Cartlidge, J., & Cliff, D. (2018). Modelling complex financial markets using real-time human-agent trading experiments. In S-H. Chen, Y-F. Kao, R. Venkatachalam, & Y-R. Du (Eds.), *Complex Systems Modeling and Simulation in Economics and Finance* (pp. 35-69). (Springer Proceedings in Complexity). Springer International Publishing AG. https://doi.org/10.1007/978-3-319-99624-0_3

Peer reviewed version

License (if available):
Other

Link to published version (if available):
[10.1007/978-3-319-99624-0_3](https://doi.org/10.1007/978-3-319-99624-0_3)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Springer Nature at https://link.springer.com/chapter/10.1007%2F978-3-319-99624-0_3 . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Modelling complex financial markets using real-time human-agent trading experiments

John Cartlidge* and Dave Cliff

Key words: Agent-based Computational Economics, Automated Trading, Continuous Double Auction, Experimental Economics, High-Frequency Trading, Human-Agent Experiments, Robot Phase Transition, Trading Agents

Abstract

To understand the impact of high frequency trading (HFT) systems on financial market dynamics, a series of controlled real-time experiments involving humans and automated trading agents were performed. These experiments fall at the interdisciplinary boundary between the more traditional fields of behavioural economics (human-only experiments) and agent based computational economics (agent-only simulations). Experimental results demonstrate that: (a) faster financial trading agents can reduce market efficiency—a worrying result given the race towards zero-latency (ever faster trading) observed in real markets; and (b) faster agents can lead to market fragmentation, such that markets transition from a regime where humans and agents freely interact, to a regime where agents are more likely to trade between themselves—a result that has also been observed in real financial markets. It is also shown that (c) realism in experimental design can significantly alter market dynamics—suggesting that, if we want to understand complexity in real financial markets, it is finally time to move away from the simple experimental economics models first introduced in the 1960s.

John Cartlidge

Department of Computer Science, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, UK, e-mail: john.cartlidge@bristol.ac.uk

Dave Cliff

Department of Computer Science, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, UK, e-mail: csdtc@bristol.ac.uk

* Contact author: john@john-cartlidge.co.uk

1 Introduction

In recent years, the financial markets have undergone a rapid and profound transformation from a highly regulated human-centred system, to a less-regulated and more fragmented computerised system containing a mixture of humans and automated trading systems (ATS)—computerised systems that automatically select and execute a trade with no human guidance or interference. For hundreds of years, financial trading was conducted by humans, for humans, via face-to-face (or latterly telephone) interactions. Today, the vast majority of trades are executed electronically and anonymously at computerised trading venues where human traders and ATS interact. Homogeneous human-only markets have become heterogeneous human-ATS markets, with recent estimates suggesting that ATS now initiate between 30% and 70% of all trades in the major US and European equity markets [25].

As computerisation has altered the structure of financial markets, so too the dynamics (and systemic risk) have changed. In particular, trading velocity (the number of trades that occur in unit time) has dramatically increased [25]; stocks and other instruments exhibit rapid price fluctuations (*fractures*) over subsecond time-intervals [36]; and wide-spread system crashes occur at astonishingly high speed. Most infamously, the *flash crash* of May 6th 2010 saw the Dow Jones Industrial Average (DJIA) plunge around 7% (\$1 trillion) in 5 minutes, before recovering most of the fall over the proceeding 20 minutes [21, 37]. Alarming, during the crash, some major company stocks (e.g., Accenture) fell to just one cent, while others (e.g., Hewlett-Packard) increased in value to over \$100,000. These dynamics were unprecedented, but are not unique. Although unwanted, flash crashes are now an accepted feature of modern financial markets.²

To accurately model financial systems, it is now no longer sufficient to consider human traders only; it is also necessary to model ATS. To this end, we take a bottom-up, agent-based experimental economics approach to modelling financial systems. Using purpose built financial trading platforms, we present a series of controlled real-time experiments between human traders and automated trading agents, designed to observe and understand the impact of ATS on market dynamics. Conducted at the University of Bristol, UK, these experiments fall at the interdisciplinary boundary between the more traditional fields of experimental economics (all human participants) and agent based computational economics (all agent simulation models) and offer a new insight into the effects that agent strategy, agent speed, human experience, and experiment design have on the dynamics of heterogeneous human-agent markets.

Results demonstrate that: (a) the speed of financial agents has an impact on market efficiency—in particular, it is shown that faster financial trading agents can lead to less efficient markets, a worrying result given the race towards zero-latency (ever faster trading) observed in real markets; (b) faster agents can lead to mar-

² Flash crashes are now so commonplace that during the writing of this chapter, a flash crash occurred in the FX rate of the British Pound (GBP). On Oct 7 2016, GBP experienced a 6% drop in two minutes, before recovering most of the losses [53]—a typical flash crash characteristic.

ket fragmentation, such that markets transition from a regime where humans and agents freely interact, to a regime where agents are more likely to trade between themselves—a result that has also been observed in real financial markets; (c) experiment design, such as discrete-time (where participants strictly act in turns) versus real-time systems (where participants can act simultaneously and at any time), can dramatically affect results, leading to the conclusion that, where possible, a more *realistic* experiment design should be chosen—a result that suggests it is finally time to move away from Vernon Smith’s traditional discrete time models of experimental economics, first introduced in the 1960s.

This chapter is organised as follows. Section 2 (Motivation) introduces the argument that financial systems are inherently complex ecosystems that are best modelled using agent-based approaches rather than neoclassical economic models. Understanding the causes and consequences of transient non-linear dynamics—e.g., *fractures* and *flash crashes* that exacerbate systemic risk for the entire global financial system—provides the primary motivation for this research. In Section 3 (Background), the agent-based experimental economics approach—i.e., human-agent financial trading experiments—is introduced and contextualised with a chronological literature review. Section 4 (Methodology) introduces the trading platform used for experiments, and details experiment design and configuration. Empirical results presented in Section 5 (Results) demonstrate market fragmentation—a significantly higher proportion of agent-only and human-only trading in markets containing super-humanly fast agents. Since the experimental market we construct is too constrained to exhibit fractures directly, in Section 6 (Discussion) we interpret this result as proxy evidence for the robot phase transition associated with fractures in real markets. In Section 7 (Conclusion), conclusions are drawn, and some avenues for future research are outlined.

2 Motivation

There exists a fundamental problem facing financial market regulators—current understanding of the dynamics of financial systems is woefully inadequate; there is simply no sound theoretical way of knowing what the systemic effect of a structural change will be [7]. Therefore, when policy makers introduce new market regulation, they are effectively trial and error testing in the live markets. This is a concerning state of affairs that has negative ramifications for us all; and it provides adequate motivation for the research presented here.

In this section, it is argued that our lack of understanding is a symptom of the dominant neoclassical economic paradigms of rational expectations and oversimplified equilibrium models. However, a solution is proposed. It has been compellingly argued elsewhere that economic systems are best considered through the paradigm of complexity [41]. Agent-based models—dynamic systems of heterogeneous interacting agents—present a way to model the financial economy as a complex system [22] and can naturally be extended to incorporate human (as living agent) inter-

actions. In addition, the converging traditions of behavioural [38] and experimental [48] economics can address non-rational human behaviours such as overconfidence and fear using controlled laboratory experiments. Here, we present a hybrid approach—mixed human-agent financial trading experiments—that we believe offers a path to enlightenment.

2.1 Complex Economic Systems

Neoclassical economics relies on assumptions such as market efficiency, simple equilibrium, agent rationality, and Adam Smith’s invisible hand. These concepts have become so ingrained that they tend to supersede empirical evidence, with many economists subliminally nurturing an implicit Platonic idealism about market behaviour that is divorced from reality. As Robert Nelson argued in his book, *Economics as Religion*, it is almost “as if the marketplace has been deified” [42]. Consequently, no neoclassical framework exists to understand and mitigate *wild* market dynamics such as flash crashes and fractures. It is necessary, therefore, to develop “a more pragmatic and realistic representation of what is going on in financial markets, and to focus on data, which should always supersede perfect equations and aesthetic axioms” [7]. Disturbingly, despite global capitalism’s existential reliance on well functioning financial markets, there exists no mature models to understand and predict issues of systemic risk [14]. Policy makers, therefore, are essentially acting in the dark; with each new regulatory iteration perturbing the market in unanticipated ways.

Fuelled by disillusionment with orthodox models, and a desire to address the inadequacies of naïve policy making, there is a trend toward alternative economic modelling paradigms: (1) *Non-equilibrium economics* focuses on non-equilibrium processes that transform the economy from within, and include the related and significantly overlapping fields of *evolutionary economics* (the study of processes that transform the economy through the actions of diverse agents from experience and interactions, using an evolutionary methodology, e.g., [43]), *complexity economics* (seeing the economy not as a system in equilibrium, but as one in motion, perpetually constructing itself anew, e.g., [2]), *circular and cumulative causation* (CCC) (understanding the real dynamic and self-reinforcing aspects of economic phenomena, e.g., [5]), and *network effects and cascading effects* (modelling the economy as a network of entities connected by inter-relationships) [3, 6, 20, 46]; (2) *Agent-based models* potentially present a way to model the financial economy as a complex system, while taking human adaptation and learning into account [7, 22, 23, 41]; (3) *Behavioural economics* addresses the effects of social, cognitive, and emotional factors on the economic decisions of individuals [3, 38]; (4) *Experimental economics* is the application of experimental methods to study economic questions. Data collected in experiments are used to test the validity of economic theories, quantify the effects, and illuminate market mechanisms [48].

Following this movement away from traditional economic models, in the research presented here, markets are modelled using an approach that straddles the interdisciplinary boundary between experimental economics and agent based computational economics. Controlled real-time experiments between human traders and automated financial trading agents (henceforth, referred to simply as *agents*; or alternatively as *robots*) provide a novel perspective on real-world markets, through which we can hope to better understand, and better regulate for, their complex dynamics.

2.2 Broken Markets: Flash Crashes and Subsecond Fractures

As algorithmic trading has become common over the past decade, automated trading systems (ATS) have been developed with truly super-human performance; assimilating and processing huge quantities of data, making trading decisions, and executing them, on subsecond timescales. This has enabled what is known as *high-frequency trading* (HFT), where ATS take positions in the market (e.g., by buying a block of shares) for a very short period of perhaps one or two seconds or less, before reversing the position (e.g., selling the block of shares); each such transaction may generate relatively small profit measured in cents, but by doing this constantly and repeatedly throughout the day, steady streams of significant profit can be generated. For accounts of recent technology developments in the financial markets, see [1, 29, 40, 44].

In February 2012, Johnson et al. [35] published a working paper—later revised for publication in Nature Scientific Reports [36]—that immediately received widespread media attention, including coverage in New Scientist [26], Wired [39], and Financial News [45]. Having analysed millisecond-by-millisecond stock-price movements over a five year period between 2006 and 2011, Johnson et al. argued that there is evidence for a step-change or *phase transition* in the behaviour of financial markets at the subsecond time-scale. At the point of this transition—approximately equal to human response times—the market dynamics switch from a domain where humans and automated *robot* (i.e., *agent*) trading systems freely interact with one another, to a domain newly-identified by Johnson et al. in which humans cannot participate and where all transactions result from robots interacting only among themselves, with no human traders involved.³ Here, we refer to this abrupt system-wide transition from mixed human-algorithm phase to a new all-algorithm phase, the *robot phase transition* (RPT).

At subsecond timescales, below the robot transition, the robot-only market exhibits *fractures*—ultrafast extreme events (UEEs) in Johnson et al.’s parlance, akin to mini flash-crashes—that are undesirable, little understood, and intriguingly ap-

³ The primary reason for no human involvement on these timescales is not because of granularity in decision making—i.e., limitations in human abilities to process information, e.g., [12]—but rather that humans are simply too slow to react to events happening, quite literally, in the blink of an eye.

pear to be linked to longer-term instability of the market as a whole. In Johnson et al.'s words, “[w]e find 18,520 crashes and spikes with durations less than 1500 ms in our dataset... We define a crash (or spike) as an occurrence of the stock price ticking down (or up) at least ten times before ticking up (or down) and the price change exceeding 0.8% of the initial price... Their rapid subsecond speed and recovery... suggests [UEEs are] unlikely to be driven by exogenous news arrival” [36].

In other words, while fractures are relatively rare events at human time scales—those above the RPT—at time scales below the RPT, fractures are commonplace, occurring many thousands of times over a five year period (equivalent to more than ten per day when averaged uniformly). This is interesting. The price discovery mechanism of markets is generally assumed to be driven by the actions of buyers and sellers acting on external information, or news. For instance, the announcement of poor quarterly profits, a new takeover bid, or civil unrest in an oil producing region, will each affect the sentiment of buyers and sellers, leading to a shift in price of financial instruments. The prevalence of ATS means that markets can now absorb new information rapidly, so it is not unusual for prices to shift within (milli)seconds of a news announcement. However, fractures are characterised by a shift in price followed by an immediate recovery, or inverse shift (e.g., a spike from \$100 to \$101; returning to \$100). To be driven by news, therefore, fractures would require multiple news stories to be announced in quick succession, with opposing sentiment (positive/negative) of roughly equal net weighting. The speed and frequency of fractures makes this highly unlikely. Therefore, fractures must be driven by an endogenous process resulting from the interaction dynamics of traders in the market. Since fractures tend to occur only below the RPT, when trading is dominated by robots, it is reasonable to conclude that they are a direct result of the interaction dynamics of HFT robot strategies.

What Johnson et al. have identified is a phase transition in the behaviour of markets in the temporal domain caused by fragmentation of market participants—i.e., at time scales below the RPT, the only active market participants are HFT robots, and the interactions between these robots directly results in fractures that are not observed over longer time scales above the RPT. Intriguingly, however, Johnson et al. also observe a correlation between the frequency of fractures and global instability of markets over much longer time scales. This suggests that there may be a causal link between subsecond fractures and market crashes. “[Further, data] suggests that there may indeed be a degree of causality between propagating cascades of UEEs and subsequent global instability, despite the huge difference in their respective timescales... [Analysis] demonstrates a coupling between extreme market behaviours below the human response time and slower global instabilities above it, and shows how machine and human worlds can become entwined across timescales from milliseconds to months... Our findings are consistent with an emerging ecology of competitive machines featuring ‘crowds’ of predatory algorithms, and highlight the need for a new scientific theory of subsecond financial phenomena” [36].

This discovery has the potential for significant impact in the global financial markets. If short term micro-effects (fractures) can indeed give some indication of longer-term macro-scale behaviour (e.g., market crashes) then it is perhaps possible

that new methods for monitoring the stability of markets could be developed—e.g., using fractures as early-warning systems for impending market crashes. Further, if we can better understand the causes of fractures and develop methods to avoid their occurrence, then long-term market instability will also be reduced. This provides motivation for our research. To understand fractures, the first step is to model the RPT.

Here, we report on using a complementary approach to the historical data analysis employed by Johnson et al. [35, 36]. We conduct laboratory-style experiments where human traders interact with algorithmic trading agents (i.e., robots) in minimal experimental models of electronic financial markets using Marco De Luca’s *OpEx* artificial financial exchange (for technical platform details, see [19, pp. 26–33]). Our aim is to see whether correlates of the two regimes suggested by Johnson et al. can occur under controlled laboratory conditions—i.e., we attempt to *synthesize* the RPT, such that we hope to observe the market transition from a regime of mixed human-robot trading, to a regime of robot-only trading.

3 Background

Experimental human-only markets have a rich history dating back to Vernon Smith’s seminal 1960’s research [48]. “Before Smith’s experiments, it was widely believed that the competitive predictions of supply/demand intersections required very large numbers of well-informed traders. Smith showed that competitive efficient outcomes could be observed with surprisingly small numbers of traders, each with no direct knowledge of the others’ costs or values” [32]. This was a significant finding, and it has spawned the entire field of experimental economics; whereby markets are studied by allowing the market equilibration process to *emerge* from the interacting population of actors (humans and/or agents), rather than assuming an *ideal* market that is trading at the theoretical equilibrium. By measuring the distance between the experimental equilibrium and the theoretical equilibrium, one can quantify the *performance* of the market. Further, by altering the rules of interaction (the market mechanism) and varying the market participants (human or agent), one can begin to understand and quantify the relative effects of each. This is a powerful approach and it is one that we adopt for our experimental research.⁴

The following sections present a detailed background. Section 3.1 introduces the continuous double auction mechanism used for experiments; Section 3.2 provides metrics for evaluating the performance of markets; and Section 3.3 presents a review of previous human-agent experimental studies.

⁴ For a more thorough background and literature review, refer to [19, pp. 6–25].

3.1 *The Continuous Double Auction*

An auction is a mechanism whereby sellers and buyers come together and agree on a transaction price. Many auction mechanisms exist, each governed by a different set of rules. In this chapter, we focus on the *Continuous Double Auction* (CDA), the most widely used auction mechanism and the one used to control all the world's major financial exchanges. The CDA enables buyers and sellers to freely and independently exchange quotes at any time. Transactions occur when a seller accepts a buyer's *bid* (an offer to buy), or when a buyer accepts a seller's *ask* (an offer to sell). Although it is possible for any seller to accept any buyer's bid, and *vice-versa*, it is in both of their interests to get the best deal possible at any point in time. Thus, transactions execute with a counter party that offers the most competitive quote.

Vernon Smith explored the dynamics of CDA markets in a series of Nobel Prize winning experiments using small groups of human participants [47]. Splitting participants evenly into a group of buyers and a group of sellers, Smith handed out a single card (an *assignment*) to each buyer and seller with a single *limit price* written on each, known only to that individual. The limit price on the card for buyers (sellers) represented the maximum (minimum) price they were willing to pay (accept) for a fictitious commodity. Participants were given strict instructions to not bid (ask) a price higher (lower) than that shown on their card, and were encouraged to bid lower (ask higher) than this price, regarding any difference between the price on the card and the price achieved in the market as profit.

Experiments were split into a number of *trading days*, each typically lasting a few minutes. At any point during the trading day, a buyer or seller could raise their hand and announce a quote. When a seller and a buyer agreed on a quote, a transaction was made. At the end of each trading day, all stock (sellers assignment cards) and money (buyer assignment cards) was recalled, and then reallocated anew at the start of the next trading day. By controlling the limit prices allocated to participants, Smith was able to control the market's supply and demand schedules. Smith found that, typically after a couple of trading days, human traders achieved very close to 100% allocative efficiency; a measure of the percentage of profit in relation to the maximum theoretical profit available (see Section 3.2). This was a significant result: few people had believed that a very small number of inexperienced, self-interested participants could effectively self-equilibrate.

3.2 *Measuring Market Performance*

An *ideal* market can be perfectly described by the aggregate quantity supplied by sellers and the aggregate quantity demanded by buyers at every price-point (i.e., the market's supply and demand schedules; see Fig. 1). As prices increase, in general there is a tendency for supply to increase, with increased potential revenues from sales encouraging more sellers to enter the market; while, at the same time, there is a tendency for demand to decrease as buyers look to spend their money elsewhere.

Fig. 1 Supply and Demand curves (here illustrated as straight lines) show the quantities supplied by sellers and demanded by buyers at every price-point. In general, as price increases, the quantity supplied increases and the quantity demanded falls. The point at which the two curves intersect is the theoretical equilibrium point; where Q_0 is the equilibrium quantity and P_0 is the equilibrium price.

At some price-point, the quantity demanded will equal the quantity supplied. This is the theoretical market equilibrium. An idealised theoretical market has a *market equilibrium* price and quantity (P_0, Q_0) determined by the intersection between the supply and demand schedules. The dynamics of competition in the market will tend to drive transactions toward this partial equilibrium point.⁵ For all prices above P_0 , supply will exceed demand, forcing suppliers to reduce their prices to make a trade; whereas for all prices below P_0 , demand exceeds supply, forcing buyers to increase their price to make a trade. Any quantity demanded or supplied below Q_0 is called *intra-marginal*; all quantity demanded or supplied in excess of Q_0 , is called *extra-marginal*. In an ideal market, all intra-marginal units and no extra-marginal units are expected to trade.

In the real world, markets are not ideal. They will always trade away from equilibrium at least some of the time. We can use metrics to calculate the *performance* of a market by how far from ideal equilibrium it trades. In this chapter, we make use of the following metrics:

Smith's Alpha

Following Vernon Smith [47], we measure the equilibration (equilibrium-finding) behaviour of markets using the coefficient of convergence, α , defined as the root mean square difference between each of n transaction prices, p_i (for $i = 1 \dots n$) over some period, and the P_0 value for that period, expressed as a percentage of the equilibrium price:

⁵ The micro-economic supply and demand model presented only considers a single commodity, *ceteris paribus*, and is therefore a partial equilibrium model. The market is considered independently from other markets, so this is not a general equilibrium model.

$$\alpha = \frac{100}{P_0} \sqrt{\frac{1}{n} \sum_{i=1}^n (p_i - P_0)^2} \quad (1)$$

In essence, α captures the standard deviation of trade prices about the theoretical equilibrium. A low value of α is desirable, indicating trading close to P_0 .

Allocative Efficiency

For each trader, i , the maximum theoretical profit available, π_i^* , is the difference between the price they are prepared to pay (their *limit price*) and the theoretical market equilibrium price, P_0 . Efficiency, E , is used to calculate the performance of a group of n traders as the mean ratio of realised profit, π_i , to theoretical profit, π_i^* :

$$E = \frac{1}{n} \sum_{i=1}^n \frac{\pi_i}{\pi_i^*} \quad (2)$$

As profit values cannot go below zero (traders in these experiments are not allowed to enter into loss-making deals; although that constraint can easily be relaxed), a value of 1.0 indicates that the group has earned the maximum theoretical profit available, π_i^* , on all trades. A value below 1.0 indicates that some opportunities have been missed. Finally, a value above 1.0 means that additional profit has been made by taking advantage of a trading counterparty's willingness to trade away from P_0 . So, for example, a group of sellers might record an allocative efficiency of, say, 1.2 if their counterparties (a group of buyers) consistently enter into transactions at prices greater than P_0 ; in such a situation, the buyers' allocative efficiency would not be more than 0.8.

Profit Dispersion

Profit dispersion is a measure of the extent to which the profit/utility generated by a group of traders in the market differs from the profit that would be expected of them if all transactions took place at the equilibrium price, P_0 . For a group of n traders, profit dispersion is calculated as the root mean square difference between the profits achieved, π_i , by each trader, i , and the maximum theoretical profit available, π_i^* :

$$\pi_{disp} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\pi_i - \pi_i^*)^2} \quad (3)$$

Low values of π_{disp} indicate that traders are extracting actual profits close to profits available when all trades take place at the equilibrium price P_0 . In contrast, higher values of π_{disp} indicate that traders' profits differ from those expected at equilibrium. Since zero-sum effects between buyers and sellers do not mask profit dispersion, this statistic is attractive [28].

Delta Profit

Delta profit is used to calculate the difference in profit maximising performance between two groups, x and y , as a percentage difference relative to the mean profit of the two groups, π_x, π_y :

$$\Delta P(x - y) = \frac{2(\pi_x - \pi_y)}{\pi_x + \pi_y} \quad (4)$$

Delta profit directly measures the difference in profit gained by two groups. In a perfect market, we expect $\Delta P(x - y) = 0$, with both groups trading at the equilibrium price P_0 . A positive (negative) value indicates that group x secures more (less) profit than group y . Using this measure enables us to determine which, if either, of the two groups competitively outperforms the other.

3.3 Human vs. Agent Experimental Economics

In 1993, after three decades of human-only experimental economics, a landmark paper involving a mix of traditional human experimental economics and software-agent market studies was published in the *Journal of Political Economy* by Gode and Sunder (G&S) [28]. G&S were interested in understanding how much of the efficiency of the CDA is due to the intelligence of traders, and how much is due to the organisation of the market. To test this, G&S introduced a very simple *Zero Intelligence Constrained* (ZIC) trading agent that generate random bid or ask prices drawn from a uniform distribution, subject to the constraint that prices generated cannot be loss-making—i.e., sell prices are equal or above limit price, buy prices are equal or below limit price. G&S performed a series of ZIC-human experiments, with results demonstrating that the simple ZIC agents produced convergence towards the theoretical equilibrium and had human-like scores for allocative efficiency (equation 2); suggesting that market convergence toward theoretical equilibrium is an emergent property of the CDA market mechanism and not the intelligence of the traders. Indeed, G&S found that the only way to differentiate the performance of humans and ZIC traders was by using their profit dispersion statistics (equation 3). These results were striking and attracted considerable attention.

In 1997, Dave Cliff [13] presented the first detailed mathematical analysis and replication of G&S's results. Results demonstrated that the ability of ZIC traders to converge on equilibrium was dependent on the shape of the market's demand and supply curves. In particular, ZIC traders were unable to equilibrate when acting in markets with demand and supply curves very different to those used by G&S. To address this issue, Cliff developed the *Zero Intelligence Plus* (ZIP) trading algorithm. Rather than issuing randomly generated bid and ask prices in the manner of ZIC, Cliff's ZIP agents contain an internal profit margin from which bid and ask prices are calculated. When a buyer (seller) sees transactions happen at a price below (above) the trader's current bid (ask) price, profit margin is raised, thus resulting in a lower

(higher) bid (ask) price. Conversely, a buyer's (seller's) profit margin is lowered when order and transaction prices indicate that the buyer (seller) will need to raise (lower) bid (ask) price in order to transact [13, p.43]. The size of ZIP's profit margin update is determined using a well established machine learning mechanism (derived from the Widrow-Hoff *Delta rule* [56]). Cliff's autonomous and adaptive ZIP agents were shown to display human-like efficiency and equilibration behaviours in all markets, irrespective of the shape of demand and supply.

Around the same time that ZIP was introduced, economists Steve Gjerstad and his former PhD supervisor John Dickhaut independently developed a trading algorithm that was later named GD after the inventors [27]. Using observed market activity—frequencies of bids, asks, accepted bids and accepted asks—resulting in the most recent L transactions (where $L = 5$ in the original study), GD traders calculate a private, subjective “belief” of the probability that a counterparty will accept each quote price. The belief function is extended over all prices by applying cubic-spline interpolation between observed prices (although it has previously been suggested that using *any* smooth interpolation method is likely to suffice [19, p.17]). To trade, GD quotes a price to buy or sell that maximises expected surplus, calculated as price multiplied by the belief function's probability of a quote being accepted at that price. Simulated markets containing GD agents were shown to converge to the competitive equilibrium price and allocation in a fashion that closely resembled human equilibration in symmetric markets, but with greater efficiency than human traders achieved [27]. A modified GD (MGD) algorithm, where the belief function of bid (ask) prices below (above) the previous lowest (highest) transaction price was set to probability zero, was later introduced to counter unwanted price volatility.

In 2001, a series of experiments were performed to compare ZIP and MGD in real-time heterogeneous markets [52]. MGD was shown to outperform ZIP. Also in 2001, the first ever human-agent experiments—with MGD and ZIP competing in the same market as human traders—were performed by Das et al., a team from IBM [15]. Results had two major conclusions: (a) firstly, mixed human-agent markets were off-equilibrium—somehow the mixture of humans and agents in the market reduce the ability of the CDA to equilibrate; (b) secondly, in all experiments reported, the efficiency scores of humans was lower than the efficiency scores of agents (both MGD and ZIP). In Das et al.'s own words, “... the successful demonstration of machine superiority in the CDA and other common auctions could have a much more direct and powerful impact—one that might be measured in billions of dollars annually” [15]. This result, demonstrating for the first time in human-algorithmic markets that agents can outperform humans, implied a future financial market system where ATS replace humans at the point of execution.

Despite the growing industry in ATS in real financial markets, in academia there was a surprising lack of further human-agent market experiments over the following decade. In 2003 and 2006, Grossklags & Schmidt [30, 31] performed human-agent market experiments to study the effect that human behaviours are altered by their knowledge of whether or not agent traders are present in the market. In 2011, De Luca & Cliff successfully replicated Das et al.'s results, demonstrating that GDX (an extension of MGD, see [51]) outperforms ZIP in agent-agent and agent-human

markets [17]. They further showed that *Adaptive Aggressive* (AA) agents—a trading agent developed by Vytelingum in 2006 that is loosely based on ZIP, with significant novel extensions including short-term and long-term adaptive components [54, 55]—dominate GDX and ZIP, outperforming both in agent-agent and agent-human markets [18]. This work confirmed AA as the dominant trading-agent algorithm. (For a detailed review of how ZIP and AA have been modified over time, see [49, 50].) More recent human-agent experiments have focused on emotional arousal level of humans, monitoring heart rate over time [57]; and monitoring human emotions via EEG brain data [8].

Complementary research comparing markets containing only humans against markets containing only agents—i.e., human-only or agent-only markets rather than markets in which agents and humans interact—can also shed light on market dynamics. For instance, Huber, Shubik, and Sunder (2010) compare dynamics of three market mechanisms (sell-all, buy-all, and double auction) in markets containing all humans against markets containing all agents. “The results suggest that abstracting away from all institutional details does not help understand dynamic aspects of market behaviour and that inclusion of mechanism differences into theory may enhance our understanding of important aspects of markets and money, and help link conventional analysis with dynamics” [33]. This research stream reinforces the necessity of including market design in our understanding of market dynamics. However, it does not offer the rich interactions between humans and ATS that we observe in real markets, and that only human-agent interaction studies can offer.

4 Methodology

In this Section, the experimental methodology and experimental trading platform (OpEx) are presented. Open Exchange (OpEx) is a real-time financial-market simulator specifically designed to enable economic trading experiments between humans and automated trading algorithms (robots). OpEx was designed and developed by Marco De Luca between 2009–2010 while he was a PhD student at the University of Bristol, and since Feb. 2012 is freely available for open-source download from SourceForge, under the terms of the Creative Commons Public License.⁶ Fig. 2 shows the *Lab-in-a-box* hardware arranged ready for a human-agent trading experiment. For a detailed technical description of the OpEx platform, refer to [19, pp. 26–33].

At the start of each experiment, 6 human participants were seated at a terminal around a rectangular table—with three buyers on one side and three sellers opposite—and given a brief introduction and tutorial to the system (explaining the human trading GUI illustrated in Fig. 3), during which time they were able to make test trades among themselves while no robots were present in the market. Participants were told that their aim during the experiment was to maximise profit by

⁶ OpEx download available at: www.sourceforge.net/projects/open-exchange

Fig. 2 The *Lab-in-a-box* hardware ready to run an Open Exchange (OpEx) human versus agent trading experiment. Six small netbook computers run human trader Sales GUIs, with three buyers (near-side) sitting opposite three sellers (far-side). Net-book clients are networked via Ethernet cable to a network switch for buyers and a network switch for sellers, which in turn are connected to a router. The central exchange and robots servers run on the dedicated hardware server (standing vertically, top-left), which is also networked to the router. Finally, an *Administrator* laptop (top table, centre) is used to configure and run experiments. Photograph: © J. Cartlidge, 2012.

Fig. 3 Trading GUI for a human buyer. New order assignments (or *permits*) arrive over time in the *Client Orders* panel (top-left); and listed in descending order by potential profit. Assignments are selected by double-clicking. This opens a *New Order* dialogue pop-up (top-centre) where bid price and quantity are set before entering the new bid into the market by pressing button *BUY*. The market *Order Book* is displayed top-right, with all bids and asks displayed. Bid orders that the trader currently has live in the market are listed in the *Orders* panel (middle); and can be amended from here by double-clicking. When an order executes it is removed from the orders panel and listed in the *Trades* history panel (bottom). For further GUI screen shots, refer to [9, Appendix C].

trading client orders (assignments; or alternatively named *permits* to distinguish that traders will simultaneously have multiple client orders to work, whereas in the traditional literature, a new assignment would only be received once the previous assignment had been completed) that arrive over time. For further details on the experimental method, refer to [9, pp. 9–11].

Trading Agents (Robots)

Agent-robots are independent software processes running on the multi-core hardware server that also hosts the central exchange server. Since agents can act at any time—there is no central controller coordinating when, or in which order, an agent can act—and since the trading logic of agents does not explicitly include temporal information, in order to stop agents from issuing a rapid stream of quotes, a sleep timer is introduced into the agent architecture. After each action, or decision to not act, an agent will *sleep* for t_s milliseconds before *waking* and deciding upon the next action. We name this the *sleep-wake* cycle of agents. For instance if $t_s = 100$, the sleep-wake cycle is 0.1 seconds. To ensure agents do not miss important events during sleep, agents are also set to wake (i.e., sleep is interrupted) when a new assignment permit is received and/or when an agent is notified about a new trade execution. The parameter t_s is used to configure the “speed” of agents for each experiment.

Trading agents are configured to use the *Adaptive Aggressive* (AA) strategy logic [54, 55], previously shown to be the dominant trading agent in the literature (see Section 3.3). AA agents have short term and long term adaptive components. In the short term, agents use learning parameters β_1 and λ to adapt their order aggressiveness. Over a longer time frame, agents use the moving average of the previous N market transactions and a learning parameter β_2 to estimate the market equilibrium price, $\hat{p}_0$. The *aggressiveness* of AA represents the tendency to accept lower profit for a greater chance of transacting. To achieve this, an agent with high (low) aggression will submit orders better (worse) than the estimated equilibrium price $\hat{p}_0$. For example, a buyer (seller) with high aggression and estimated equilibrium value $\hat{p}_0 = 100$ will submit bids (asks) with price $p > 100$ (price $p < 100$). Aggressiveness of buyers (sellers) increases when transaction prices are higher (lower) than $\hat{p}_0$, and decreases when transaction prices are lower (higher) than $\hat{p}_0$. The Widrow-Hoff mechanism [56] is used by AA to update aggressiveness in a similar way that it is used by ZIP to update profit margin (see Section 3.3). For all experiments reported here, we set parameter values $\beta_1 = 0.5$, $\lambda = 0.05$, $N = 30$, and $\beta_2 = 0.5$. The convergence rate of bids/asks to transaction price is set to $\eta = 3.0$.

Exploring the Effects of Agent Speed on Market Efficiency: April–June 2011

All experiments were run at the University of Bristol between April and July 2011 using postgraduate students in non-financial but analytical subjects (i.e., students

Table 1 Permit schedule for market efficiency experiments. Six permit types are issued to each market participant, depending on their role. For each role (e.g., *Buyer 1*), there are two traders: one human (*Human Buyer 1*) and one robot (*Robot Buyer 1*). Thus, there are 12 traders in the market. Permit values show *limit price*—the maximum value at which to buy, or minimum value at which to sell—and the time-step they are issued (in parentheses). The length of each time-step is 10s, making one full permit cycle duration 170s. During a 20-minute experiment there are 7 full cycles.

	1	2 ^a	3	4	5	6
Buyer 1	350 (0)	250 (4)	220 (7)	190 (09)	150 (14)	140 (16)
Buyer 2	340 (1)	270 (3)	210 (8)	180 (10)	170 (12)	130 (17)
Buyer 3	330 (2)	260 (4)	230 (6)	170 (11)	160 (13)	150 (15)
Seller 1	50 (0)	150 (4)	180 (7)	210 (09)	250 (14)	260 (16)
Seller 2	60 (1)	130 (3)	190 (8)	220 (10)	230 (12)	270 (17)
Seller 3	70 (2)	140 (4)	170 (6)	230 (11)	240 (13)	250 (15)

^a Type 2 permits were accidentally issued to Buyer1/Seller1 at time-step 4 rather than time-step 5.

with skills suitable for a professional career in finance, but with no specific trading knowledge or experience). Participants were paid £20 for participating and a further £40 bonus for making the most profit, and £20 bonus for making the second highest profit. Moving away from the artificial constraint of regular simultaneous replenishments of currency and stock historically used, assignment permits were issued at regular intervals. AA agents had varying sleep-wake cycle: $t_s = 100$, and $t_s = 10,000$. We respectively label these agents AA-0.1 to signify a sleep-wake cycle of 0.1s, and AA-10 to signify a sleep-wake cycle of 10s. A total of 7 experiments were performed, using the assignment permit schedules presented in Table 1. The supply and demand curves generated by these permits are shown in Fig. 4. We can see that for all experiments, $P_0 = 200$ and $Q_0 = 126$. Since each human only participates in one experiment, and since trading agents are re-set at the beginning of each run, traders have no opportunity to learn the fixed value of P_0 over repeated runs. For further details of experimental procedure, see [11].

Exploring the Robot Phase Transition (RPT): March 2012

Twenty-four experiments were run on 21st March, 2012, at Park House Business Centre, Park Street, Bristol, UK. Participants were selected on a first-come basis from the group of students that responded to adverts broadcast to two groups: (1) students enrolled in final year undergraduate and postgraduate module in computer science that includes coverage of the design of automated trading agents; (2) members of the Bristol Investment Society, a body of students interested in pursuing a career in finance. We assume that these students have the knowledge and skills to embark on a career as a trader in a financial institution. Volunteers were paid £25 for participating, and the two participants making the greatest profit received an iPad valued at £400. To reduce the total number of participants required, each group were used in a session of six separate experiments. Therefore, 24 experiments were run

Fig. 4 Stepped supply and demand curves for permit schedule defined in Table 1. Curves show the aggregate quantity that participants are prepared to buy (demand) and sell (supply) at every price point. The point at which the two curves intersect is the theoretical equilibrium point for the market: $P_0 = 200$ is the equilibrium price; and Q_0 is the equilibrium quantity. As there are two traders in each role—one human and one robot—each permit cycle $Q_0 = 2 \times 9 = 18$; and over the seven permit cycles of one full experiment, $Q_0 = 18 \times 7 = 126$. The market is symmetric about P_0 .

using the 24 participants. Between experiments, human participants rotated seats, so each played every role exactly once during the session of 6 experiments. Human roles were purposely mixed between experiment rounds to reduce the opportunity for collusion and counteract any bias in market role. Once again, agents used the AA algorithm with varying sleep-wake cycle; and assignment orders were released into the market at regular intervals.

Table 2 presents the assignment permit schedules used for each experiment, and the full supply and demand curves generated by these permits are plotted in Fig. 5. At each price point—i.e., at each *step* in the *permit* schedule—two assignment permits are sent simultaneously to a human trader and to a robot trader, once every replenishment cycle. For all experiments, permits are allocated in pairs symmetric about P_0 such that the equilibrium is not altered; and the inter-arrival time of permits is 4s. Cycles last 72s and are repeated 8 times during a 10 minute experiment. Therefore, over a full experiment there are $2 \times 8 = 16$ permits issued at each price point. The expected equilibrium number of trades for the market, Q_0 , is 144 intra-marginal units. Each experiment, P_0 is varied in the range 209–272 to stop humans from learning the equilibration properties of the market between experiments. Agents are reset each time and have no access to data from previous experiments. In *cyclical* markets, permits are allocated in strict sequence that is unaltered between cycles. In *random* markets, the permit sequence across the entire run is randomised. For further details on experimental procedure, see [9, 10].

Table 2 Permit schedule for RPT experiments. Six permit types are issued to each market participant, depending on their role. For each role, there is one human and one robot participant. Permit values show *limit price* $- P_0$. Thus, for e.g., if $P_0 = 100$, a permit of type 4 to Buyer1 would have a limit price of 91. For buyers, limit prices are the maximum value to bid; and for sellers, limit prices are the minimum value to ask. Numbers in brackets show the time-step sequence in which permits are allocated. Thus, after 11 time-steps, Buyer2 and Seller2 each receive a permit of type 4. For all experiments, the inter-arrival time-step between permits is 4 seconds. Permits are always allocated in pairs, symmetric about P_0 . In cyclical markets, the sequence is repeated 8 times: the last permits are issued to Buyer3 and Seller3 at time 576s; and the experiment ends 24s later. In non-cyclical or ‘random’ markets, the time-step of permits is randomised across the run. Participants receive the same set of permits in both cyclical and random markets, but in a different order.

	1	2	3	4	5	6
Buyer 1	77 (1)	27 (4)	12 (7)	-9 (10)	-14 (13)	-29 (16)
Buyer 2	73 (2)	35 (5)	8 (8)	-5 (11)	-22 (14)	-25 (17)
Buyer 3	69 (3)	31 (6)	16 (9)	-1 (12)	-18 (15)	-33 (18)
Seller 1	-77 (1)	-27 (4)	-12 (7)	9 (10)	14 (13)	29 (16)
Seller 2	-73 (2)	-35 (5)	-8 (8)	5 (11)	22 (14)	25 (17)
Seller 3	-69 (3)	-31 (6)	-16 (9)	1 (12)	18 (15)	33 (18)

Fig. 5 Stepped supply and demand curves for an entire run of the RPT experiments, defined by the permit schedules shown in Table 2. Curves show the aggregate quantity that participants are prepared to buy (demand) and sell (supply) at every price point. The point at which the two curves intersect is the theoretical equilibrium point for the market: $Q_0 = 144$ is the equilibrium quantity; and P_0 is the equilibrium price. Each experiment the value of P_0 is varied in the range 209–272 to avoid humans learning a fixed value of P_0 over repeated trials. The market is symmetric about P_0 .

Fig. 6 Time series of quote and trade prices from a cyclical market containing AA-0.1 agents. The dotted horizontal line represents the theoretical market equilibrium, P_0 . Vertical dotted lines indicate the start of each new permit replenishment cycle.

5 Results

Here, we present empirical results from the two sets of experiments: (a) exploring the robot phase transition, performed in March 2012; and (b) exploring the effects of agent speed on market efficiency, performed in April–June 2011. Throughout this section, for detecting significant differences in location between two samples we use the nonparametric Robust Rank-Order (RRO) test and critical values reported by Feltovich [24]. RRO is particularly useful for small sample statistics of the kind we present here, and is less sensitive to changes in distributional assumptions than the more commonly known Wilcoxon-Mann-Whitney test [24].

5.1 Exploring the Robot Phase Transition: March 2012

Experiments were run using AA agents with sleep-wake cycle times (in seconds) $t_s = 0.1$ (AA-0.1), $t_s = 1$ (AA-1), $t_s = 5$ (AA-5), and $t_s = 10$ (AA-10). Of the 24 runs, one experienced partial system failure, so results were omitted. Runs with agent sleep time 5s (AA-5) are also omitted from analysis where no significant effects are found. For further detail of results, see [9, 10].

5.1.1 Market Data

OpEx records time-stamped data for every exchange event. This produces rich datasets containing every quote (orders submitted to the exchange) and trade (orders that execute in the exchange) in a market. In total, we gathered 4 hours of trading data across the four one-hour sessions, but for brevity we explore only a small set of indicative results here; however, for completeness, further datasets are presented in [9, Appendix A]. Fig. 6 plots time series of quotes and trades for a cyclical market containing AA-0.1 agents. The dotted horizontal line represents the theoretical market equilibrium, P_0 , and vertical dotted lines indicate the start of each new permit replenishment cycle (every 72s). We see the majority of trading activity (denoted by filled markers) is largely clustered in the first half of each permit-replenishment cycle; this correlates with the phase in which intra-marginal units are allocated and trades are easiest to execute. After the initial *exploratory* period, execution prices tend toward P_0 in subsequent cycles. In the initial period, robots (blue diamonds for sellers; blue inverted triangle for buyers) explore the space of prices. In subsequent periods, robots quote much closer to equilibrium. Agent quotes are densely clustered near to the start of each period, during the phase that intra-marginal units are allocated. In contrast, humans (red squares for sellers; red triangles for buyers) tend to enter exploratory quotes throughout the market’s open period.

5.1.2 Smith’s α

We can see the equilibration behaviour of the markets more clearly by plotting Smith’s α for each cycle period. In Fig. 7 we see mean α ($\pm 95\%$ confidence interval) plotted for cyclical and random markets. Under both conditions, α follows a similar pattern, tending to approx 1% by market close. However, in the first period, cyclical markets produce significantly greater α than random markets (RRO, $p < 0.0005$). This is due to the sequential order allocation of permits in cyclical markets, where limit prices farthest from equilibrium are allocated first. This enables exploratory shouts and trades to occur far from equilibrium. In comparison, in random markets, permits are not ordered by limit price, thus making it likely that limit prices of early orders are closer to equilibrium than they are in cyclical markets.

5.1.3 Allocative Efficiency

Tables 3 and 4 display the mean allocative efficiency of agents, humans, and the whole market grouped by agent type and market type, respectively. Across all groupings, $E(\text{agents}) > E(\text{humans})$. However, when grouped by robot type (Table 3), the difference is only significant for AA-0.1 and AA-5 (RRO, $0.051 < p < 0.104$). When grouped by market type (Table 4), $E(\text{agents}) > E(\text{humans})$ is significant in cyclical markets (RRO, $0.05 < p < 0.1$), random markets (RRO, $0.05 < p < 0.1$),

Fig. 7 Mean α ($\pm 95\%$ confidence interval) plotted using log scale for results grouped by market type. In cyclical markets, α values are significantly higher than in random markets during the initial period (RRO, $p < 0.0005$). In subsequent periods all markets equilibrate to $\alpha < 1\%$ with no statistical difference between groups.

Table 3 Efficiency and profit for runs grouped by robot type. Agents achieve greater efficiency $E(\text{agents}) > E(\text{humans})$, and greater profit $\Delta P(\text{agents} - \text{humans}) > 0$, under all conditions.

Robot Type	Trials	$E(\text{agents})$	$E(\text{humans})$	$E(\text{market})$	$\Delta P(\text{agents} - \text{humans})$
AA-0.1	6	0.992	0.975	0.984	1.8%
AA-1	5	0.991	0.977	0.984	1.4%
AA-5	6	0.990	0.972	0.981	1.8%
AA-10	6	0.985	0.981	0.983	0.4%
All	23	0.989	0.976	0.983	1.34%

and across all 23 runs (RRO, $0.01 < p < 0.025$). These results suggest that agents outperform humans.

In Table 3, it can be seen that as sleep time increases the efficiency of agents decreases (column 3, top-to-bottom). Conversely, the efficiency of humans tends to increase as sleep time increases (column 4, top-to-bottom). However, none of these differences are statistically significant (RRO, $p > 0.104$). In Table 4, efficiency of agents, humans, and the market as a whole are all higher when permit schedules are issued cyclically rather than randomly, suggesting that cyclical markets lead to greater efficiency. However, these differences are also not statistically significant (RRO, $p > 0.104$). Finally, when comparing $E(\text{agents})$ grouped by robot type using only data from cyclical markets (data not shown), AA-0.1 robots attain a significantly higher efficiency than AA-1 (RRO, $p = 0.05$), AA-5 (RRO, $p = 0.05$), and AA-10 (RRO $p = 0.1$); suggesting that the very fastest robots are most efficient in cyclical markets.

Table 4 Efficiency and profit for runs grouped by market type. Agents achieve greater efficiency $E(\text{agents}) > E(\text{humans})$, and greater profit $\Delta P(\text{agents} - \text{humans}) > 0$, under all conditions.

Market Type	Trials	$E(\text{agents})$	$E(\text{humans})$	$E(\text{market})$	$\Delta P(\text{agents} - \text{humans})$
Cyclical	12	0.991	0.978	0.985	1.32%
Random	11	0.987	0.974	0.981	1.36%
All	23	0.989	0.976	0.983	1.34%

Table 5 Profit dispersion for runs grouped by market type. Profit dispersion in random markets is significantly lower than in cyclical markets for agents $\pi_{disp}(\text{agents})$, humans $\pi_{disp}(\text{humans})$, and the whole market $\pi_{disp}(\text{market})$.

Market Type	Trials	$\pi_{disp}(\text{agents})$	$\pi_{disp}(\text{humans})$	$\pi_{disp}(\text{market})$
Cyclical	12	89.6	85.4	88.6
Random	11	50.2	57.2	55.6
All	23	70.0	71.9	72.8

5.1.4 Delta Profit

From the right-hand columns of Tables 3 and 4, it can be seen that agents achieve greater profit than humans under all conditions, i.e., $\Delta P(\text{agents} - \text{humans}) > 0$. Using data across all 23 runs, the null hypothesis $H_0 : \Delta P(\text{agents} - \text{humans}) \leq 0$ is rejected (t-test, $p = 0.0137$). Therefore, the profit of agents is significantly greater than the profit of humans, i.e., agents outperform humans across all runs. Differences in $\Delta P(\text{agents} - \text{humans})$ between robot groupings and market groupings are not significant (RRO, $p > 0.104$).

5.1.5 Profit Dispersion

Table 5 shows the profit dispersion of agents $\pi_{disp}(\text{agents})$, humans $\pi_{disp}(\text{humans})$, and the whole market $\pi_{disp}(\text{market})$, for runs grouped by market type. It is clear that varying between cyclical and random permit schedules has a significant effect on profit dispersion, with random markets having significantly lower profit dispersion of agents (RRO, $0.001 < p < 0.005$), significantly lower profit dispersion of humans (RRO, $0.025 < p < 0.05$), and significantly lower profit dispersion of the market as a whole (RRO, $0.005 < p < 0.01$). These results indicate that traders in random markets are extracting actual profits closer to profits available when all trades take place at the equilibrium price, P_0 ; i.e., random markets are trading closer to equilibrium, likely due to the significant difference in α during the initial trading period (see Section 5.1.2. When grouping data by robot type (not shown), there is no significant difference in profit dispersion of agents, humans, or markets (RRO, $p > 0.104$).

Fig. 8 Box plot showing the percentage of homogeneous counterparty executions (i.e., trades between two humans, or between two agents). In a fully mixed market, there is an equal chance that a counterparty will be agent or human; denoted by the horizontal dotted line, H_0 . When agents act and react at time scales equivalent to humans (i.e., when sleep time is 1s or 10s), counterparties are selected randomly—i.e., there is a mixed market and H_0 is not rejected ($p > 0.1$). However, when agents act and react at super-human timescales (i.e., when sleep time is 0.1s), counterparties are more likely to be homogeneous— H_0 is rejected ($p < 0.0005$). This result suggests that, even under simple laboratory conditions, when agents act at super-human speeds the market fragments.

5.1.6 Execution Counterparties

Let aa denote a trade between agent buyer and agent seller, hh a trade between human buyer and human seller, ah a trade between agent buyer and human seller, and ha a trade between human buyer and agent seller. Then, assuming a fully mixed market where any buyer (seller) can independently and anonymously trade with any seller (buyer), we generate null hypothesis, H_0 : the proportion of trades with homogeneous counterparties— aa trades or hh trades—should be 50%. More formally:

$$H_0 : \frac{\Sigma aa + \Sigma hh}{\Sigma aa + \Sigma hh + \Sigma ah + \Sigma ha} = 0.5$$

In Fig. 8, box-plots present the proportion of homogeneous counterparty trades for markets grouped by robot type (AA-0.1, AA-1, and AA-10); the horizontal dotted line represents the H_0 value of 50%. It can clearly be seen that the proportion of homogeneous counterparty trades for markets containing AA-0.1 robots is significantly greater than 50%; and H_0 is rejected (t-test, $p < 0.0005$). In contrast, for markets containing AA-1 and AA-10 robots, H_0 is not rejected at the 10% level of significance. This suggests that for the fastest agents (AA-0.1) the market tends to fragment, with humans trading with humans and robots trading with robots more than would be expected by chance. There also appears to be an inverse relation-

Fig. 9 Mean Smith's α ($\pm 95\%$ confidence interval) plotted using log scale for results grouped by robot type. In markets containing fast AA-0.1 robots, α values are significantly higher than in markets containing slow AA-10 robots. After the initial period, all markets equilibrate to $\alpha < 10\%$.

ship between robot sleep time and proportion of homogeneous counterparty trades. RRO tests show that the proportion of homogeneous counterparty trades in AA-0.1 markets is significantly higher than AA-1 markets ($p < 0.051$) and AA-10 markets ($p = 0.0011$); and for AA-1 markets the proportion is significantly higher than AA-10 ($p < 0.104$). For full detail of RRO analysis of execution counterparties, see [9, Appendix A.2.1].

5.2 Effect of Agent Speed on Market Efficiency: April–June 2011

Experiments were run using AA agents with sleep-wake cycle times (in seconds) $t_s = 0.1$ (AA-0.1) and $t_s = 10$ (AA-10). A total of 8 experiments were performed. However, during one experiment, a human participant began feeling unwell and could no longer take part; so results for this trial are omitted. Here, we present results of the remaining 7 experiments. For further detail of results, see [11].

5.2.1 Smith's α

Fig. 9 plots mean α ($\pm 95\%$ confidence interval) for each permit replenishment cycle, grouped by robot type: AA-0.1 and AA-10. Under both conditions, $\alpha > 10\%$ in the initial period, and then equilibrates to a value $\alpha < 10\%$. For every period, i , mean value α_i is lower for markets containing AA-10 robots than it is for markets containing AA-0.1 robots. Using RRO, this difference is significant at every period: α_1 ($p < 0.029$), α_2 ($p < 0.029$), α_3 ($p < 0.029$), α_4 ($p < 0.029$), α_5 ($p < 0.114$), α_6 ($p < 0.057$), α_7 ($p < 0.029$). This suggests that markets with slower agents (AA-10) are able to equilibrate better than markets with faster agents (AA-0.1).

Table 6 Efficiency and profit for runs grouped by robot type. Agents achieve greater efficiency $E(\text{agents}) > E(\text{humans})$, and greater profit $\Delta P(\text{agents} - \text{humans}) > 0$, when robots are fast AA-0.1. In contrast, humans achieve greater efficiency $E(\text{humans}) > E(\text{agents})$, and greater profit $\Delta P(\text{agents} - \text{humans}) < 0$, when robots are slow AA-10.

	Trials	$E(\text{agents})$	$E(\text{humans})$	$E(\text{market})$	$\Delta P(\text{agents} - \text{humans})$
AA-0.1	3	0.966	0.906	0.936	3.2%
AA-10	4	0.957	0.963	0.960	-0.3%

Table 7 Profit dispersion for runs grouped by market type. In markets with slow AA-10 robots, profit dispersion of agents $\pi_{disp}(\text{agents})$, humans $\pi_{disp}(\text{humans})$, and the whole market $\pi_{disp}(\text{market})$ is significantly lower than in markets with fast AA-0.1 robots.

	Trials	$\pi_{disp}(\text{agents})$	$\pi_{disp}(\text{humans})$	$\pi_{disp}(\text{market})$
AA-0.1	3	105	236	185
AA-10	4	100	164	139

5.2.2 Allocative Efficiency and Delta Profit

Table 6 presents mean allocative efficiency and delta profit for runs grouped by robot type. The efficiency of agents is similar under both conditions, with no statistical difference (RRO, $p > 0.114$). However, runs with slow AA-10 robots result in significantly higher efficiency of humans, (RRO, $0.114 < p < 0.029$), and significantly higher efficiency of the whole market, (RRO, $0.114 < p < 0.029$). In markets containing slower AA-10 robots, humans are able to secure greater profit than agents $\Delta P(\text{agents} - \text{humans}) < 0$; whereas in markets containing fast AA-0.1 robots, agents secure more profit than humans $\Delta P(\text{agents} - \text{humans}) > 0$. However, the difference in delta profit between the two groups is not significant (RRO, $p > 0.114$).

These data provide evidence that markets containing fast AA-0.1 robots are less efficient than markets containing slow AA-10 robots. However, this does not imply that AA-10 outperform AA-0.1, as their efficiency shows no significant difference. Rather, we see that humans perform more poorly when competing in markets containing faster trader-agents, resulting in lower efficiency for the market as a whole.

5.2.3 Profit Dispersion

Table 7 presents profit dispersion for runs grouped by robot type. In markets containing fast AA-0.1 robots, profit dispersion is significantly higher for agents (RRO, $p < 0.114$), humans (RRO, $p < 0.114$), and the market as a whole (RRO, $0.029 < p < 0.057$). These data provide evidence that fast AA-0.1 agents result in higher profit dispersion than slow AA-10 agents; an undesirable result.

6 Discussion

Here we discuss results presented in the previous section. First, in Section 6.1 we summarize the main results that hold across all our market experiments presented in Section 5.1. Subsequently, in Section 6.2 we discuss potentially conflicting results from experiments presented in Section 5.2. Finally, in Section 6.3 we discuss results that demonstrate significant differences between cyclical and random markets.

6.1 Evidence for the Robot Phase Transition (RPT)

Results in Section 5.1.2 show that, across all markets, α values start relatively high ($\alpha \approx 10\%$) as traders *explore* the space of prices, and then quickly reduce, with markets tending to an equilibration level of $\alpha \approx 1\%$. This suggests that the market's price-discovery is readily finding values close to P_0 . Further, in Sections 5.1.3 and 5.1.4, agents are shown to consistently outperform humans, securing greater allocative efficiency $E(\text{agents}) > E(\text{humans})$, and gaining greater profit $\Delta P(\text{agents} - \text{humans}) > 0$. These results demonstrate a well-functioning robot-human market trading near equilibrium, with robots out-competing humans. This is an interesting result, but for our purpose of exploring the RPT described by [35, 36] it only serves as demonstrative proof that our experimental markets are performing as we would expect. The real interest lies in whether we can observe a phase transition between two regimes: one dominated by robot-robot interactions, and one dominated by human-robot interactions. We seek evidence of this by observing the proportion of homogeneous counterparties within a market; that is, the number of trade executions that occur between a pair of humans or a pair of robots, as a proportion of all market trades. Since traders interact anonymously via the exchange, there can be no preferential selection of counterparties. Therefore, every buyer (seller) has an equal opportunity to trade with every seller (buyer), as long as both have a pending assignment. The experimental market is configured to have an equal number of robot traders and human traders, and an equal number of identical assignments are issued to both groups. Hence, in the limit, we should expect 50% of trade counterparties to be homogeneous (both robot, or both human), and 50% to be heterogeneous (one robot and one human), as traders execute with counterparties drawn at random from the population.

From Section 5.1.6, our results demonstrate that for markets containing AA-0.1 robots (with sleep-wake cycle $t_s = 100\text{ms}$; faster than human response time), the proportion of homogeneous counterparties is significantly higher than we would expect in a mixed market; whereas for markets containing robots AA-1 ($t_s = 1,000\text{ms}$; a similar magnitude to human response time) and AA-10 ($t_s = 10,000\text{ms}$; slower than human response time), the proportion of homogeneous counterparties cannot be significantly differentiated from 50%. We present this as tentative first evidence for a robot-phase transition in experimental markets with a boundary between 100 milliseconds and 1 second; although, in our experiments the effects of increasing

robot speed appear to give a progressive response rather than a step-change. However, we feel obliged to caveat this result as non-conclusive proof until further experiments have been run, and until our results have been independently replicated.

The careful reader may have noticed that the results presented have not demonstrated *fractures*—ultrafast series of multiple sequential up-tick or down-tick trades that cause market price to deviate rapidly from equilibrium and then just as quickly return—phenomena that [35, 36] revealed in real market data. Since we are constraining market participants to one role (as buyer, or seller) and strictly controlling the flow of orders into the market and limit prices of trades, the simple markets we have constructed do not have the capacity to demonstrate such fractures. For this reason, we use the proportion of homogeneous counterparties as proxy evidence for the robot phase transition.

6.2 Fast Agents and Market Efficiency

Results presented in Section 5.2 compare markets containing fast AA-0.1 robots to markets containing slower AA-10 robots. It is shown that markets containing fast AA-0.1 robots have higher α (Section 5.2.1), lower allocative efficiency (Section 5.2.2), and higher profit dispersion (Section 5.2.3). Together, these facts suggest that when agents act at super-human speeds, human performance suffers, causing an overall reduction in the efficiency of the market. The reason for this could be, perhaps, that the presence of very fast acting agents causes confusion in humans, resulting in poorer efficiency. If an analogous effect occurs in real financial markets, it may imply that high frequency trading (HFT) can reduce market efficiency.

However, these findings largely contradict the findings presented in Section 5.1 and discussed in Section 6.1; where market equilibration α (Section 5.1.2), market efficiency (Section 5.1.3), and profit dispersion (Section 5.1.5) are shown to be unaffected by robot speed. The reason for this disparity is primarily due to an unanticipated feature (a bug) in the behaviour of AA agents used in the experiments of Section 5.2, that was not discovered at the time (for details see [9, pp. 25–26] and [50, p. 8]). These AA agents included a *spread jumping* rule such that agents will execute against a counterparty in the order-book if the relative spread width (the difference in price between the highest bid and the lowest ask, divided by the mean of the highest bid and lowest ask) is below a relative threshold of $MaxSpread = 15\%$. This is a large, unrealistic threshold; and it was reduced to $MaxSpread = 1\%$ for experiments presented in Section 5.1.

It is reasonable to infer that the spread jumping behaviour of AA agents is directly responsible for the higher α values presented in Section 5.2.1; compared with results for non spread jumping agents shown in Section 5.1.2.⁷ However, when considering

⁷ Some of the variation in α between results presented in Section 5.1.2 and Section 5.2.1 may be explained by the different permit schedules used for the two experiments (compare Tables 1 and 2). However, previous results from a direct comparison using an identical permit schedule to Table 2 show that $MaxSpread = 15\%$ results in higher α than $MaxSpread = 1\%$ [9, Appendix

market efficiency (Section 5.2.2), the explanation is not quite so simple. Despite the *bug* existing in the agent, efficiency for agents is largely unaffected when agent speed is increased; whereas human efficiency drops from a level comparable with agents when sleep-wake cycle time is 10s, to 6% lower than agents when agents act with 0.1s sleep-wake cycle time. Therefore, the effect of the bug on efficiency only affects humans, and does so only when agents act at super-human speeds. We also see a similar affect on the magnitude of profit dispersion (Section 5.2.3), such that $\pi_{disp}(\text{humans})$ is 76% higher in AA-0.1 markets compared with AA-10 markets, whereas $\pi_{disp}(\text{agents})$ is only 5% higher.

This slightly counter-intuitive result is perhaps again further evidence for the RPT. When agents act at super-human speeds, fragmentation in the market means that agents are more likely to trade with other agents. While agents that execute a trade due to the spread jumping bug will lose out, the agent counterparty to the trade will gain; thus cancelling out the negative efficiency effects for agents overall. Human efficiency, however, is negatively affected by the resulting market dynamics, as the market trades away from equilibrium. A similar phenomena has also been observed in a recent pilot study performed at the University of Nottingham Ningbo China (UNNC) in July 2016, using a different agent (ZIP) and performed on a different experimental platform (ExPo2—details forthcoming in future publication).⁸ In the pilot study, agents were allowed to submit loss-making orders into the market (i.e., agent quote price was not constrained by assignment limit price). Interestingly, when agents acted at human-speeds (10s sleep-wake cycle), markets equilibrated as expected. However, when agents acted at super-human speeds (0.1s sleep-wake cycle), the market did not equilibrate to P_0 . This demonstrates that when agents act on human timescales, i.e., above the RPT, the equilibration behaviour of humans can dampen idiosyncratic agent behaviour. However, at super-human timescales (i.e., below the RPT), the cumulative effects of agent behaviour dominate the market. Therefore, as we move from timescales above the RPT to below the RPT, the market transitions from a more efficient to a less efficient domain.

We see similar effects occur in real markets, for example Knight Capital’s fiasco, previously dubbed elsewhere as the *Nightmare on Wall Street*. On August 1st 2012, Knight Capital—formerly the largest US equities trader by volume, trading an average of 128,000 shares per second—started live trading their new Retail Liquidity Provider (RLP) market making software on NYSE. Within 45 minutes, RLP executed 4 million trades across 154 stocks; generating a pre-tax loss of \$440 million. The following day, Knight’s share price collapsed over 70%. Knight subsequently went into administration, before being acquired by Getco, a smaller rival, forming KCG Holdings (for further details, see [4]). It is widely accepted that Knight’s failure was due to repurposing, and inadvertently releasing, deprecated test code that began executing trades deliberately designed to move the market price. In the live markets, and at high frequencies well above the RPT, this resulted in Knight’s RLP effectively trading with itself, but at a loss on either side of the trade. The parallel

B]. Although, a more recent study [16] suggests the opposite result, so there is some uncertainty around this effect.

⁸ ExPo: The Exchange portal: www.exchangeportal.org

here with spread-jumping AA agents is clear; if RLP acted at much lower frequencies, below the RPT, it is likely, perhaps, that the market could have dampened the instability caused. Of further interest is that the market perturbation caused by RLP percolated across a large number of stocks as other automated trading systems reacted to the behaviour. This demonstrates how single stock fractures below the RPT can have wider market impact over longer timescales.

6.3 Realism in Market Experiments: Artefacts or Evidence?

The cyclical-replenishment permit schedules presented in Section 4 approximate real-world markets more poorly than random-replenishment permit schedules. In real markets, demand and supply does not arrive in neat price-ordered cycles. For that reason, where results from cyclical markets (presented in Section 5.1) show a significant effect that is not also present in random markets, we interpret it as an indication that introducing artificial constraints into experimental markets for ease of analysis runs the risk of also introducing artefacts that, because they are statistically significant, can be misleading.

The following relationships were all observed to be statistically significant in cyclical markets and not statistically significant in random markets; providing further support for the argument for *realism* in artificial-market experiment design, previously advanced at length in [19]:

1. Cyclical-replenishment markets have significantly greater α in the first period of trade (see Section 5.1.2). This is a direct consequence of cyclical-replenishment allocating orders in a monotonically decreasing sequence from most profitable to least profitable. As such, the first orders allocated into the market have limit prices far from equilibrium. Since the market is empty, there is no mechanism for price discovery other than trial-and-error exploration; leading to large α . In random-replenishment markets, the initial orders entering the market are drawn at random from the demand and supply schedules. This leads to lower bounds on limit prices and hence lower α . Subsequently, price discovery is led by the order book, resulting in lower α that is statistically similar in both cyclical and random markets.
2. In cyclical-replenishment markets, the efficiency of AA-0.1 robots is significantly higher than the efficiency of the other robot types (see Section 5.1.3). While there is some evidence of an inverse relationship between robot sleep time and robot efficiency across all markets, we infer that this difference is an artefact of cyclical replenishment until further experimental trials can confirm otherwise.
3. In cyclical-replenishment markets, profit dispersion is significantly higher for agents, humans, and the market as a whole (see Section 5.1.5). Since lower profit dispersion is a desirable property of a market, this suggests that the relatively high profit dispersion observed in previous cyclical-replenishment experiments [11, 19] is an artefact of the experimental design.

7 Conclusion

We have presented a series of laboratory experiments between agent traders and human traders in a controlled financial market. Results demonstrate that, despite the simplicity of the market, when agents act on super-human timescales—i.e., when the sleep-wake cycle of agents is 0.1s—the market starts to fragment, such that agents are more likely to trade with agents, and humans are more likely to trade with humans. In contrast, when agents act on human timescales—i.e., when the sleep-wake cycle of agents is 1s, or above—the markets are well mixed, with agents and humans equally likely to trade between themselves and between each other. This transition to a fragmented market from a mixed market intriguingly appears to be linked to market inefficiency, such that below the threshold of human reaction times (i.e., at 0.1s timescale) any idiosyncratic agent behaviours can adversely perturb the market; whereas above the threshold (i.e., at timescales of 1s and above) human interactions help to dampen market perturbations, ensuring better equilibration and efficiency.

This behaviour has parallels with the real financial markets, and in particular, we present this as tantalising evidence for the robot phase transition (RPT), discovered by Johnson et al. [35, 36]. In Johnson et al.’s words, “a remarkable new study by Cliff and Cartledge provides some additional support for our findings. In controlled lab experiments, they found when machines operate on similar timescales to humans (longer than 1s), the ‘lab market’ exhibited an efficient phase (c.f. few extreme price-changes events in our case). By contrast, when machines operated on a timescale faster than the human response time (100 milliseconds) then the market exhibited an inefficient phase (c.f. many extreme price-change events in our case)” [36].

In the final quarter of 2016, a new exchange node containing the first ever intentional delay was introduced in the United States. To achieve a delay of 350 microseconds in signal transmission, the exchange embedded a 38-mile coil of fibre optic cable. The desired intention is to “level out highly asymmetric advantages available to faster participants” in the market [34]. However, the impact this might have at the system level are unknown. To address this, Johnson declares that more academic studies need to focus on subsecond resolution data; and he identifies the work we have reported here as one of the few exceptions in the literature that attempts to understand subsecond behaviours [34].

This work is presented as a demonstration of the utility of using experimental human-agent laboratory controlled markets: (a) to better understand real-world complex financial markets; and (b) to test novel market policies and structures before implementing them in the real world. We hope that we are able to encourage the wider scientific community to pursue more research endeavour using this methodology.

Future Work

For results presented here, we used De Luca’s OpEx experimental trading software, running on the *Lab-in-a-box* hardware; a self-contained wired-LAN containing networked exchange server, netbooks for human participants, and an administrator’s laptop. This platform is ideally suited for controlled real-time trading experiments, but is designed for relatively small-scale, synchronous markets where participants are physically co-located. If experiments are to be scaled up, to run for much longer periods and to support large scale human participation, an alternative platform architecture is required. To this end, development began on ExPo—the *Exchange Portal*—in 2011. ExPo has a Web service architecture, with humans participating via interaction through a Web browser (see [50]). This enables users to connect to the exchange via the Internet, and participate remotely. Immediately, ExPo negates the requirement for specific hardware, and enables long-term and many-participant experimentation, with users able to leave and return to a market via individual account log-in. Currently, an updated version—ExPo2—is under development at UNNC, in collaboration with Paul Dempster. As with OpEx and ExPo, ExPo2 will be released open-source to encourage replication studies and engagement in the wider scientific community.

In [8] a detailed proposal for future research studies is presented. In particular, future work will concentrate on relaxing some experimental constraints; such as enabling agents to trade on their own account, independent of permit schedules. This relaxation—effectively changing the function of agents from an agency trader (or “broker”) design, to a proprietary “prop” trader design—should enable the emergence of more realistic dynamics, such as Johnson et al.’s UEE price swing fractures. If we are able to reproduce these dynamics in the lab, this will provide compelling evidence for the RPT. Further, market structures and regulatory mechanisms such as financial circuit breakers, intentional network delays, and periodic (rather than real-time) order matching at the exchange, will be tested to understand the impact these have on market dynamics. In addition, preliminary studies to monitor human emotional responses to market shocks, using EEG brain data, are underway. Hopefully these studies can help us better understand how emotional reactions can exacerbate market swings, and how regulatory mechanisms, or trading interface designs, can be used to dampen such adverse dynamics.

Acknowledgements The experimental research presented in this chapter was conducted in 2011–2012 at the University of Bristol, UK, in collaboration with colleagues Marco De Luca (the developer of OpEx) and Charlotte Szostek. They both deserve a special thanks. Thanks also to all the undergraduate students and summer interns (now graduated) that helped support related work, in particular Steve Stotter and Tomas Gražys for work on the original ExPo platform. Finally, thanks to Paul Dempster and the summer interns at UNNC for work on developing the ExPo2 platform, and the pilot studies run during July 2016. Financial support for the studies at Bristol was provided by EPSRC grants EP/H042644/1 and EP/F001096/1, and funding from the UK Government Office for Science (Go-Science) Foresight Project on *The Future of Computer Trading in Financial Markets*. Financial support for ExPo2 development at UNNC was provided by FoSE summer internship funding.

References

1. Angel, J., Harris, L., Spratt, C.: Equity trading in the 21st century. Working Paper FBE-09-10, Marshall School of Business, Univ. Southern California, CA (February 2010). Available via SSRN: <https://ssrn.com/abstract=1584026> Accessed 22 Mar 2017
2. Arthur, W.B.: Complexity and the Economy. Oxford University Press (2014)
3. Battiston, S., Farmer, J.D., Flache, A., Garlaschelli, D., Haldane, A.G., Heesterbeek, H., Hommes, C., Jaeger, C., May, R., Scheffer, M.: Complexity theory and financial regulation: Economic policy needs interdisciplinary network analysis and behavioral modeling. *Science* **351**(6275) (2016) 818–819
4. Baxter, G., Cartlidge, J.: Flying by the seat of their pants: What can high frequency trading learn from aviation? In Brat, G., Garcia, E., Moccia, A., Palanque, P., Pasquini, A., Saez, F.J., Winckler, M., eds.: Proc. 3rd Int. Conf. Appl. and Theory of Autom. in Command and Control Syst. (ATACCS), Naples, Italy, ACM: IRIT Press (May 2013) 64–73
5. Berger, S., ed.: The Foundations of Non-Equilibrium Economics. Routledge, NY (2009)
6. Bisiyas, D., Flood, M., Lo, A.W., Valavanis, S.: A survey of systemic risk analytics. *Annu. Rev. of Financial Econ.* **4** (October 2012) 255–296
7. Bouchaud, J.P.: Economics needs a scientific revolution. *Nature* **455**(7217) (October 2008) 1181
8. Cartlidge, J.: Towards adaptive ex ante circuit breakers in financial markets using human-algorithmic market studies. In: Proc 28th Int. Conf. Artif. Intell. (ICAI), Las Vegas, NV, CSREA Press (July 2016) 77–80
9. Cartlidge, J., Cliff, D.: Exploring the ‘robot phase transition’ in experimental human-algorithmic markets. In: Future of Computer Trading. Government Office for Science, London, UK (October 2012) DR25. Available via GOV.UK: <https://www.gov.uk/government/publications/computer-trading-robot-phase-transition-in-experimental-human-algorithmic-markets> Accessed 22 Mar 2017
10. Cartlidge, J., Cliff, D.: Evidencing the robot phase transition in human-agent experimental financial markets. In Filipe, J., Fred, A., eds.: Proc. 5th Int. Conf. on Agents and Artif. Intell. (ICAART). Volume 1., Barcelona, Spain, SciTePress (February 2013) 345–352
11. Cartlidge, J., De Luca, M., Szostek, C., Cliff, D.: Too fast too furious: Faster financial-market trading agents can give less efficient markets. In Filipe, J., Fred, A., eds.: Proc. 4th Int. Conf. Agents and Artif. Intell. (ICAART). Volume 2., Vilamoura, Portugal, SciTePress (February 2012) 126–135
12. Chen, S.H., Du, Y.R.: Granularity in economic decision making: An interdisciplinary review. In Pedrycz, W., Chen, S.M., eds.: Granular Computing and Decision-Making: Interactive and Iterative Approaches. Springer (2015) 47–72
13. Cliff, D., Bruten, J.: Minimal-intelligence agents for bargaining behaviours in market-based environments. Tech. Rep. HPL-97-91, Hewlett-Packard Labs., Bristol, UK (August 1997)
14. Cliff, D., Northrop, L.: The global financial markets: an ultra-large-scale systems perspective. In: Future of Computer Trading. Government Office for Science, London, UK (September 2011) DR4. Available via GOV.UK: <https://www.gov.uk/government/publications/computer-trading-global-financial-markets> Accessed 22 Mar 2017
15. Das, R., Hanson, J., Kephart, J., Tesauero, G.: Agent-human interactions in the continuous double auction. In Nebel, B., ed.: Proc. 17th Int. Jt. Conf. Artif. Intell. (IJCAI), Seattle, WA, Morgan Kaufmann (August 2001) 1169–1176
16. De Luca, M.: Why robots failed: Demonstrating the superiority of multiple-order trading agents in experimental human-agent financial markets. In Loiseau, S., Filipe, J., Duval, B., van den Herik, J., eds.: Proc. 7th Int. Conf. on Agents and Artif. Intell. (ICAART). Volume 1., Lisbon, Portugal, SciTePress (January 2015) 44–53
17. De Luca, M., Cliff, D.: Agent-human interactions in the continuous double auction, redux: Using the OpEx lab-in-a-box to explore ZIP and GDX. In Filipe, J., Fred, A., eds.: Proc 3rd Int. Conf. Agents and Artif. Intell. (ICAART). Volume 2., SciTePress (January 2011) 351–358

18. De Luca, M., Cliff, D.: Human-agent auction interactions: Adaptive-aggressive agents dominate. In Walsh, T., ed.: Proc. 22nd Int. Jt. Conf. Artif. Intell. (IJCAI), AAAI Press (July 2011) 178–185
19. De Luca, M., Szostek, C., Cartledge, J., Cliff, D.: Studies of interactions between human traders and algorithmic trading systems. In: Future of Computer Trading. Government Office for Science, London, UK (September 2011) DR13. Available via GOV.UK: <https://www.gov.uk/government/publications/computer-trading-interactions-between-human-traders-and-algorithmic-trading-systems> Accessed 22 Mar 2017
20. Easley, D., Kleinberg, J.: Networks, Crowds, and Markets: Reasoning about a Highly Connected World. Cambridge University Press, NY (July 2010)
21. Easley, D., Lopez de Prado, M., O'Hara, M.: The microstructure of the 'flash crash': Flow toxicity, liquidity crashes and the probability of informed trading. *J. of Portfolio Management* **37**(2) (Winter 2011) 118–128
22. Farmer, J.D., Foley, D.: The economy needs agent-based modelling. *Nature* **460**(7256) (August 2009) 685–686
23. Farmer, J.D., Skouras, S.: An ecological perspective on the future of computer trading. In: Future of Computer Trading. Government Office for Science, London, UK (September 2011) DR6. Available via GOV.UK: <https://www.gov.uk/government/publications/computer-trading-an-ecological-perspective> Accessed 22 Mar 2017
24. Fel'tovich, N.: Nonparametric tests of differences in medians: Comparison of the Wilcoxon-Mann-Whitney and Robust Rank-Order tests. *Exp. Econ.* **6** (2003) 273–297
25. Foresight: The future of computer trading in financial markets. Final project report, The Government Office for Science, London, UK (October 2012). Available via GOV.UK: <http://www.cftc.gov/ido/groups/public/@aboutcftc/documents/file/tacfuturecomputertrading1012.pdf> Accessed 22 Mar 2017
26. Giles, J.: Stock trading 'fractures' may warn of next crash. *New Scientist* (2852) (February 2012). Available Online: <http://www.newscientist.com/article/mg21328525.700-stock-trading-fractures-may-warn-of-next-crash.html> Accessed 22 Mar 2017
27. Gjerstad, S., Dickhaut, J.: Price formation in double auctions. *Games and Econ. Behav.* **22**(1) (1998) 1–29
28. Gode, D., Sunder, S.: Allocative efficiency of markets with zero-intelligence traders: Markets as a partial substitute for individual rationality. *J. Political Econ.* **101**(1) (1993) 119–137
29. Gomber, P., Arndt, B., Lutat, M., Uhle, T.: High frequency trading. Technical report, Goethe Universität, Frankfurt Am Main, Germany (2011) Commissioned by Deutsche Börse Group.
30. Grossklags, J., Schmidt, C.: Artificial software agents on thin double auction markets: a human trader experiment. In Liu, J., Faltings, B., Zhong, N., Lu, R., Nishida, T., eds.: Proc. IEEE/WIC Conf. Intell. Agent Technol. (IAT), Halifax, Canada, IEEE Press (October 2003) 400–407
31. Grossklags, J., Schmidt, C.: Software agents and market (in)efficiency: a human trader experiment. *IEEE Trans. Syst., Man and Cybern., Part C (Appl. and Rev.)* **36**(1) (2006) 56–67
32. Holt, C.A., Roth, A.E.: The Nash equilibrium: A perspective. *Proc. Nat. Acad. Sci. USA (PNAS)* **101**(12) (2004) 3999–4002
33. Huber, J., Shubik, M., Sunder, S.: Three minimal market institutions with human and algorithmic agents: Theory and experimental evidence. *Games and Econ. Behav.* **70**(2) (2010) 403–424
34. Johnson, N.: To slow or not? Challenges in subsecond networks. *Science* **355**(6327) (2017) 801–802
35. Johnson, N., Zhao, G., Hunsader, E., Meng, J., Ravindar, A., Carran, S., Tivnan, B.: Financial black swans driven by ultrafast machine ecology. Working paper published on arXiv repository (feb 2012)
36. Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J., Tivnan, B.: Abrupt rise of new machine ecology beyond human response time. *Scientific Reports* **3**(2627) (September 2013) 1–7
37. Joint CFTC-SEC Advisory Committee on Emerging Regulatory Issues: Findings regarding the market events of May 6, 2010. Report, CFTC-SEC, Washington, DC (September 2010)

- 2010). Available via SEC: <https://www.sec.gov/news/studies/2010/marketevents-report.pdf> Accessed 22 Mar 2017
38. Kahneman, D.: *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, NY (2011)
 39. Keim, B.: Nanosecond trading could make markets go haywire. *Wired* (February 2012). Available Online: <http://www.wired.com/wiredscience/2012/02/high-speed-trading> Accessed 22 Mar 2017
 40. Leinweber, D.: *Nerds on Wall Street*. John Wiley (2009)
 41. May, R.M., Levin, S.A., Sugihara, G.: Complex systems: Ecology for bankers. *Nature* **451** (February 2008) 893–895
 42. Nelson, R.H.: *Economics as religion: from Samuelson to Chicago and beyond*. Penn State Univ. Press, University Park, PA (2001)
 43. Nelson, R.R., Winter, S.G.: *An evolutionary theory of economic change*. Harvard University Press (1982)
 44. Perez, E.: *The Speed Traders*. McGraw-Hill (2011)
 45. Price, M.: New reports highlight HFT research divide. *Financial News* (February 2012). Available Online: <https://www.fnlondon.com/articles/hft-reports-highlight-research-divide-cornell-20120221> Accessed 22 Mar 2017
 46. Schweitzer, F., Fagiolo, G., Sornette, D., Vega-Redondo, F., White, D.R.: Economic networks: what do we know and what do we need to know? *Adv. in Complex Syst.* **12**(04n05) (2009) 407–422
 47. Smith, V.: An experimental study of comparative market behavior. *J. Political Econ.* **70** (1962) 111–137
 48. Smith, V.: *Papers in Experimental Economics*. Cambridge University Press, Cambridge, UK (2006)
 49. Stotter, S., Cartlidge, J., Cliff, D.: Exploring assignment-adaptive (ASAD) trading agents in financial market experiments. In Filipe, J., Fred, A., eds.: *Proc. 5th Int. Conf. Agents and Artif. Intell. (ICAART)*. Volume 1., Barcelona, Spain, SciTePress (February 2013) 77–88
 50. Stotter, S., Cartlidge, J., Cliff, D.: Behavioural investigations of financial trading agents using Exchange Portal (ExPo). In Nguyen, N.T., Kowalczyk, R., Fred, A., Joaquim, F., eds.: *Transactions on Computational Collective Intelligence XVII. Volume 8790 of Lecture Notes in Computer Science (LNCS)*. Springer, Berlin, Heidelberg (November 2014) 22–45
 51. Tesauero, G., Bredin, J.: Strategic sequential bidding in auctions using dynamic programming. In Castelfranchi, C., Johnson, W.L., eds.: *Proc. 1st Int. Jt. Conf. Auton. Agents and Multiagent Syst. (AAMAS)*, Bologna, Italy, ACM (July 2002) 591–598
 52. Tesauero, G., Das, R.: High-performance bidding agents for the continuous double auction. In: *Proc. ACM Conf. Electron. Commer. (EC)*, Tampa, FL (October 2001) 206–209
 53. Treanor, J.: Pound’s flash crash ‘was amplified by inexperienced traders’. *The Guardian* (January 2017). Available Online: <https://www.theguardian.com/business/2017/jan/13/pound-flash-crash-traders-sterling-dollar> Accessed 22 Mar 2017
 54. Vytelingum, P.: *The Structure and Behaviour of the Continuous Double Auction*. PhD thesis, Sch. Electron. and Comput. Sci., Univ. Southampton, UK (2006)
 55. Vytelingum, P., Cliff, D., Jennings, N.: Strategic bidding in continuous double auctions. *Artif. Intell.* **172** (2008) 1700–1729
 56. Widrow, B., Hoff, M.E.: Adaptive switching circuits. *Institute of Radio Engineers, Western Electronic Show and Convention, Convention Record, Part 4* (1960) 96–104
 57. Zhang, S.S.: *High Frequency Trading in Financial Markets*. PhD thesis, Karlsruher Institut für Technologie (KIT), Germany (2013)