

Human-Agent Auction Interactions: Adaptive-Aggressive Agents Dominate

Marco De Luca and Dave Cliff*

Department of Computer Science, University of Bristol
Bristol, BS8 1UB, UK
{deluca, dc}@cs.bris.ac.uk

Abstract

We report on results from experiments where human traders interact with software-agent traders in a real-time asynchronous continuous double auction (CDA) experimental economics system. Our experiments are inspired by the seminal work reported by IBM at IJCAI 2001 [Das *et al.*, 2001], where it was demonstrated that software-agent traders could consistently outperform human traders in real-time CDA markets. IBM tested two trading-agent strategies, ZIP and a modified version of GD, and in a subsequent paper they reported on a new strategy called GDX that was demonstrated to outperform GD and ZIP in agent vs. agent CDA competitions, on which basis it was claimed that GDX “...may offer the best performance of any published CDA bidding strategy.” [Tesauro and Bredin, 2002]. In this paper, we employ experiment methods similar to those pioneered by IBM to test the performance of “Adaptive Aggressive” (AA) algorithmic traders [Vytelingum, 2006]. The results presented here confirm Vytelingum’s claim that AA outperforms ZIP, GD, and GDX in agent vs. agent experiments. We then present the first results from testing AA against human traders in human vs. agent CDA experiments, and demonstrate that AA’s performance against human traders is superior to that of ZIP, GD, and GDX. We therefore claim that, on the basis of the available evidence, AA may offer the best performance of any published bidding strategy.

1 Introduction

At the 2001 International Joint Conference on Artificial Intelligence (IJCAI), a team of researchers from IBM reported re-

sults from a series of experiments where human traders were pitted against software-agent traders in an experimental version of the continuous double auction (CDA) market mechanism, the mechanism that underlies electronic trading in most of the world’s financial markets. In this context, the CDA is a set of rules in which traders (buyers and sellers) may asynchronously post quotes (bids and offers) to an electronic system that provides a real-time display back to all traders, showing an indication of the current outstanding bids and offers. In some cases the traders may see only the current best bid and offer; in other cases they may see more “depth” in the market, showing price-ordered lists of the top n , or perhaps all, of the outstanding quotes. Exactly this mechanism is familiar to traders in the global financial markets for equities, commodities, currencies, and derivatives. The stunning result published by IBM [Das *et al.*, 2001], which made news headlines around the world, was that they found that two types of software agents (employing strategies known as ZIP and GD) consistently and robustly outperformed the human traders. The IBM team concluded their paper with:

“...the successful demonstration of machine superiority in the CDA and other common auctions could have ... direct and powerful impact – one that might be measured in billions of dollars annually.”

Almost a decade has passed since then, and yet, to the best of our knowledge, the IBM experiment has never been replicated¹. We set up a similar facility to that used by IBM by spending only a few thousand dollars on cheap “netbook” PCs which can act as the trader-terminals for human subjects in the experiments, and a mid-spec home server that runs the software trading agents and acts as the “exchange”. Thus, in this paper, we describe results that replicate IBM’s 2001 human vs. agent experiments on a network of netbooks. Also, we present the first results from testing a strategy called GDX

*The authors are very thankful to the UK Engineering and Physical Sciences Research Council (EPSRC: <http://www.epsrc.ac.uk>) for funding the equipment used for this research, and to the many students and staff at University of Bristol and Hewlett Packard Labs in Bristol who took part in the experiments. Both authors are members of, and receive financial support from, the UK Large-Scale Complex IT Systems (LSCITS) Initiative. Further details at <http://www.lscits.org>.

¹The most noteworthy study we found is [Grossklags and Schmidt, 2006], where human subjects and software agents trade on a simulated futures market. Although some of the experimental conditions described in their work are similar to ours & IBM’s, the majority of them are substantially different. To name a few: traders act both as buyers and sellers; orders for multiple units are allowed; limit prices change every round; no spread improvement rule is applied; both private and public information are distributed; and the trading agents employed are simple strategies. Their results are thus hardly comparable with those of IBM.

against human traders: IBM reported on GDX in [Tesauro and Bredin, 2002], and demonstrated that GDX outperforms both GD and ZIP in agent vs. agent CDA competitions. We then go on compare the performance of a new, fourth, agent strategy designed by Vytelingum [Vytelingum, 2006], called “Adaptive Aggressive” or AA. We find that AA outperforms ZIP, GD, and GDX in both human vs. agent and agent vs. agent contexts, and we therefore claim that AA may offer the best performance of any published bidding strategy.

2 Background

In his Nobel-prize-winning work [Smith, 1962], Vernon Smith ran several experiments with human traders to study the dynamics of CDA-based markets. In his experiments, Smith assigned one unit to sell(buy) at no less(more) than a specific price to each of the traders. The price of the unit, known as a *limit price* for a buyer, or a *cost price* for a seller, represents the maximum amount of money l a buyer can spend to buy the unit, or the minimum value c for which a seller can sell the unit. As a consequence, buyers make a profit $l - p$ if they buy at a price p that is less than their limit price, whereas sellers make a profit $p - c$ if they sell for a price p higher than their limit price. The limit prices are private, each trader knowing only her limit. The traders interact by quoting the price at which they are willing to trade their units. In Smith’s early experiments this happened by speaking the number out loud, thus the public quotes in a CDA are often referred to as *shouts*. A random player is selected every turn to make a shout, and the game finishes after a fixed number of turns. Following the rules of the CDA, a trade occurs when the outstanding bid is greater than or equal to the outstanding ask. Smith measured the performance of a trader in terms of *allocative efficiency*, which is the total profit earned by the trader divided by the *maximum theoretical profit* of that trader, expressed as a percentage. The maximum theoretical profit of a trader is the profit that trader could have made if all the market participants would have traded their units at the theoretical competitive market *equilibrium price* p^* .

Formally, let I be the set of buyers and J the set of sellers in the market. Let $L_i = \{l_{i,1}, l_{i,2}, \dots, l_{i,N_i}\}$ be the set of limit prices of the units owned by buyer i , and $C_i = \{c_{j,1}, c_{j,2}, \dots, c_{j,M_j}\}$ the set of cost prices of the units owned by seller j . The market equilibrium price is given by:

$$p^* = \arg \max_p \left\{ \sum_{i \in I} \sum_{n=1}^{N_i} \max(0, l_{i,n} - p) + \sum_{j \in J} \sum_{m=1}^{M_j} \max(0, p - c_{j,m}) \right\} \quad (1)$$

The maximum theoretical profit Π_i^* of buyer b_i is given by:

$$\Pi_i^* = \sum_{n=1}^{N_i} \max(0, l_{i,n} - p^*) \quad (2)$$

Denoting with $p_{i,n}$ the price at which buyer i actually trades the unit with limit price $l_{i,n}$, the actual profit Π_i earned

by buyer i is $\Pi_i = \sum_{n=1}^{N_i} \max(0, l_{i,n} - p_{i,n})$. Therefore, the allocative efficiency E_i of buyer i is:

$$E_i = \frac{\Pi_i}{\Pi_i^*} \quad (3)$$

The calculations above also apply to sellers, replacing limit prices with cost prices. A further measure of the performance of a market is the *profit dispersion*: this is defined as the cross-sectional root mean squared difference between the actual profits and the maximum theoretical profits of individual traders. Using the notation introduced above, for a group of T traders the profit dispersion is given by:

$$\sqrt{\frac{1}{T} \sum_{k=1}^T (\Pi_i - \Pi_i^*)^2} \quad (4)$$

Smith demonstrated that markets governed by the CDA can reach close-to-optimal efficiency. Also, he proved that transaction prices converge to the market’s theoretical competitive equilibrium price. Furthermore, he found that if the supply and demand of markets suddenly changed, the transaction prices would rapidly converge to the new equilibrium price. [Gode and Sunder, 1993] introduced “zero intelligence” (ZI) automated trading agents that submit random bids and offers, and found that they exhibit similar (although slightly lower on average) levels of allocative efficiency to those achieved by humans in Smith’s experiments: indeed, they concluded that allocative efficiency is almost entirely a product of market structure, rather than an indication of the traders’ negotiation skills.

This was proven to be wrong by [Cliff and Bruten, 1997], who demonstrated that ZI agents converged to equilibrium price only in certain special cases, and introduced a slightly more intelligent trading agent that they named Zero Intelligence Plus (ZIP). ZIP agents, which employed simple adaptive mechanisms, showed convergence independently from the supply and demand curves, and generated better results than ZI agents in terms of both allocative efficiency and profit dispersion. Several studies on automated trading agents followed: more details can be found in the sections below.

2.1 ZIP

In 1996, Cliff created the Zero-Intelligence Plus (ZIP) algorithm to investigate the minimum level of intelligence required to achieve convergence to market equilibrium price [Cliff and Bruten, 1997]. ZIP has been used in several subsequent studies ([Das *et al.*, 2001], [Tesauro and Das, 2001], [Tesauro and Bredin, 2002], [Vytelingum *et al.*, 2008]) as a benchmark for evaluation of strategy efficiency.

Each ZIP trader agent maintains a real-valued *profit margin* and employs simple heuristic mechanisms to adjust their margin using market data. In this context, the profit margin represents the difference between the agent’s limit price and the shout price, that is the price that the agent sends to the market to buy or sell the commodity. By observing market events, ZIP buyers increase their profit margin, and therefore make cheaper bids, when a trade at a lower price than their

current shout price occurs. Conversely, ZIP buyers that observe an accepted offer at a price higher than the one they have put on the market move towards that price by lowering their profit margin (that is, bidding a higher price). The same applies to buyers that witness a rejected bid at a higher price than the one they are advertising. Symmetrically similar rules are followed by ZIP sellers.

The profit-margin adaptation rule used in the ZIP algorithm to dynamically respond to the market conditions is based on the Widrow-Hoff “delta rule” with an additional noise-smoothing “momentum” term. The profit margin of the ZIP traders is adjusted by a small random quantity, proportional to the learning rate of the individual agent.

2.2 GD/MGD/GDX

At about the same time as Cliff and Bruten, but working independently of them, Gjerstad and Dickhaut [Gjerstad and Dickhaut, 1998] introduced a bidding algorithm, which we shall refer to as GD, centred on “belief” functions that agents form on the basis of observed market data. GD agents collect the orders (rejected shouts) and trades (accepted shouts) occurred during the last M trades, and store them in a history H . When a GD agent prices an order, from the history H it builds the belief function $f(p)$, which represents the probability that an order at price p will result in a trade. For example, the belief function for a GD buyer is:

$$f(p) = \frac{TBL(p) + AL(p)}{TBL(p) + AL(p) + RBG(p)} \quad (5)$$

where $TBL(p) = \sum_{d \leq p} TB(d)$ represents the number of accepted bids found in H at price $\leq p$, $AL(p) = \sum_{d \leq p} A(d)$ is the number of asks in H with price $\leq p$, and $RBG(p) = \sum_{d \geq p} RB(d)$ is the number of rejected bids in H at price $\geq p$. Here, $TB(d)$, $A(d)$ and $RB(d)$ are the taken bids, the offers, and the rejected bids at price d respectively. Note that $f(p)$ depends on H , and therefore it can potentially change every time a market participant sends an order to the market. Because $f(p)$ is defined only for some values of p , the function is interpolated to provide values over the domain of all the valid prices. Finally, the price p that maximises the product of the interpolated $f(p)$ and the profit function of the agent (equal to $l - p$ for buyers and $p - l$ for sellers) is chosen as the order price.

The original GD algorithm was modified by Das, Hanson, Kephart and Tesauro to fit in MAGENTA, the real-time asynchronous framework they described in [Das *et al.*, 2001]. Unlike the CDA mechanism Gjerstad and Dickhaut adopted to develop GD, the CDA implemented in MAGENTA allows persistent orders, and therefore the concept of “rejected order” becomes unclear. This problem was addressed by not adding unmatched orders to the history H as soon they are entered, but only after a time equal to a “grace period” τ_p has expired. Also, the parameter M was increased to a much larger value and the simple count terms in Equation 1 were replaced by exponentially weighted sums that emphasise the most recent terms and dim the old ones. Finally, the modified GD agent could handle multiple units to trade sequentially

during the auction period, by maintaining their limit prices in a vector.

Different changes to GD are described by Tesauro and Das in their work [Tesauro and Das, 2001]. Their version of GD, which they named “MGD”, maintains the highest and lowest prices of the last trading period in order to them as constraints for the belief function: MGD buyers (sellers) assume that the probability that an order is accepted at a price greater than the highest price in the last period is 1 (0), and that the probability that an order is accepted at a price less than the lowest price in the last period is 0 (1). These constraints are added to the belief function after interpolation, together with the constraints deriving from the current bid-ask spread. MGD agents can also deal with multiple tradable units, and are allowed to trade the least valuable unit if there is more than one unit available.

Yet another set of changes were made to GD by Tesauro and Bredin in [Tesauro and Bredin, 2002], that resulted in their GDX algorithm. GDX substantially differs from GD in that it makes use of Dynamic Programming (DP) to price orders. The pricing function takes into account both the effect of trading the current unit immediately, and the effect of trading it in the future, discounting the latter by a parameter γ . As a result, GDX agents do not just maximise the immediate profit, but instead optimise the pricing process in order to achieve overall higher returns over the entire trading period:

$$V(m, n) = \max_p \left\{ \overbrace{f(p)s_m(p)}^{\text{...if I trade now...}} + \underbrace{\gamma f(p)V(m-1, n-1)}_{\text{money I will make later}} + \underbrace{\gamma[1-f(p)]V(m, n-1)}_{\text{same units, 1 less time unit}} \right\} \quad (6)$$

2.3 AA

The Adaptive Aggressiveness (AA) trading algorithm was developed by Vytelingum for his 2007 PhD thesis [Vytelingum, 2006], and it is the most recent automated trading strategy explored here. In Vytelingum’s work, aggressiveness represents the agent’s tendency to trade off profit in favour of higher chances of transacting: the more (less) aggressive the agent, the better (worse) the offers it submits are than what the competitive equilibrium price the agent believes to be. Similarly to Cliff’s ZIP automated traders, AA agents also monitor market signals and adjust internal parameters using a learning mechanism. The innovation introduced with the latter consists of updating two individual components, a short-term and a long-term one. The short-term learning mechanism updates the aggressiveness of the agent on the basis of the market data observed, in order to react promptly to the market fluctuations. The long-term learning process, on the other hand, captures the slower market trends that develop through the time so that agents can take those into account when making their bidding decisions.

The heuristics employed as learning rules are analogous to those used in ZIP, except they control the aggressiveness of an agent instead of regulating its profit margin. Each time

a trade occurs, an AA agent adjusts its aggressiveness according to the transaction price and its own current target price. A buyer will become more aggressive (and therefore it will shout higher prices) if the transaction price is higher than its current target price, whereas it will become less aggressive (by submitting cheaper bids) if the transaction price is lower than its target price. Equivalently, a seller will become less aggressive (and therefore it will shout higher prices) if the transaction price is higher than its current target price, whereas it will become more aggressive (by submitting cheaper offers) if the transaction price is lower than its target price. An AA buyer (seller) also increases its aggressiveness when a bid (an offer) is rejected at a price higher (lower) than its current target price. The aggressiveness is updated according to the Widrow-Hoff rule, that is backprojecting a fraction of the error between the desired and the current degree onto the new degree of aggressiveness.

3 Experimental Setup

We ran our experiments on Open Exchange (OpEx), an experimental algorithmic trading platform developed by De Luca [De Luca and Cliff, 2011]. OpEx was designed to resemble closely the structure and the behaviour of modern commercial financial-market electronic trading systems, and to be generic enough to support experimental economics simulations of arbitrary complexity. All of our human vs. robot experiments involved 6 human traders and 6 robot traders, both equally split into 3 buyers and 3 sellers, the structure described by IBM in their original work [Das *et al.*, 2001]. Before each experiment, the human subjects were briefed about the rules, and were given some time to familiarise with the Sales Trading GUI (briefing and tutorial typically took less than 30 minutes). Figure 1 shows a screenshot of the OpEx Sales Trading GUI.

The subjects had no previous professional experience in electronic trading, and they were only allowed to participate in one experiment. We motivated all 6 participants by giving each of them a token worth £20, plus a bonus of £40 and £20 to the first and the second best trader, respectively. An experiment consisted of 10 consecutive “rounds” 3 minutes long. At the beginning of a round, each of the 12 players received a fresh supply of 13 units to buy or to sell during that round, according to their role. At the end of the round the unused units were discarded, without any profit or loss for the traders. Players had to trade their units sequentially, and the sequence of their limit prices was arranged in an arithmetic progression. Only 3 “generator” sequences were actually used to produce the prices for all the players: a human and his/her robot counterparty had the same limit prices; and buyers and sellers share the same values except for the order, that is increasing for sellers and decreasing for buyers. The progressions had the same slope, and they were chosen so that each player had approximately the same maximum theoretical surplus in a given round. In line with [Das *et al.*, 2001], we introduced market shocks, which periodically altered the limit prices adding or subtracting a constant to them every 2 rounds. Thus, a 30 minute simulation was constituted by 5 consecutive trading periods at different equilibrium prices.




Figure 1: OpEx Sales Trading GUI. In the top section, from left to right: the order panel, where orders are entered by the traders; the market orderbook; the info panel, displaying market open/close time, PNL, and remaining/executed orders. In the mid and bottom section, the order blotter and the trade blotter, showing the status of the trader’s orders and the details of her trades, respectively.

Table 1 presents 3 sample sequences that we used in the experiments described here.

Unit no. (B)	Unit no. (S)	Seq. 1	Seq. 2	Seq. 3
1	13	381	383	385
2	12	371	373	375
3	11	361	363	365
4	10	351	353	355
5	9	341	343	345
6	8	331	333	335
7	7	321	323	325
8	6	311	313	315
9	5	301	303	305
10	4	291	293	295
11	3	281	283	285
12	2	271	273	275
13	1	261	263	265

Table 1: Generator sequences used in our experiments. The limit prices of the units assigned to the first human buyer (HB1) are those of Sequence 1, in the order shown in the leftmost column. The first robot buyer (AB1) is assigned identical limit prices. The cost price of the units assigned to the first human seller (HS1) and the first robot seller (AS1) are also those of sequence 1, but in the order presented in the second column from the left. This logic applies seamlessly to the remaining players: HB2, AB2, HS2 and AB2 follow sequence 2, and HB3, AB3, HS3 and AB3 follow sequence 3.

4 Experimental Results

4.1 Agents vs. Humans

We ran a total of nine human vs. agent experiments, three for each of ZIP, GDX and AA. Table 2 presents the mean values

Experiment	Trades			Performance			Market	
	A-A	A-H	H-H	Eff(A)	Eff(H)	Δ Profit(A-H)	Eff	Profit Disp
AA	41%	32%	27%	1.078	0.867	27%	0.978	793
GDX	33%	42%	25%	1.040	0.866	23%	0.954	568
ZIP	39%	30%	31%	1.014	0.941	9%	0.978	418

Table 2: Summary of the nine human vs. agent experiments. For each strategy, the table displays: the strategy employed by all six agents; the percentage of trades made between two Agents, an Agent and a Human, and two Humans; the average efficiency of Agents and Humans; the percentage difference between Agent surplus and Human surplus; the market efficiency and the profit dispersion. The mean maximum theoretical profit per trader per simulation is 2107. Lower profit dispersion and higher mean efficiency values are better. All numerical values are mean values over three experiments.

of the results we obtained for all the strategies².

The first noteworthy finding is that the agents performed better than the humans as a group: the mean percentage difference between the profit made by the agents and the humans is 20% over the nine experiments. In addition to that, the allocative efficiency achieved by the agents is greater than 100% for every strategy, which proves that all of the strategies were successful in exploiting human errors. Second, the trades between agents and humans were, on average, 35% of the trades made in an experiment: despite the fact that the automated traders were faster, our subjects were very well integrated in the market.

Moreover, we analysed the timing of trades and found that they are more likely to happen in the first minute of the trading period. Because of the distribution of the limit prices we employed, buyers have higher values and sellers have lower costs for their first units, resulting in a wider spectrum of acceptable prices by buyers and sellers. Such price spectrum narrows as time passes, as a consequence of both more demanding limit prices and the New York Stock Exchange (NYSE) spread-improvement rule³, leading to barely any trading activity towards the end of the trading period. We narrowed down our timing analysis to distinguish, for each strategy, among the three groups of trades: between agents (A-A); between humans (H-H); and between humans and agents (A-H).

Figure 2 presents a chart of the probability distribution function of the trade time of the three groups of trades for every strategy. The decreasing trend just described is displayed more evidently by the A-A series, confirming that the agents were faster than our subjects in taking advantage of the early very profitable trading opportunities. The shape of the A-H series is similar although smoother. The trading activity between humans, on the other hand, is distributed more evenly over the time and generally exhibits more significant values during the first minute.

Furthermore, analysing the rankings of the players' efficiency, we discovered that the best 6 players were either mostly buyers or mostly sellers, consistently throughout the 9 experiments. In more detail: in 4 experiments the best 6




Figure 2: Empirical probability distribution function of the time of the trades between two agents (A-A), between two humans (H-H) and between one human and one agent (A-H) for the three strategies. Data are aggregated in 5 seconds bins.

players were a homogeneous group of either buyers or sellers; in 4 experiments, 5 out of the best 6 players were on the same side; and in the remaining one 4 sellers were in the best 6. Interestingly, this pattern was not found in the numerous robot vs. robot experiments we ran, nor is there mention of it in previous agent vs. agent work related to this. Thus, in line with [Das *et al.*, 2001], we speculate that this asymmetry is due to the heterogeneous nature of our market: the agents follow a rigorous algorithm to price their offers, while the human players form a belief about the market that includes an irrational psychological component. The data from the experiments that favoured buyers show that the human sellers, observing low (although increasing) trade prices at the start of the trading period, offered cheaper and cheaper deals to their counterpart until they matched their bids. This behaviour was followed by the robots, which sensed that trades happened at low prices and adjusted accordingly. Although unable to make passive (i.e. more expensive than the best price) of-

²To 6 decimals, AA-human=0.977797, ZIP-human=0.977814.

³The NYSE spread-improvement rule is typically found in CDA studies such as those cited in this paper, and requires that new bids (or asks) must be higher (or lower) than the current best bid (or ask) in the market.




Figure 3: Trade price time series for a human vs. AA experiment. The vertical lines represent the start of a new round. The 10 rounds of 3 minutes each were divided into 5 phases, each of which with its own set of limit prices. The theoretical equilibrium price p^* for each phase is indicated by the horizontal dashed lines. Trades between two humans are marked with open squares, between two agents with open circles, and between an agent and a human with solid circles. Mean efficiency per phase (vertical bars) and per rounds are shown for Agent Buyers (AB), Agent Sellers (AS), Human Buyers (HB) and Human Sellers (HS).

fers because of the NYSE rule, the humans could nevertheless have waited until the market settled on a price, and then they could have started offering: this way, the agent sellers would have crossed the narrow spread to beat the human prices and they would have soon exhausted their intramarginal units, giving the human sellers control over the price. We interviewed one of our subjects that we noticed was playing this strategy during an experiment that displayed similar dynamics to those just described, and he confirmed that he was actually following that strategy. However, because he was the only human seller to do so, the tactic turned in his disfavour: every other human (and agent) seller kept underselling while he was waiting, thus when he started trading most of the intramarginal buyers' units had been already traded, and he could only make few underpriced deals.

AA Agents vs. Humans (AA-v-H)

The trade price time series of a typical AA-v-H experiment is shown in Figure 3. We will refer to this specific experiment, although the observations we made on the other AA-v-H experiments are very similar. The dashed vertical lines separate the trading periods, whereas the dashed horizontal lines mark the theoretical equilibrium price p^* . The shape of the time series indicates robust and recurrent convergence to p^* . Every trading period begins with a fast trading phase where the market price settles, thanks to both the NYSE spread improvement rule and the particular sequence of limit prices we

employed. During this phase, the most profitable units are consumed while the spread between intramarginal sellers and buyers decreases. As a consequence, the amplitude of the oscillations drops, and prices move neatly towards p^* . As soon as a new period starts, because new profitable units are available, the buyers start over bidding low prices, and so forth.

Also, it is worth noticing that the efficiency achieved by the buyers in each trading period is consistently higher than that of sellers, as evidence of the unbalance introduced by the human component.

GDX Agents vs. Humans (GDX-v-H)

Four trading periods of the trade price time series of the GDX-v-H experiment ran on 3rd September 2010 are shown in Figure 4. We will refer to this specific sample, as it is quite representative of the data we gathered from the three GDX-v-H experiments. The time series exhibits a recurring pattern of convergence towards a price that is often somewhat lower than p^* . Most of the trades were made at lower prices than p^* , since buyers closed deals at reasonably lower prices than their limit prices, and therefore kept higher profit margins than their seller counterparts. This is confirmed by the fact that, in this experiment, the five best traders in terms of mean allocative efficiency are buyers: this is in line with the above mentioned asymmetry we detected throughout the experiments.

A more detailed analysis of the efficiency per trading pe-




Figure 4: Four periods of a human vs. GDX experiment, compared to four periods of a human vs. ZIP experiment.

riod reveals that the discrepancy between buyers and sellers is accentuated by the raising of the equilibrium price (e.g. between trading periods 6 and 7), and attenuated by the drop (e.g. between trading periods 4 and 5). We explained this by looking at the first few trades made in the trading period following the market shock: their prices tend to remain close to the previous value of p^* , resulting in better opportunities for buyers or for sellers, if there was a raise or a drop of p^* respectively. This confirms that the GDX strategy requires a few samples before it can adapt to the new market condition.

Furthermore, the DP approach employed in GDX is revealed by the small tail of the distribution in Figure 2: when they wake up during the last few seconds, due to the iterative nature of the algorithm, the GDX agents estimate the probability of making a deal immediately to be higher than that of profiting from a postponed order. As a result, GDX degenerates into GD and the agents price the final units by maximising the product of profit and belief function, rather than by holding them in the hope of higher future returns.

ZIP Agents vs. Humans (ZIP-v-H)

Figure 4 also illustrates the first four trading periods of a typical ZIP-v-H experiment. By visual inspection, it can be verified that human-ZIP markets display excellent capabilities of tracking the equilibrium price, as convergence to p^* is more pronounced than in human-GDX markets.

It can be noted that, qualitatively, the shape of the time series is reasonably consistent across the trading periods, and that the curve presents a higher price excursion in a shorter time than GDX before converging to p^* . Indeed, our time analysis confirms that the mean trade time relative to the trading period is 45s for ZIP-human and 69s for GDX-human markets.

By isolating the trades between two agents (A-A), between two humans (H-H), and between a human and an agent (A-H), we found that the mean trade time of the three types of trades is consistently higher in GDX than in ZIP: this is qual-




Figure 5: Shout price series of one AA and one ZIP vs. human experiment.

itatively confirmed by Figure 2. Also, the mean trade time of A-A trades is the smallest and that of H-H trades is the largest consistently across trading periods in the experiments involving ZIP, while this relationship does not hold for some trading periods of the GDX-v-H experiments.

Although the trade price series of the AA-human and the ZIP-human markets look similar, AA agents managed to extract substantially more profit from the market than what ZIP agents did. To explain this, we analysed the shout price time series, that is the series of prices submitted by buyers and sellers over the time, including both accepted and rejected orders. Figure 5 represents one trading period of a ZIP-human market as compared to one trading period of an AA-human market. The chart outlines how some AA agents are reluctant to trade the first units during the initial “price settling” period, as they rather increase their aggressiveness gradually. As a consequence, a number of underpriced sales and overpriced purchases are made by the human players, and while this happens AA buyers keep their positions having bought most of the units during the initial phase, whereas AA sellers trade at the now higher price. Similar observations can be made for AA markets that favour buyers. On the other hand, the shout prices of the ZIP sample period are clustered quite closely around the trade price trajectory, with the exception of an initial exploration of decreasingly high prices by agent buyers. Thus, although ZIP-human markets exhibit a lower profit dispersion, the group efficiency for ZIP agents against humans is lower than for AA agents.

4.2 Agents vs. Agents

We ran further experiments to investigate the performance of AA, ZIP and GDX in pure agent vs. agent markets. The simulations were performed using OpEx Discrete Event Simulator (DES), and the supply and demand curves employed are identical to those we used for the human vs. robot markets. The results of five sets of experiments are presented in Table 3. The second column of the table displays the percentage

Agents	Experiments	Rounds
GDX ($\gamma = 0$) vs. ZIP	98%	9.015(± 0.123)
GDX ($\gamma = 0.9$) vs. ZIP	100%	9.001(± 0.029)
AA vs. GDX ($\gamma = 0.9$)	94%	9.941(± 0.235)
AA vs. GDX ($\gamma = 0$)	92%	9.924(± 0.265)
AA vs. ZIP	95%	9.046(± 0.209)

Table 3: Summary of five agent vs. agent experiments, each one repeated over 1000 times. For each set of experiments, the table shows: the agents involved (the winner first), the percentage of experiments won, and the mean number of rounds per experiment won.

of experiments in which the first group of agents made more profit than the second one. The third column shows the mean number of rounds won, again in terms of profit extracted. The data clearly indicate that each of the competitions resulted in the undisputed supremacy of one of the strategies.

Qualitatively in line with [Tesauro and Bredin, 2002], GDX visibly outperforms ZIP, both when run in optimal mode ($\gamma = 0.9$) and when degenerated to GD ($\gamma = 0$); in particular, the performance of GDX in terms of profit extracted improves slightly for $\gamma = 0.9$, although the mean number of rounds won by GD is mildly higher. Similar observations can be made on the AA-GDX markets. AA beat both GDX and GD by far, GDX being a slightly better opponent. Also, the profit extracted per round by AA is much higher than that by GD/GDX, as shown by the number of rounds won, which is very close to 10. AA defeated ZIP by far as well, although interestingly its victory was not as overwhelming as one may naively expect from the results just described for AA vs. GDX and GDX vs. ZIP. These findings are consistent with [Vytelingum *et al.*, 2008].

5 Conclusions

The competition between algorithmic trading strategies has been of interest to researchers for over a decade. Three main adaptive strategies have been studied and led to substantial literature: namely ZIP, the GD class, and AA.

ZIP had been shown to perform worse than GD, MGD and GDX in agent vs. agent markets, and also Das *et al.*, in their seminal human vs. agent study, showed that their variant of GD, and also ZIP, both outperformed human traders. More recently, Vytelingum designed AA and proved that its performance is better than both ZIP and the GD class in agent vs. agent contests only: he ran no tests against humans.

We have designed and implemented, for a total cost of a few thousand dollars, an experimental economics laboratory network trading system, where “trader terminal” netbooks communicate with a central “exchange” server, with the potential for multiple instruments to be traded simultaneously in varying quantities, and with every quote in the marketplace, and details of all transactions, written to a database as a single “consolidated tape” record of the trading events (to sub-second timestamp accuracy) over the course of a trading experiment. We employed this trading system, called *OpEx*, to investigate the behaviour of these strategies in agent vs. agent markets, and to pit human subjects against ZIP and, for the

first time ever, GDX and AA traders.

AA proved to outperform every other strategy to a great degree when competing against each other, and to perform sensibly better than the other ones in mixed agent-human markets. We therefore claim that AA may offer the best performance of any published strategy.

Finally, it would be interesting to test our algorithmic traders in an additional scenario: one where the “reaction times” of the agents are restricted to be comparable to those of the human traders. This would reveal in what measure the superiority of the agents is due to their speed. We intend to explore this, and other issues, in our future work.

References

- [Cliff and Bruten, 1997] D. Cliff and J. Bruten. Minimal intelligence agents for bargaining behaviors in market-based environments. *Technical report, HPL-97-91, HP Labs*, 1997.
- [Das *et al.*, 2001] R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In *Proc. IJCAI*, pages 1169–1176, August 2001.
- [De Luca and Cliff, 2011] M. De Luca and D. Cliff. Agent-human interactions in the continuous double auction, redux: Using the opex lab-in-a-box to explore zip and gdx. In *Proc. ICAART*, 2011.
- [Gjerstad and Dickhaut, 1998] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22(1):1–29, 1998.
- [Gode and Sunder, 1993] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *J. Polit. Econ.*, 101(1):119–137, 1993.
- [Grossklags and Schmidt, 2006] J. Grossklags and C. Schmidt. Software agents and market (in)efficiency: a human trader experiment. *IEEE Trans. on Systems, Man & Cybernetics*, 36(1):56–57, January 2006.
- [Smith, 1962] V. Smith. An experimental study of comparative market behavior. *J. Polit. Econ.*, 70:111–137, 1962.
- [Tesauro and Bredin, 2002] G. Tesauro and J. L. Bredin. Strategic sequential bidding in auctions using dynamic programming. In *Proc. AAMAS*, pages 591–598, 2002.
- [Tesauro and Das, 2001] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proc. Third ACM Conf. on e-Commerce*, pages 206–209, 2001.
- [Vytelingum *et al.*, 2008] P. Vytelingum, D. Cliff, and N. R. Jennings. Strategic bidding in continuous double auctions. *Artificial Intelligence*, 172:1700–1729, 2008.
- [Vytelingum, 2006] P. Vytelingum. *The Structure and Behaviour of the Continuous Double Auction*. PhD thesis, University of Southampton, Southampton, UK, 2006.