section below).
- Vertical integration: When the markets are owned and controlled by the same interests that drive the vast majority of trading, anticompetitive pressures will be at work. For example, Chi-X which has the largest trading volume among all European exchanges was established in 2007 by Instinet which is a subsidiary of Nomura Holdings and is now owned by a private consortium also including BNP Paribas, Citadel, Citigroup, Credit Suisse, Fortis, GETCO Europe Ltd, Goldman Sachs, Merrill Lynch, Morgan Stanley, Optiver, Société Générale and UBS. Menkveld (2011) suggests a symbiotic relationship between Chi-X and a single large HFT market maker which we strongly suspect is also one of the owners. While Chi-X has effectively exerted competitive pressure on incumbent exchanges, we are not sure its business model would have been as successful had it not been owned by its own main customers. This example illustrates, that while competition may exist between various dominant market participants, fair opportunities for outsiders to enter may be lacking.
- Monopsony: A single HFT customer can be responsible for as much as half the business of an exchange (see e.g. the case of Chi-X above) suggesting monopsonistic pressures. For example, the monopsonist can pressure exchanges to develop new products which benefit it at the expense of other customers. This may be one force at work behind the appearance of co-location services which in effect give an advantage to technologically

intensive computer trading which is something large exchange customers certainly specialize in. This also suggests that competition among exchanges cannot be effectively analyzed without also considering broader features of the market ecology.

- Positive externalities and monopoly power: Traders co-ordinate where they trade to benefit from the liquidity they generate for each other. Such externalities of liquidity together with switching costs give primary exchanges market/monopoly power. This is exacerbated by a variety of regulatory privileges that stem from being the primary exchange including for example the right to accept new firms for listing or to set market trading hours.
- Barriers to entry: There are large barriers to entry especially for HFT. Arnuk and Saluzzi (2009) cite a TABB report according to which expenditures on colocation and facilities for fast access amount to \$1.8 billion per year (it was not clear to us whether this number referred only to expenses in US equity markets or whether it is broader). In addition, markets are also spending huge sums and the NYSE alone is investing in facilities at a cost of \$500 million. According to Price (2009), the cost of in-house solutions for competitive data feeds is of the order of \$260,000 per month and a start-up cost of \$270,000 per data center. Just to begin recording the kind of data that is necessary as a first step in designing a HFT strategy, a potential entrant would need several million per year. Furthermore, brokers typically require very large amounts of capital under management in order to provide competitive commission schedules (of the order of several tens of millions). Finally, the industry is very opaque in terms of information regarding fees and latencies of competing brokers and connectivity providers and it is a very difficult task to compare such information across a large menu of available options.
- Incumbent advantage: The problems of large barriers to entry and an oligopolistic structure are exacerbated because the benefits of privileged access are very difficult to quantify without owning privileged access. One reason for this is that an institution without co-location does not have access to historical data time-stamped at a co-location center with which to back-test HFT strategies. In our experience it is very difficult to independently commercialize a good computer trading idea without passing ownership of the underlying intellectual property to an HFT firm.
- Waste and negative sums arms races: Some modes of competition such as competition in terms of latency among HFT are negative sums arms races which can be expected to have a winner-takes-it-all outcome, suggesting that certain niches in the trading ecology may be dominated by a single uncontestable player. We would like to emphasize that in our opinion the resources being wasted on trading activities in general and computer trading in particular are huge, especially in terms of human capital. One quantitative indication of this is the study of Phillippon and Reshef (2009) which shows a huge shift in US skilled labour - especially PhDs - towards financial services from other industries since the deregulation of the 1980s.
- Path-dependency, lock in and multiple equilibria: are likely key features of the current computer trading ecology and indeed, Skouras and Farmer (2011) argue that the role speed plays in the current ecology may be the outcome of such features. This can lead to market ecologies that can be made more efficient with appropriate regulatory intervention.
- Price discrimination: As an example, the Deutsche Börse was recently allowed to offer AT cheaper prices than other types of traders. Volume based discounts and loyalty enhancing rebates for certain types of orders are extremely common. While the issue of whether rebates should be treated as abusive per se is controversial, there is broad consensus about their potential anticompetitive effects (e.g. Motta 2009)
- Legal loopholes and legal innovation: What constitutes a financial instrument is becoming increasingly ill-defined. For example, betting websites offer the possibility of bets contingent on the behaviour of key market indexes, while with the necessary incentives

legal innovation will devise contracts that serve to substitute standard financial investments. For example, most investment banks offer significant UK traders "contracts-for-difference" which are contracts that precisely simulate equity ownership while circumventing UK taxes on transactions ("Stamp Duty").

- **Regulator handicaps:** Regulators and academic researchers have several huge handicaps. They have incomparably fewer resources, personnel and incentives to collect and analyze market data than commercial traders, especially in HFT where strategies are quantitatively driven. At the same time, as discussed earlier, the theoretical tools for analyzing many issues of concern are lagging reality and markets are changing faster than they can be analyzed. In such an opaque setting, the immense political pressures that are routinely exerted by the "elite" lobbies of a few large financial institutions can skew policy (Johnson, 2009) and encourage regulatory capture. The problem is exacerbated by incentives for financial institutions to pursue "regulatory arbitrage", i.e. move their activities to the jurisdiction with the most favourable regulations. We do not mean to suggest that any of the above observations prove the existence of anticompetitive or even unfair behaviour, which would require a very detailed empirical analysis. Our intention is simply to show that the conditions for anticompetitive and unfair behaviour are present, and that this deserves further study. Any competition policy that attempts to deal with these issues of the market microstructure ecology should be founded on data and empirical evidence which at the moment is not even being collected. Calls by politicians for transaction taxes and bans on short selling despite at best very scant evidence that such measures are helpful (see e.g. Matheson 2011 and Saffi and Sigurdson 2008 respectively) are cases in point.

The extreme complexity of the interaction of effects we have described and the multi-platform nature of the economics of market microstructure imply that traditional economics approaches to competition policy might be intractable or impractical. Until theory catches up, a simulation-based evolutionary ecology perspective provides an attractive modeling approach on which to build an evidence-based competition policy. Current approaches to regulatory research are not well-suited for getting the needed answers, and in the absence of any hard evidence provides a good excuse for making decisions based on political expediency. Viewed from an ecology perspective such proposals are potentially dangerous as they would indiscriminately kill large and varied niches of trading species with potentially huge consequences for the entire ecology (see Section 4.2.3 below on the role of diversity in ecologies). While there are many indications pointing at certain types of computer trading as potentially problematic (as may be the case with HFT), there may be much more subtle ways to contain them which build on a detailed understanding of the microstructure ecology (see e.g. Skouras and Farmer 2011). In the medium run, we hope that policymakers will collect the data and develop the infrastructure and expertise that will make an evidence-based approach feasible.

4.2 Possible regulation to discourage market instabilities

4.2.1 Leverage and systemic risk in the ecology

As already mentioned, computer trading strategies often use high leverage. The recent crisis is widely regarded to be closely associated with too much leverage, followed by a liquidity freeze, with not enough leverage. Such a leverage cycle has been discussed extensively by Geanakoplos (2003, 2009) from an equilibrium point of view. We now present an evolution-based argument that gives a different intuition for what drives leverage cycles, discuss why this is particularly a problem for computer trading, and why this deserves careful study.

Many financial decision models lead to the conclusion that optimal leverage in an investment should be proportional to expected returns and inversely proportional to a measure of risk such

as the standard deviation or variance of returns. This can be made less subjective using the Kelly criterion, which has the advantage of yielding a quantitative result without an unknown parameter for risk aversion. The Kelly criterion is based on the desire to maximize long-range wealth without consumption along the way. Under idealizations including log-normal market returns and the ability to continuously and costlessly adjust leverage, the optimal leverage is

$$\lambda = \frac{\mu}{\sigma^2}$$

where μ is the expected unleveraged return of a given strategy and σ is its unleveraged standard deviation (MacLean et al., 2011, Peters, 2009). Note that this is a dimensionless number. It behaves in a sensible way: The optimal leverage increases with increasing unleveraged returns and decreases with risk.

As a strategy evolves along the path toward efficiency and becomes more widely used, its unleveraged expected return will decrease while its standard deviation increases (this is, after all, what it means to become more efficient). The formula above implies that the optimal leverage also decreases. This is the opposite of what usually happens, for a number of reasons. As unleveraged returns go down fund managers acting as agents under limited liability arrangements have institutionalized incentives to increase their leverage in an attempt to keep leveraged returns constant. In the early days, when unleveraged returns are high, fund managers do not need to use excessive leverage, but as time goes on, they must use more and more leverage to attract investors. This makes the market more vulnerable to crashes. This has been clearly demonstrated in a model of leveraged value investors (Turner et al., 2009)

In simulations it is also possible to show that underleveraging results in a sub-optimal rate of growth of investor capital, but a smooth growth trajectory for the capital and market prices. Over-leveraging, in contrast, also results in a sub-optimal rate of growth, but gives rise to a very bumpy ride for market prices, with dramatic losses for investors. Over-leveraging results in an unstable market, and can lead to "mass extinction" of many types of agents (see also Khandani and Lo, 2009).

4.2.3 Correlated strategies / crowded trades and systemic risk

A few algos (of a handful of institutions) and their interactions may dominate market dynamics. Fads in algos are more dangerous than fads in broader trading principles, because by involving extremely systematic trading patterns they can become very widespread and can lead to systemic risk because they are highly correlated across many participants. One reason for this is that it has proved very difficult in practice to protect intellectual property in the AT space so there is a lot of homogeneity and interdependence in algorithmic trading.

Indeed, this may well have been underlying the 1987 crash (Shiller 1988), the 2007 quant crisis (Khandani and Lo, 2009) and the 2010 flash crash - see e.g. the Flash-crash narrative as modified by Nanex (2010) which suggests it was due to an aggressive sell-off by highly correlated automated market making algos which were overwhelmed with the data task associated with the resulting surge in activity. Similarly, Gurliacci et al (2009) find that during a turbulent period in November 2008 the ecology of algos being used in US equity markets shifted dramatically. Khandani and Lo (2009) find that the "quant crash" of 2007 was due to crowded trades, which is equivalent to saying that these trading strategies were suffering from overpopulation.

4.2.4 Encouraging diversity for the sake of stability

In ecology there is a long-running debate about the relationship between diversity and stability. Are diverse ecologies more or less stable than non-diverse ecologies? In biology this debate

has now been raging for at least a half-century, and is considered one of the central questions in the field. In economics, surprisingly, this question has received little attention, although it is likely just as important.

In biology there is a basic intuition that complex, diverse ecosystems should be more resistant to perturbation because they have more redundancy and can therefore have more paths to recover from trauma. This intuition was seriously challenged by the work of May (1972), who studied randomly generated networks of interacting species (under a variety of strong assumptions) and showed that in this framework diverse ecosystems were actually less stable. In ecology this debate continues – it is not clear whether May's conclusions are the result of unrealistic assumptions or whether the core intuition that originally motivated the hypothesis that diversity favours stability is simply wrong.

In agent-based simulations of economic systems, in contrast, the evidence seems to favour the hypothesis that diversity is stabilizing. For example, LeBaron (2002) has simulated systems in which investors can choose to be either fundamentalists or trend followers, and shown that booms are driven by trend followers. Crashes follow booms, and occur when trend followers become too dominant in the population. One can also construct other examples, such as the order book imbalance/market maker ecology discussed above, in which an overpopulation of one type of strategy drives instabilities. Currently, however, despite the appeal of such arguments, there is little empirical evidence for or against this hypothesis, due to the difficulty of obtaining data with information about identity. One exception is Gurliacci et al (2009) and the SEC flash crash report which both suggest that the mechanism driving the flash crash involved the sudden coincident disappearance of a group of algos performing automated market making.

A further serious worry is that algorithmic strategies may displace humans performing similar functions (indeed, see IBM 2006 which predicts a 90% reduction of employment in human trading). But humans are still much better than computers at dealing with unexpected events and Barclay et al show that in less active markets where liquidity is lower and order matching is harder, humans tend to be responsible for a much larger portion of trading. But over a longer term, if humans have been permanently displaced by algos they will not be available to fill a gap created by algos that choose not to participate (Gurliacci et al 2009).

It is also worth noting that as the market ecology becomes less diverse, the possibility of "financial warfare" becomes more realistic. By this we mean that a malicious market participant such as a rich enemy state may be able to design and cause a major disruption in international financial markets by following an algorithmic "virus" strategy designed to interact catastrophically with the rest of the ecology. Because such a strategy is likely to be a strategy that intentionally loses a lot of money, the ecology may not be adapted to it which is why it is not too far-fetched that the potential for catastrophe may be present.

4.2.5 Inability of the AT ecology to cope with unexpected events

These patterns suggest a trend in systemic risk due to AT which is likely to be even worse because computer trading is likely to be subject to intense normalization of deviance effects whereby institutions tend to become too comfortable with respect to "deviant" high risk low-likelihood events (Cliff et al, 2011). In some ways this parallels the fact that technology tends to create situations where with small probability a catastrophe will occur. In the words of Douglas Adams, "The major difference between a thing that might go wrong and a thing that cannot possibly go wrong is that when a thing that cannot possibly go wrong goes wrong it usually turns out to be impossible to get at or repair." Such concepts are related to the principle of Highly Optimized Tolerance (Carlson and Doyle, 1999), which states that complex systems that

are optimized either through natural selection or by evolutionary design (such as the internet) to provide robust performance in complex environments may nonetheless be fragile to unusual perturbations.

Unfortunately, while in principle there may be profits to be made from strategies that stabilize extreme events it is unlikely that there is any business incentive to build such strategies which would involve an ongoing maintenance cost without any ability to verify performance or reward their designers until an extreme event occurs.

4.2.6 Controlling financial innovation: new financial products as invasive species

The question of diversity and stability becomes of central relevance when one tries to anticipate the effect of introducing exotic species. We know too well that introducing rabbits to Australia, or kudzu in North America, can have drastic and deleterious consequences. Ecologists have devoted considerable effort to trying to understand how to control such pests, and how to manage ecosystems so that they are as stable as possible against invasions by exotic species.

The analogy in financial markets is to financial innovation and new financial products. When a new financial product is introduced, will it be stabilizing or destabilizing? Because of the rapid pace of evolution of computer trading products, and because of the rigidity inherent in their programming, this problem is particularly acute. Without a deeper ecological understanding of markets and how they evolve through time, we will continue to have little ability to anticipate the effects of the introduction of new financial products. To obtain such an understanding we need better data gathering, better simulation models, and more theoretical investigations that develop these ideas.

4.2.7 Multiple equilibria in the ecology

It is well understood that even highly stylized simple models of ecologies and multi-platform markets can have multiple equilibria with more - and more diverse -- equilibria being present as the models become richer. In the finance context, multiplicity of equilibria implies the possibility of sudden switches between very different market ecologies. For example, Foucault et al (2009) find that monitoring decisions and market fee structures can interact to create multiple equilibria with different liquidity levels, trading rates and clustering of trades between which markets may switch.

The effect of switching of market organization can be harmful but it is not clear that this would lead to market crashes. A more worrying prospect is that computer trading may lead to price dynamics that have multiple equilibria. While many market participants suggest the flash crash of 2010 is an indication such dynamics are present we are sceptical about the possibility of equilibrium switching in the equity market due to errant algos, first because the dynamics of prices driven by errant algos are easy to spot and second because we are not convinced the equity market has multiple price equilibria. We are far more worried by what could happen if computer trading dominates markets that are much more likely to have multiple equilibria such as the sovereign debt market where noise created by errant algos might be much more dangerous. For example, what would have happened if computer algos had caused a large but plausible spike in Spanish interest rates? We suspect the global implications would have been far more wide-reaching than the clearly erroneous behaviour of equity markets during the Flash Crash.

5. On measuring the current ecology of computer trading

5.1 The need for measurements

In order to map out a real market ecology it is necessary to have data with detailed information about the identity of market participants which would allow most (if not all) market events to be classified in terms of the source of the decision that caused them. Most institutions are afraid that if the identity of their trades were known (even via anonymous codes), someone could potentially back-engineer their trading strategies. In addition there are very substantial technical difficulties to systematically monitoring information about identity. Thus there have been few studies with identity information and a focus on the market ecology.

We expect analysis of such data could be used to design better markets, better regulations and better market monitoring and warning of incipient crashes which could be used e.g. as a trigger of brief halts in trading (see Section 5.3). Furthermore collection of such data is a necessary first step towards the ecology simulation proposed in Section 3.4 though we emphasize it is a step which will be of great value even if an ecology simulation is never pursued. In Section 6 we will see how extrapolating long term trends in computer trading would be much easier if we had good measurements of the current state of the ecology.

Note that data with partial information about identity has already been shown to be extremely useful in an ecology context. Data from the London and Spanish stock exchanges contains brokerage information. This has been used to classify trading behaviors (Lillo et al) and to study the statistical properties and market impact of large institutional trades (Vaglica et al, Moro et al., Toth et al.)

5.2 On the type of data that is necessary

There are at least three major issues that proper systematic data collection needs to deal with:

1. In most countries with major financial markets, data is not collected on the level of the behaviour of particular strategies or even institutions. Currently, with enough effort regulators can track the legal entities behind ad hoc suspicious trades but this is a process that can take months and may not even be possible for orders that do not lead to trades (AFT 2011) to say nothing of the fact that beneficial owners of the legal entities may be unknown, that strategies may be disaggregated across separate legal entities or that many legal entities will typically aggregate many strategies into an overwhelming flow of orders in which undesirable or manipulative algos may be impossible to detect in practice.
2. Such data should cover all interrelated market venues in a uniform format. It is revealing that most of the analysis in this paper and in research we refer to is based almost exclusively on equity markets, while much trading is only sensible if viewed as part of a portfolio inventory that may extend across asset classes and regulatory jurisdictions.
3. Datasets collected by the exchanges themselves (and then made available to researchers or regulators) are not time stamped or measured with the level of accuracy that is considered essential by serious AT firms.

While limited types of data with some form of identity information exists, it is rarely available openly or commercially but rather is made available to particular researchers for particular purposes meaning that even the few empirical results that have been reported cannot be independently replicated as is usual in empirical finance. Furthermore, the data usually is not detailed enough to support the goals we propose here, though there are initiatives to begin collecting more information close to what would be necessary.

In particular, some countries require that all market events are assigned to a particular trading account which is distinguished in terms of the human trader responsible for it. If the Office of Financial Research fulfils its mandate as envisioned by the committee that conceived it, they will have the authority to collect financial data in real time including account level information. With this kind of data it would be possible to classify trading activity according to its ecological type, study how such ecologies evolve over time, and monitor them in real time. In Europe, MiFID2 could initiate such data collection requirements (and if the Americans do this as well, they could probably do so without fear of business going elsewhere). To begin historical studies, they might be able to convince other countries that have already been collecting such data to selectively release their data for detailed pilot studies.

In our opinion the usefulness of the kind of data we describe here is so great, that its lack should be viewed in itself as a significant lacuna of the current international financial system.

5.4 Early warning diagnostics for market crashes based on patterns in the market ecology

Our ecology perspective suggests that certain patterns in the ecology could provide warning signals for market crashes if observed in real-time. If real-time data close to the strategy level were available, an authority mandated with financial oversight could for example implement market halts if alarming measurements of the following were observed:

- Diversity the population of strategies.
- Interactions of the AT ecology with human market participants.
- The leverage of each group of strategies and their ability to withstand a performance shock or large fund withdrawal.

6. A long term view: How might the ecology evolve over the next decade?

In this section we turn our attention to the extremely challenging and risky task of conjecturing what markets might look like in the medium-run future. We emphasize our microstructure ecology perspective and attempt to build on our previous observations and to extrapolate the long term trends we believe will remain stable.

Here are some long term trends that we consider particularly important:

- Regulation will continue to shape markets: Paradoxically, the role of policy is so significant in how markets change that it becomes difficult to analyze to the point that it is neglected in scientific research. Nevertheless, regulators have been and will continue to be key agents driving change in the market ecology. Successful market participants are often biased in the way they think about regulation, confounding proposed changes in regulation with increases in regulation, so that they are readily dismissed on the basis of any sweeping standard anti-regulation argument. For example, MiFID clearly encouraged competition among European exchanges and created profitable opportunities for all sorts of AT and HFT strategies which were symbiotic with the new exchanges and which exploited and encouraged fragmentation of liquidity. For obvious reasons, it is natural that those who benefitted from MiFID the most are also the most wary of MiFID 2. We believe that the shape of markets a decade from now will depend greatly on the sophistication and enlightenment of policy on the global level.
- High frequency trading is infrastructure intensive. Researchers and regulators need substantial infrastructure in the form of data gathering capabilities, hardware and software to understand what is going on in financial markets. For market participants this hardware

is absolutely necessary to make their living. For researchers and regulators, in contrast, the substantial resources needed to fund such infrastructure can be difficult to obtain. As an example, consider the fact that the sampling frequency at which empirical analyses are conducted can lead to very different conclusions. The large literature empirically analyzing inter-market price discovery has been conducted on a minute-by-minute time interval making it largely irrelevant for today's markets where we need to go to sub-second measurements (see e.g. Muravyev, Pearson and Broussard 2011 and Axioglou 2010). See our Figure 3 as an example of how misleading analyses can be if they are based on too coarse-grained timestamps. Unfortunately commercial vendors do not pursue accurate time-stamping so only a handful of researchers with privileged relationships with market participants can now properly study themes that used to be mainstream topics in empirical finance.

This trend of a widening knowledge gap is likely to be exacerbated if left untreated because of the huge incentive gap between researchers and market participant to conduct tedious and difficult market data analysis.

- Decreasing trading costs?: Another important trend which we believe is often somewhat misinterpreted has been the universal reduction in spreads, commissions and exchange fees in the last few decades. According to Table 2, spreads on NYSE are now around 3.5 basis points for highly liquid stocks while Jones (2002) reports spreads of 70 basis points in 1900. While many commentators view this trend towards decreasing spreads, commissions and exchange fees as caused by the coincident trend in computer trading, not only is the direction or existence of any causality difficult to establish, but the trend itself may well be less significant than it seems.

First, note that spreads only began falling around the 1970s so clearly major improvements in technology that occurred from 1900 to 1970 do not always result in a decrease in spreads - so again policy and market organization evidently have played a key role.

Second, while costs for making a small roundtrip trade have certainly decreased, it is not clear that the cost of buying a significant portion of a company (e.g. 5%) has decreased nor have measures of the total trading costs generated by trading activity decreased. Indeed, market trading costs viewed as the product of turnover with the cost per trade have not really dropped in the period 1950-2000 (Jones 2002). Our own calculations based on more recent data from the London Stock Exchange suggests that the last decade has led to an increase in aggregate costs of total trading as well. This illustrates that some apparently beneficial trends may in fact be much more neutral than they seem because they are correctly interpretable only together with other broader trends.

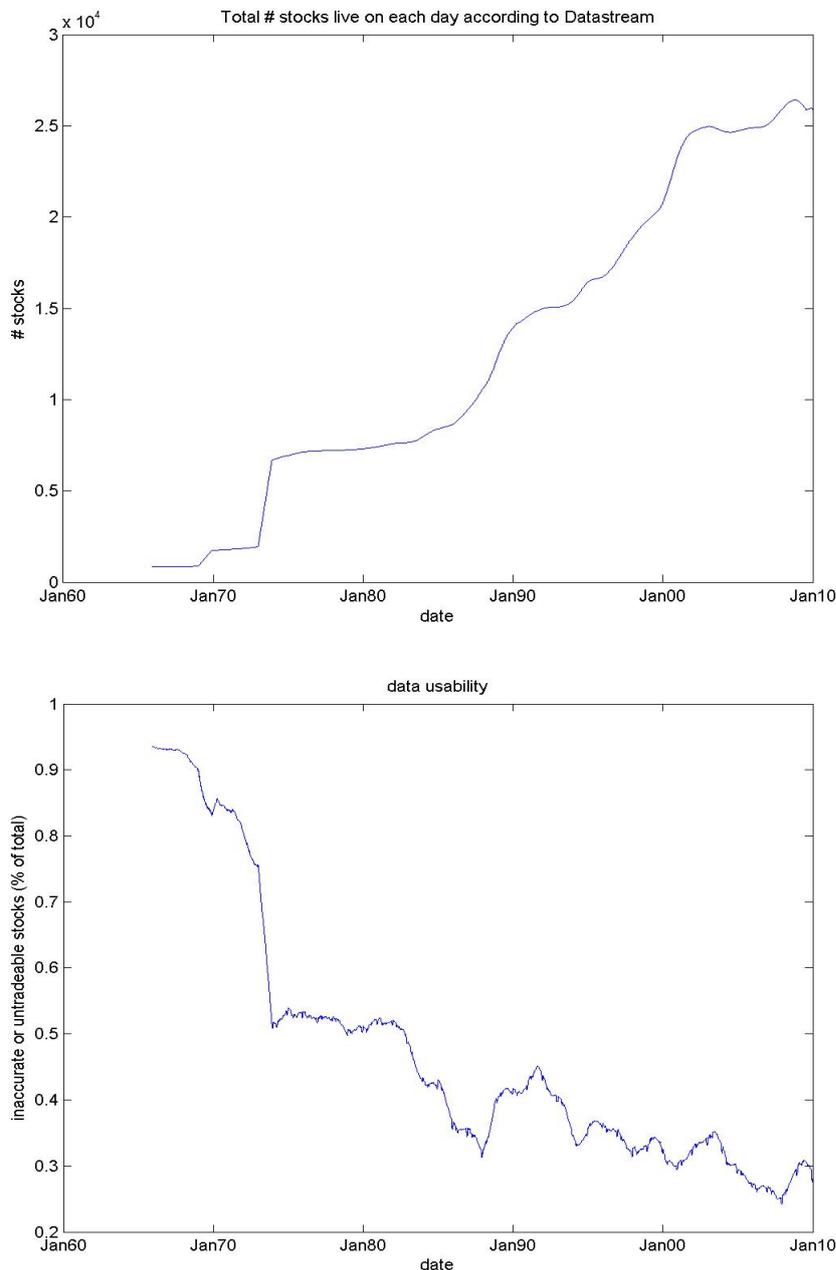
- More crowdedness and systemic risk: Another important trend broadly recognized by market participants (for which we do not have any original evidence) is the increasing crowdedness of computer trading. Evidence in this direction is implicit in several disparate research efforts. For example, Khandani and Lo (2007) demonstrate a very significant increase in the crowdedness of what they call "quant trading" which is in fact largely computer trading. It is also clear e.g. from the SEC flash-crash report that the HFT market making space is highly crowded with potentially systemic consequences for risk. This suggests crowdedness monitoring can be a good way of tracking one form of systemic risk generated by computer trading.
- Mobility of financial centres, employment and policy co-ordination: Increased competition together with the effects of reduced inter-market latencies afforded by the agglomeration of markets has led to consolidated market centres with the closure of many regional exchanges internationally and even with the migration of primary exchanges across borders (e.g. the primary market of Portuguese stocks is Euronext which used to be

An ecological perspective on the future of computer trading

located in Paris until it recently migrated to London). Importantly, the role of humans in markets may dwindle (IBM (2006) forecasts a 90% reduction of employment in trading related positions because of replacement by computers) making markets even more mobile. This suggests financial market policies should be co-ordinated at the G-20 level to discourage regulatory arbitrage which will become increasingly easy for markets and increasingly tempting to governments.

- More markets, more trading, more data: One often overlooked effect of computer trading and generally technology in markets has been the fact that the integrity of market data has improved vastly over time (see figure 6 below).

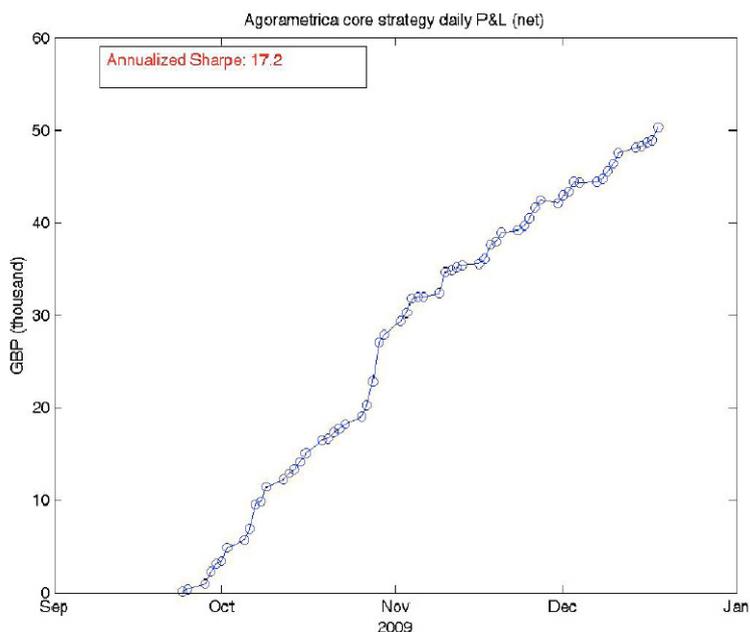
Figure 6: The universe of stocks which can be empirically analyzed with standard commercial datasets has vastly expanded over time, not just because the number of stocks has increased (top figure) but also because computer trading has facilitated the collection of cleaner data (bottom figure). The top figure illustrates the number of stocks that were both tradable each day ("live") and that are now in the Thomson-Datastream global database (which provides an upper bound on the number of stocks that would have been available in databases in real-time). The bottom figure illustrates what proportion of these stocks passes a range of weak integrity filters and had at least one trade, plotted for each date covered by the database (see Landis and Skouras 2011).



No doubt the availability of high quality market data, driven by the electronification of exchanges and improved database technologies has greatly facilitated the expansion of AT. Similarly, computers have driven the creation and trading of increasingly complex financial products. It will be increasingly difficult to operate in markets without sophisticated analytics and technology.

- More speed: A controversial trend has been the vast reduction in market latencies. For example, the execution speed of the NYSE in the 1980s was 5-10 seconds (see e.g. Hendershott and Moulton, 2010) whereas now it is in the region of microseconds. An effective way to measure this market speed-up is by analyzing the timeframes over which related markets interact. Figure 3 below leaves no doubt that interactions between markets and intermarket price discovery has become much faster than it has been in the past.

Figure 3: This plot illustrates the real audited track record of an automated market making strategy developed by a firm co-founded by the authors which was active in European equity markets in 2009.



Based on our previous discussion of HFT as an element in the market ecology, we believe that an effort to characterize speed as good or bad is a false dichotomy: speed is good in some ways and bad in others, so that its net effect may be quite sensitive to the specifics of a particular market at any point in time.

- Trends in the ecology: Unfortunately, detailed long-term microstructure data from which to infer long term trends in market ecologies are not readily available to us and indeed there have been few studies using such data. We can deduce some broad features however by comparing international markets which differ in their maturity (see Table 1)

Table 1: This table illustrates differences across market structure globally. Presenting markets in decreasing order of maturity we see that maturity leads to smaller trade sizes, smaller trade-to-quote ratios and less competition among exchanges. The numbers are 20 day moving averages ending on April 27 2011 for all stocks in all markets covered with valuations in the region US\$64-128 billion. All reported currencies are in US\$. Market share means the proportion of total volume traded on each

An ecological perspective on the future of computer trading

exchange and Frequency of BBO represents the fraction of times on which quotes with the best bid or offer were on each exchange.

US											
	Consolidated	BATS	ARCA	CBOE	TRF	BX	CHX	NSX	NSDQ/INT	MSDQ/NMS	NYSE
Trade Size		8,813	10,071	13,352	13,471	7,816	19,801	10,856	9,607	9,373	15,829
Quote Size	691,643										
Spread (b.p.)	3.5										
Market share		9.66%	13.85%	0.03%	32.55%	2.43%	0.05%	0.95%	12.17%	11.97%	16.36%
Frequency of BBO		14.32%	20.53%	0.04%	3.6%	0.04%	1.42%	18.04%	18.04%	17.75%	24.26%
Europe											
	Primary Exchange	CHIX	TRC	BATS	Arca	Off market	Commen				
Trade Size	21,512	10,618	9,084	9,084							
Quote Size	153,215	77,506	37,316	53,085							
Spread (b.p.)	4.68	5.73	6.63	6.19							
Market share	34.75%	8.44%	1.75%	2.52%	0.02%	52.54%					
Frequency of BBO	19.66%	5.43%	0.57%	0.97%			73.37%				
Asia											
	Consolidated	Tokyo	Sydney	Hong Kong	Korea						
Trade Size		25,675	8,351	22,500	314						
Quote Size	1,056,675										
Spread (b.p.)	12.05										

Noting that Asian markets are less developed than European markets which are less developed than US markets we observe that the trend is towards smaller trade sizes and larger ratios of trade-to-order sizes, more competition among exchange venues and as we have already noted smaller spreads and tick sizes. This suggests more mature markets have more order-splitting, including splitting across several exchanges. There also exists significant variation in the ecology of even interacting closely concentrated markets in Europe as illustrated in Table 2. This suggests that there is significant hysteresis or path dependency across market ecologies which may allow heterogeneous ecologies to persist across similar and connected markets indefinitely into the future.

Table 2: This table illustrates the differences in number of trades, order book events and the ratio of the two across primary markets and across primary markets and Chi-X for the 50 stocks in the STOXX50 index for May 2009. Numbers in parentheses are standard errors. Evidently, there are big differences across markets and the relationship between behaviour of the primary market and Chi-X differs a lot depending on the market.

	London	Netra	Milan	Stockholm	Helsinki	Luxembourg	Madrid	Swiss	All
# trades per stock day (10000)	4.98 (0.17)	4.46 (0.15)	7.57 (0.60)	11.63 (0.85)	2.66 (0.13)	4.13 (0.14)	3.44 (0.20)	3.07 (0.10)	4.51 (0.09)
trades/orders (% of) (std)	64.89 (0.30)	61.96 (0.54)	87.32 (0.54)	59.99 (1.54)	70.70 (1.15)	56.48 (0.62)	99.83 (0.86)	74.83 (0.87)	66.93 (0.45)
# order book events per stock day (10000)	129.73 (3.13)	142.21 (4.61)	106.43 (6.21)	34.75 (2.28)	128.76 (6.25)	138.86 (5.85)	67.64 (2.74)	57.83 (1.57)	118.26 (2.25)
trades/orders (% of) (std)	67.76 (0.25)	66.84 (0.41)	69.64 (0.73)	65.87 (0.73)	88.81 (0.62)	64.57 (0.57)	83.56 (0.56)	63.15 (0.56)	67.23 (0.26)
# trades per order book event (std)	3.92 (0.08)	3.70 (0.09)	7.02 (0.27)	1.82 (0.07)	2.13 (0.13)	3.34 (0.09)	5.27 (0.13)	5.41 (0.13)	4.14 (0.06)
trades (% of)	0.59 (0.06)	0.70 (0.11)	4.85 (0.26)	0.57 (0.21)	1.30 (0.18)	1.23 (0.11)	4.98 (0.12)	2.06 (0.21)	12.26 (0.09)

- Increasing specialization of the trading taxonomy: Trading patterns and behaviours are increasingly sensitive to the fine print of market design as illustrated for example by Figure 2. We have also already noted trends in trading styles such as market making and pairs trading which used to be conducted by humans but have now become automated

and increasingly fast and sophisticated. This corresponds to standard patterns seen in ecologies over time.

7. Conclusion: How to prepare for the future

Our analysis suggests that computer trading will put several forces at work on financial markets, potentially in opposite directions. On the positive side, the adoption of more sophisticated execution technologies by markets and algorithmic trading by investors seems to have had a very positive effect on markets viewed over long time scales. However, some types of computer trading such as HFT may push markets towards increasingly anticompetitive or unfair directions and may be dangerous for market stability with implications that may be felt on the level of the entire international financial system and the global economy.

Our ecological perspective emphasizes the co-evolution of computer trading, human trading, markets and regulators. Decisions by regulators and the academic research that drives them can have a lasting important impact on financial markets and economic well-being. For example, regulations that encouraged ECNs and MTFs to compete with primary markets have dramatically changed the trading landscape as has MiFID. The classic but controversial finding of collusion on NASDAQ by Christie and Schultz (1992) almost certainly played a role in changing US market regulations that surely sped up the global dominance of the open electronic order book. Phillippon and Reshef (2008) provide strong evidence that the trend towards deregulation of the financial services industry that began in the 1980s is responsible for the trend for using increasingly skilled human capital in financial services which no doubt also contributed to rapid adoption of new technologies in this sector. Interestingly, this suggests one way to contain the AT trends in computer trading might be to incentivize skilled human capital to pursue other activities such as basic science research.

The future of computer trading and financial markets is likely to be shaped very significantly by current policy decisions, just as the current state has been shaped to a much larger extent than many market participants acknowledge by previous regulations. Just as the market ecology evolves, it is also natural that regulation evolves. Regulation must be constantly modified as the market ecology changes and as our understanding of how markets function improves. We believe that it is essential that researchers and regulators develop a full understanding of the functioning of markets, firmly grounded in empirical observation. This should include a detailed understanding of the market ecology:

1. Who are the players, what are they doing, and how is this changing in time. Our proposals are the following, in increasing order of difficulty of implementation:
2. Measure the market ecology and make the corresponding data broadly available (see Section 5). This is an essential ingredient for all subsequent goals and should help to avoid any policy that is not evidence-based. Particular attention needs to be taken to avoid policies that will be disruptive of the entire market ecology in ways that cannot be predicted without detailed data collection and analysis.
3. As discussed in Section 4, any new policies need to be designed to cope with the complex multi-platform nature of modern computer trading environments. In parallel a new regulatory environment should emerge that can minimize the potential for instabilities generated by computer trading. In our separate driver review on speed in financial markets (Skouras and Farmer 2011) we discuss these issues specifically in relation to speed and HFT, and discuss the desirability of coordination at the G-20 level.
4. Develop a sophisticated simulation approach to modeling the market ecology (see Section 3 for a detailed discussion). The market ecology approach is in our opinion a key

An ecological perspective on the future of computer trading

ingredient of any large scale simulation effort (e.g. as envisioned by Cliff et al 2011) which we believe may be the most fruitful approach to analyzing many important issues in market microstructure. Furthermore, it will generate positive feedbacks with our ongoing effort to develop a more fully articulated theoretical framework for thinking about market microstructure from an ecology perspective.

More work certainly needs to be done to develop a deep understanding of computer trading from the perspective of market microstructure ecology. However, computer trading and related electronic commerce technology already extends well beyond financial markets and will likely form an increasing portion of economic activity in years to come. Investing in understanding computer trading may lead to great returns as such work would be a very useful building block for understanding the new shapes economic activity will take in the future.

References

- AFM, 2011, "High Frequency Trading: The application of advanced trading technology in the European marketplace", Netherlands Authority for the Financial Markets (AFM).
- Narang, R.K., 2009, Inside the black box: The simple truth about quantitative trading, Wiley Finance.
- Axioglou, C, 2010, "Sub-second price discovery across European limit order books", Athens University of Economics and Business working paper
- Barclay, M.J., Hendershott, T. and K. Kotz, 2006, Automation versus intermediation: evidence from trasuries going off the run, Journal of Finance LXI, 5, October 2006
- Brogaard, J.A., 2010, High frequency trading and its impact on market quality, working paper
- BMO Capital Markets, 2009, The Impact of High Frequency Trading on the Canadian Market, July 22, 2009
- Carlson, J.M. and J. Doyle, 1999, "Highly optimized tolerance: A mechanism for power laws in designed systems", Physical Review E., Vol. 60, 2, 1412 - 1427.
- Cantillon, E. and P-L Yin, 2010, "Competition between Exchanges: A research agenda", International Journal of Industrial Organization, forthcoming
- Cartea A, and J. Penalva, 2010, "Where is the value in high frequency trading", working paper
- Chaboud, A., Hjalmarsson E., Vega, C. and B. Chiquoine, 2009, "Rise of the Machines: Algorithmic trading in the foreign exchange market", FRB International Finance Discussion Paper No. 980
- Cherkashin, Dmitriy, J. Doyne Farmer, and Seth Lloyd, 2009, "The Reality Game." Journal of Economic Dynamics and Control, Vol. 33, 5, 1091-1115.
- Christie, W.G., and P. H. Schultz, 1994, "Why do NASDAQ market makers avoid odd eighth quotes", XLIX, 5, December
- Cliff, D., and L. Northrop, 2011, "The global financial markets: an ultra-large-scale systems perspective", Forecast Driver Review DR4.
- Cont, R. and XXX, 20XX, "Running for the exit: Short selling and endogenous correlation in financial markets, <http://ssrn.com/abstract=1722508>.
- Cvitanic J. and A.A. Kirilenko, 2011, High Frequency Traders and Asset Prices, working paper
- Domowitz, I, 1993, A taxonomy of automated trade execution systems, Journal of International Money and Finance, 12, 607-631.
- Easley D., Lopez de Prado, M.M. and M. O'Hara, 2010, "The microstructure of the flash crash: flow toxicity, liquidity crashes and the Probability of Informed Trading", working paper
- Easley, D., Hendershott, T. and T. Ramadorai, 2009, Levelling the trading field, working paper
- Easley, D., M. O'Hara and L. Yang, 2011, "Differential access to price information in financial markets, working paper
- Economides, N., 2008, "Net Neutrality, Non-discrimination, and Digital distribution of content through the internet", I/S: A Journal of Law and Policy for the Information Society, 4(2), 209-233
- Evans, D.S., 2003, "The antitrust Economics of multi-sided platform markets", Yale Journal on Regulation, 325.

- Farmer, J.D., 2002. "Market force, ecology and evolution", *Industrial and Corporate Change*, Vol. 11, 5, 895 - 953.
- Farmer, J.D., and J. Geanakoplos. 2009, "The Virtues and Vices of Equilibrium and the Future of Financial Economics." *Complexity*, Vol. 14, 3, 11-38.
- Foucault T, O. Kadan and E. Kandel, 2009, "Liquidity Cycles and Make/Take Fees in Electronic Markets, working paper.
- Geanakoplos, John. 2003. *Advances in Economics and Econometrics: Theory and Applications*, Eighth World Congress, eds. Dewatripont, M., L.P. Hansen and S.J. Turnovsky, 170-205. Cambridge: Cambridge University Press, 2.
- Geanakoplos, John. 2009. "The Leverage Cycle.", Tech. rep. National Bureau of Economic Research Macro-Annual.
- Gurliacci, M., Jeria, D. and G. Sofianos, 2009, "When the going gets tough, the algos get going, *Journal of Trading*, Vol 4, 2, 34-45
- Hasbrouck, J., 1995, "One security, many markets: determining the contributions to price discovery", *Journal of Finance*, 50, 1175-1199.
- Hasbrouck, J. and G. Saar, 2011, "Technology and Liquidity provision: The blurring of traditional definitions", *Journal of Financial Markets*, 12(2), May, 143-172
- Hendershott, T. and R. Riordan, 2009, "Algorithmic trading and information", working paper
- Hendershott T, Jones, C.M. & A.J. Menkveld, 2011, "Does algorithmic trading improve liquidity?", working paper
- Hendershott, T. and P.C. Moulton., 2010, "Automation, speed and stock market quality: The NYSE's Hybrid", working paper
- Jacobs B.I., K.N. Levy and H.M. Markowitz, 2004, "Financial Market Simulation", *Journal of Portfolio Management*, 30th Anniversary Issue
- Jain P.K. & W.F. Johnson, 2008, "Trading Technology and Stock Market Liquidity: A Global Perspective", in, *Stock Market Liquidity*, F-S. Lhabitant and G.N. Gregoriou, John Wiley and Sons Inc.
- Jain P.K, 2005, "Financial market design and equity premium: Electronic versus floor trading", *Journal of Finance LX*, 6, December
- Jones, C., 2002, "A century of stock market liquidity and trading costs", working paper
- Jarrow R.A. and P. Protter, 2011, "A dysfunctional role of high frequency trading in electronic markets", Johnson School Research Paper Series No. 08-2011
- Kearns, M., Kulesza, A. and Y. Nevmyvaka, 2010, *Empirical Limitations on High Frequency Trading Profitability*, working paper
- Khandani, A.E., A.W. Lo, 2007, "What happened to the quants in August 2007", *Journal of Investment Management*, 5(4), 5-54.
- Kim G. and H. Markowitz, 1989, "Investment rules, margin and market volatility", *Journal of Portfolio Management*, 16(1), 45-52
- LeBaron, B., 2002, "Short-memory traders and their impact on group learning in financial markets", *Proceedings of the National Academy of Science*, vol 99 suppl. 3, 7201-7206

- Lim, K.P. and R.D. Brooks, 2011, "The evolution of stock market efficiency over time: a survey of the empirical literature", *Journal of Economic Surveys*, 25(1), 69-108
- MacLean, L.C., Thorp, E.O. and Ziemba, W.T., 2011, "The Kelly Capital Growth Investment Criterion", World Scientific.
- Masson, P.R., 1999, "Multiple equilibria, contagion and the emerging market crises", IMF working paper.
- Matheson, T., 2011, "Taxing Financial Transactions: Issues and Evidence", IMF working paper 11/54
- May, R.M., 1972, "Will a Large Complex System be Stable?", *Nature* 238, 413 - 414.
- Muravyev, D., Pearson, N.D. and J.P. Broussard 2011, 2011, "Is there price discovery in equity options?", working paper
- Naik, N.Y. and P.K. Yadav, 1999, "The effects of market reform on trading costs of public investors: Evidence from the London Stock Exchange", London Business School working paper
- Naik, N.Y. and P.K. Yadav, 2003, "Do dealer firms manage inventory on a stock-by-stock or a portfolio basis?", *Journal of Financial Economics*, 69(2), 325-353
- Nanex, 2010, "May 6th 2010 Flash Crash Analysis: Final Conclusion", October 14, 2010.
- Peters, O., 2009, "Optimal leverage from non-ergodicity", working paper.
- Price, T. 2009, *Market Data Management: In-house vs outsourced solutions*, Towergroup, February.
- Phillippon, T. and A. Reshef, 2008, "Wages and Human Capital in the U.S. Financial Industry: 1909-2006", working paper
- QSG, 2009, "QSG study proves higher trading costs incurred for VWAP algorithms vs Arrival Price algorithms, High Frequency Trading Contributing Factor", November 24, 2009
- Reiss, P.C. and I.M. Werner, 1998, "Does risk sharing motivate interdealer trading", *Journal of Finance*, 53(5), October
- Saffi P.A.C and K. Sigurdson, 2008, "Price efficiency and short selling", working paper
- Shiller, R., 1988, "Portfolio Insurance and Other Investor Fashions as Factors in the 1987 Stock Market Crash", *NBER Macroeconomics Annual*, Vol. 3, pp. 287-297
- Skouras, S. and J.D. Farmer, 2011, "Speed in Markets", *Foresight Driver Review*
- Toth, B., Eisler, Z., Lillo, F., Bouchaud, J-P., Kockelkoren, J. and J.D. Farmer, 2011, "How does the market react to your order flow?", working paper.
- Turner, S., Farmer, J.D. and Geankoplos, J., "Leverage causes fat tails and clustered volatility", working paper.
- Tradeworx, 2010, "Public commentary on SEC market structure concept release", April 21.
- Zhang, F, 2010, "High frequency trading, stock volatility and price discovery", Yale university working paper
- Zhang, F, and S.B. Powell, 2011, "The impact of high-frequency trading on markets", *CFA magazine*, March-April.

© Crown copyright 2010
Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk
URN 11/1225