

Agent-Human Interactions in the Continuous Double Auction

Rajarshi Das, James E. Hanson, Jeffrey O. Kephart and Gerald Tesauro

Institute for Advanced Commerce

IBM T.J. Watson Research Center

30 Saw Mill River Road, Hawthorne, NY 10532, USA

Abstract

The Continuous Double Auction (CDA) is the dominant market institution for real-world trading of equities, commodities, derivatives, etc. We describe a series of laboratory experiments that, for the first time, allow human subjects to interact with software bidding agents in a CDA. Our bidding agents use strategies based on extensions of the Gjerstad-Dickhaut and Zero-Intelligence-Plus algorithms. We find that agents consistently obtain significantly larger gains from trade than their human counterparts. This was unexpected because both humans and agents have approached theoretically perfect efficiency in prior all-human or all-agent CDA experiments. Another unexpected finding is persistent far-from-equilibrium trading, in sharp contrast to the robust convergence observed in previous all-human or all-agent experiments. We consider possible explanations for our empirical findings, and speculate on the implications for future agent-human interactions in electronic markets.

1 Introduction

We envision a future in which economically intelligent and economically motivated software agents will play an essential role in electronic commerce. Among the present-day glimmerings of such a future are simple bidding agents offered by auction sites such as eBay and Amazon and by third-party bidding services such as eSnipe¹, pricebots such as buy.com that automatically undercut the competition, and shopbots such as DealTime that minimize the total cost of a bundle of goods by partitioning it across one or more vendors, taking shipping cost schedules into account. It is natural to expect continued growth in the variety and sophistication of automated economic decision-making technologies such as these. In addition, we anticipate the emergence of an even larger and more diverse class of agents for which economic decision-making capabilities are still essential, but ancillary to the primary function of providing information goods and services

¹eSnipe automates a common practice among eBay bidders of waiting until a few seconds before an auction's close to place a bid.

to humans or other agents [Kephart *et al.*, 2000]. (Throughout this paper we use the term “agent” to refer exclusively to a software agent, as opposed to a human economic agent.) Whether their main business is ontology translation, match-making, network service provision, or anything else, these agents will charge a fee for their goods or services, and will negotiate both as buyers and as sellers with other agents. Thus they will have to be economically intelligent, capable of making effective decisions about pricing, purchasing, or bidding.

If this vision is to be realized, then it must be demonstrated that, within their domain of application, agents can attain a level of economic performance that rivals or exceeds that of humans on average, without introducing undue risk. Otherwise, people would not entrust agents with making economic decisions.

The purpose of this paper is to provide such a demonstration. Through a series of controlled laboratory experiments in which humans and agents participate simultaneously in a realistic auction (a Continuous Double Auction, or CDA), we show that software agents can consistently obtain greater gains from trade than their human counterparts. In a sense, this work can be viewed as another chapter in the venerable AI tradition of human vs. machine challenges. Already, machine supremacy has been demonstrated in two-player games such as backgammon, checkers, and chess, and a serious attack is now being made on games such as bridge and poker [Schaeffer, 2000], in which there are slightly more than two players who play in a well-defined sequence. In contrast, the number of players in the CDA is typically much greater (we limit it to 12 in this report), and the moves by individual players are completely independent and asynchronous. These and other features make game-theoretic analysis of the CDA intractable. Another notable difference is that the successful demonstration of machine superiority in the CDA and other common auctions could have a much more direct and powerful financial impact—one that might be measured in billions of dollars annually.

This paper is organized as follows. Section 2 defines the CDA and discusses the relationship of our work to previous studies by economists and computer scientists. Section 3 provides a brief overview of the technological infrastructure used in the experiments. Section 4 describes the agent environment and the individual agent strategies. Section 5 provides details of the market rules and experimental parameters, and

section 6 presents the results for two different agent strategies. We conclude with a brief summary and discussion of implications and future directions in section 7.

2 Background on the CDA

Our laboratory study of economic interactions between agents and humans utilizes a simplified model of a Continuous Double Auction (CDA) market. The CDA is one of the most common exchange institutions, and is in fact the primary institution for trading of equities, commodities and derivatives in markets such as NASDAQ and the NYSE. In the CDA, there is a fixed-duration trading period, and buy orders (“bids”) and sell orders (“asks”) may be submitted at any point during the period. If at any time there are open bids and asks that are compatible in terms of price and quantity of good, a trade is executed immediately. Typically, an announcement is broadcast immediately to all participants when orders are placed or trades are executed.

In our model CDA, multiple units of a single hypothetical commodity can be bought or sold. Participants are assigned a fixed role of either Buyer (only submits bids) or Seller (only submits asks). There are several periods of trading; at the start of each period, participants are given a list of “limit prices” (values for Buyers and costs for Sellers) for each unit to be bought or sold. The limit prices are held constant for several periods and periodically shifted by random amounts to test responsiveness to changing market conditions. Each participant’s objective is to maximize “surplus,” defined as (limit price - trade price) for buyers and (trade price - limit price) for sellers.

The assumptions of fixed roles and fixed limit prices conform to extensive prior studies of the CDA, including experiments involving human subjects [Smith, 1962; 1982] and automated bidding agents [Cliff and Bruten, 1997; Gjerstad and Dickhaut, 1998]. Under such assumptions, a market consisting of rational players will eventually converge to steady trading at an equilibrium price p^* , at which there is a balance between Supply (the total number of units that can be sold for positive surplus) and Demand (the total that can be bought for positive surplus). For each participant, one can define a theoretical surplus as the total surplus that would be obtained if all units traded at p^* . One can also define a participant’s efficiency μ as the ratio of actual surplus to theoretical surplus. In human subject studies [Smith, 1962; 1982], convergence close to equilibrium was found within several periods, with the approach towards p^* exhibiting a “scallop” shape (i.e., a decelerating curved trajectory) of progressively smaller amplitude in each successive period. Robust convergence to equilibrium was also found in homogeneous populations of agents [Cliff and Bruten, 1997; Gjerstad and Dickhaut, 1998], with smaller-amplitude scallops than in the all-human experiments. Both all-human and all-agent CDA studies claimed very high population efficiency, ranging between 0.95 and 1.0.

Our work differs from prior CDA studies in two significant ways. First, we are interested in studying and understanding interactions between agent and human bidding strategies. Second, we focus primarily on measuring and understanding

the performance of individual agents, instead of global measures of aggregate market behavior. As agent designers, we would like to understand the principles by which robust bidding strategies can be designed that perform well against both human and computerized opposition. Our focus on competition in heterogeneous bidder populations is similar to that of the agent vs. agent competition held at the Santa Fe Double Auction Tournament (SFDAT) [Rust *et al.*, 1992]. The SFDAT was an intrinsically discrete-time auction, with non-persistent orders, synchronized bidding of all agents at every time step, and a coarse time step size deliberately chosen to allow all agents enough time to calculate and place their bids. Thus the conclusions may not apply to trading in normal real-time CDA markets. Another market-based tournament for bidding agents, the Trading Agent Competition (TAC) [Wellman *et al.*, 2001] was held in conjunction with ICMAS-00. This competition was much more realistic in design, and required the agents to simultaneously participate in multiple markets, each of which required a different bidding strategy. Our study incorporates a degree of realism in market dynamics and messaging/communication similar to that of TAC, while preserving a classical CDA design. This allows both agents and humans to participate, as well as facilitating comparisons with prior all-human and all-agent CDA studies.

3 Overview of the Experiments

For our experiments with humans and agents, we developed a hybrid system that combined GEM, a special-purpose distributed system for experimental economics developed by members of the IBM Watson Experimental Economics Laboratory (WEEL), with Magenta, a prototype agent environment developed at IBM Research. The hybrid configuration is illustrated in Figure 1.

A real-time, asynchronous CDA was administered by a GEM auctioneer process running on a Windows NT workstation, which communicated with all bidders and executed trades when appropriate. To ensure that agents and humans could interact seamlessly with one another, and that there would be no subtle bias in their treatment, humans and agents used the same set of messages to communicate with the GEM auctioneer. Each human bidder was given a Windows NT workstation running a GEM client process that interpreted messages from the GEM auctioneer, and encoded messages to be sent back to it. This GEM client offered a GUI that permitted its user to view the order queue, the trade history, and his/her assigned parameters. It also permitted the user to enter bids or asks. Each Magenta bidder agent participated in the auction through a modified version of the GEM client that forwarded messages via TCP/IP to a UNIX workstation running the Magenta environment and the agents themselves. Each agent received messages forwarded to it by its modified GEM client. Whenever it wished, the agent could send messages to its GEM client, which forwarded them as quickly as possible to the GEM auctioneer. Thus the Magenta agents and the human bidders had access to identical streams of data from the auctioneer, and the auctioneer could not distinguish orders placed by humans from those placed by agents.

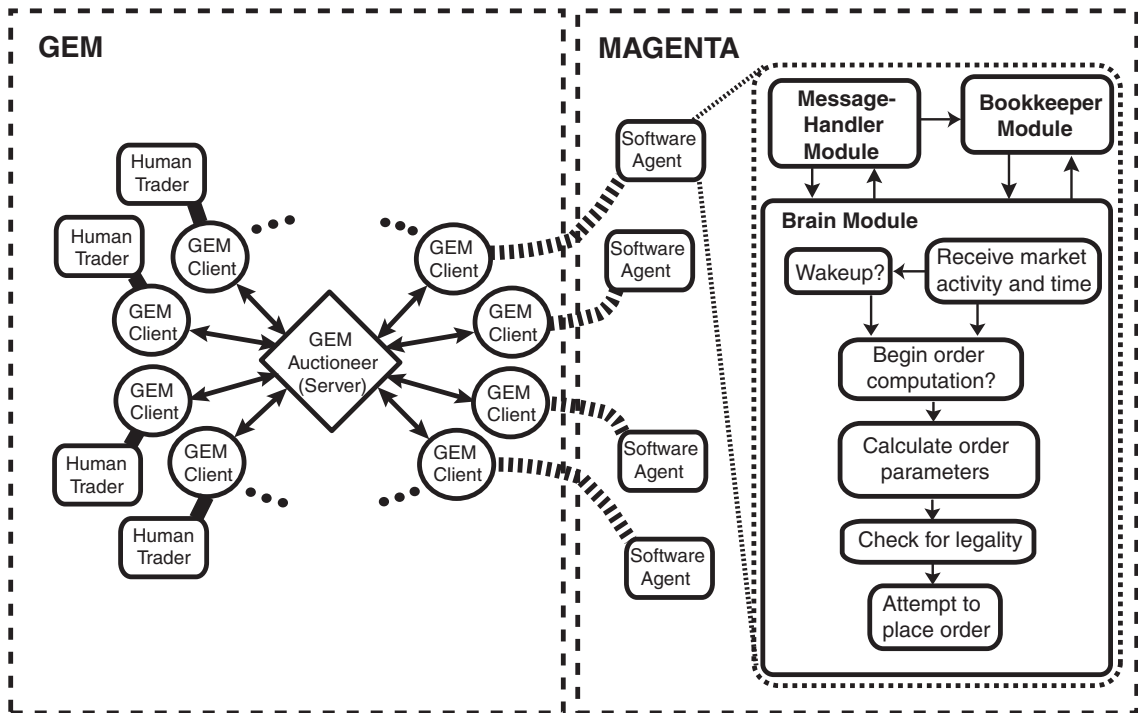


Figure 1: Hybrid GEM-Magenta configuration. At the left is the GEM system, showing a GEM auctioneer communicating with a set of GEM clients. Each GEM client either presents a GUI to a human trader or communicates via TCP/IP with one of the Magenta bidding agents shown at right. Also at right is an expanded view of an agent’s architecture, showing the three internal modules described in the text and sketching the control flow within the Brain module.

4 Agents for the CDA

4.1 Agent environment and architecture

The bidder agents were implemented on top of Magenta, which provides messaging, naming, and directory services, and supports both one-shot messaging and conversations (correlated sequences of messages). Each bidder agent contained a MessageHandler module, a Bookkeeper module, and a Brain module, all of which were specifically tailored for the CDA.

The MessageHandler was responsible for formatting and sending outgoing messages, and for receiving and parsing incoming messages. Incoming messages included initialization information, notifications of bids and asks placed by other agents, of trades, and of time remaining in an auction period. Upon receipt of a message, the MessageHandler would parse it and take an appropriate action, such as handing the record of the bid, ask, or trade to the Bookkeeper module.

The Bookkeeper maintained the agent’s internal representation of the market state (e.g., the history of orders and trades, current open orders, time remaining in the period, etc.) and of the agent’s internal state (e.g., orders the agent had placed, the agent’s inventory and available funds, etc.).

The Brain module was responsible for placing bids and asks. Each Brain placed orders on its own thread of execution, with its own activation schedule. Each Brain contained a module for a bidding strategy that determined order parameters based on the information stored in the Bookkeeper, and a timing module that governed the circumstances under which

the brain’s execution thread would wake up, apply the bidding strategy and (possibly) place an order. Outgoing order messages were formatted and sent back to the GEM client by the MessageHandler. Optionally, the timing module would awaken the brain’s execution thread whenever a trade occurred and/or when a particularly attractive bid or ask was placed by another player.

4.2 Agent strategies

The Magenta bidding agents face the following task in a live auction: given the time remaining T in the trading period, a number N of tradeable units, a vector \vec{L} of N limit prices (one for each unit), and the history H of previous activity in the market, calculate an order time τ and a price p .

The timing of orders was governed by a simple heuristic based on a sleep-wake cycle. The sleep time was set to a fixed interval of s seconds, with small random jitter $\pm(0 - 25)\%$ added. In our experiments, “fast” agents were defined by setting $s = 1$ and by permitting wakeup on all orders and trades. “Slow” agents were defined by setting $s = 5$ and allowing wakeups only on trades. When an agent wakes up, it computes an order price using its pricing algorithm. The agent will then submit the order, provided that the computed price is compatible with the market rules and its understanding of the current market conditions. If the agent currently has an open order, the order will be a replacement order; otherwise it will be a new one.

The pricing algorithms employed in these experiments are based on algorithms originally introduced for simpler ver-

sions of the CDA that lacked a persistent order queue, and made assumptions about market dynamics that are inconsistent with the notion of real-time, independent agents. We now describe modifications that we made to these algorithms to tailor them to our version of the CDA.

Zero-Intelligence-Plus (ZIP) Strategy

Cliff [Cliff and Bruten, 1997] proposed an algorithm called “Zero-Intelligence-Plus” (ZIP) to explore the minimum degree of agent intelligence required to reach market equilibrium in a simple version of the CDA. The market dynamics studied in [Cliff and Bruten, 1997] were unrealistic in that there was no explicit notion of time, no definite period length, and no persistent orders: submitted orders were either traded or removed instantaneously. We have modified ZIP to function in a real-time market with a definite period length and persistent open orders. The primary modification has to do with outbidding or undercutting existing orders. This now happens when orders remain open for a certain amount of real time without being traded. Our modifications turn out to be related to those independently proposed by Preist & van Tol [Preist and van Tol, 1998].

In our ZIP implementation, each agent maintains a vector of internal price variables \vec{p} ; the i -th component of \vec{p} , p_i , is used to set the order price when trading the i -th unit. At the start of trading, \vec{p} is initialized to random positive-surplus values, and is adjusted during the period according to the observed trading action.

When a trade occurs at trade price p_T , each p_i is adjusted by a small random increment in the direction of p_T . If the adjustment is in the direction of increasing profit margin (i.e. raising p_i for sellers and lowering p_i for buyers), the change is always made regardless of whether or not the i -th unit has already been traded. However, for adjustments in the direction of decreasing profit margin, the change is made only for units that are “active,” i.e., have not yet been traded. The size of the adjustment is proportional to a learning rate parameter, similar to that used in Widrow-Hoff or in back-propagation learning. The difference between p_i and p_T is also stored for use at the next trade, when a further adjustment in the same direction is made, proportional to a separate “momentum” parameter. This is analogous to the use of momentum to speed up convergence in back-propagation learning.

If a sufficiently long time has passed without a trade taking place (1.0 seconds in our implementation), ZIP buyers and sellers adjust p_i in the direction of improving upon the best open competing bid or ask, if the i -th unit is still active. Finally, there is a global constraint that each p_i must always correspond to non-negative agent surplus, i.e. it must always be below the buyer’s value, or above the seller’s cost.

In all-agent tests, we find that homogeneous populations of ZIP traders achieve robust convergence to theoretical equilibria with high efficiency. Depending on the precise market rules and initial conditions, efficiencies ranging from 0.980 to 0.999 have been obtained with this strategy.

Gjerstad-Dickhaut (GD) Strategy

Gjerstad and Dickhaut [Gjerstad and Dickhaut, 1998] introduced a more sophisticated trading algorithm for buyers and sellers in the CDA, which we shall term “GD”. They showed

via simulation that a homogeneous population of such agents could attain high allocative efficiency and rapid convergence to the theoretical equilibrium price. A GD agent constructs an order and trade history H , consisting of all orders and trades occurring since the earliest order contributing to the M th most recent trade. From the history H , a GD buyer or seller forms a subjective “belief” function $f(p)$ that represents its estimated probability for a bid or ask at price p to be accepted. For example, for a seller,

$$f(p) = \frac{AAG(p) + BG(p)}{AAG(p) + BG(p) + UAL(p)}, \quad (1)$$

where $AAG(p)$ is the number of accepted asks in H with price $\geq p$, $BG(p)$ is the number of bids in $H \geq p$, and $UAL(p)$ is the number of unaccepted asks in H with price $\leq p$. Interpolation is used to provide values for $f(p)$ for prices at which no orders or trades are registered in H . The GD agent then chooses a price that maximizes its expected surplus, defined as the product of $f(p)$ and the gain from trade at that price (equal to $p - l$ for sellers and $l - p$ for buyers, where l is the seller cost or buyer value). Thus the algorithm does not require the knowledge or estimation of other agents’ costs or valuations.

The original GD algorithm was developed for a market in which there was no queue, i.e. old bids or asks were erased as soon as there was a more favorable bid/ask or a trade. In our version of the CDA, unmatched orders can be retained in a queue, and therefore the notion of an unaccepted bid or ask becomes ill-defined. We addressed this problem by introducing into the GD algorithm a “grace period” τ_g . Unmatched orders were not included in H unless at least the grace period τ_g had passed since that order had been placed. Another modification to GD addressed the need to handle a vector of limit prices, since the original algorithm assumed a single tradeable unit.

We also found empirically that the original GD algorithm could behave pathologically, particularly for “fast” agents, which placed orders whenever an order or trade had been placed in the market. This often resulted in rapid bursts of orders and trades. If the last $2M$ orders resulted in M successful trades, then there were no unsuccessful orders in the history H . Laboring under the false assumption that *any* price would be accepted, the agents would then place absurdly low bids or high asks, gradually lowering them until trades finally began to occur once again, often in another burst. This cyclical behavior was associated with high trade price volatility.

To greatly reduce the chances of this occurring, we used a softer form of history truncation. All of the simple tally terms in Eq. 1 (and the analogous expression for buyers) were replaced by weighted sums that placed exponentially more emphasis on events that occurred most recently, and the truncation parameter M was increased to a much larger value. As hoped, soft truncation led to more sensible and stable pricing. It allowed the desired responsiveness to recent events, but also permitted information from old events to be used whenever there was insufficient information from recent events.

Homogeneous populations of modified GD agents also achieve robust convergence to equilibrium, with efficiencies comparable to those obtained by ZIP agents.

5 Experimental setup

Our experiments used the following CDA market rules:

1. The “NYSE” spread-improvement rule was in effect, requiring that new bids be priced higher than the current best bid (and the equivalent for asks). This conforms to prior CDA studies and is believed to facilitate convergence to equilibrium.
2. All orders were for a single unit only, and a player could have at most one open order. This was meant to simplify the task for both agents and humans, and again conforms to prior CDA studies.
3. Submitted orders remained open until they were traded or the period ended.
4. Submitted orders could be modified (subject to the NYSE rule), but not withdrawn.
5. Trades occurred when the best bid and best ask matched or crossed in price. If they crossed (i.e., bid price exceeded the ask price), the trade price was the price of the order submitted first.
6. At the start of each period, players were given a fresh supply of cash or commodity.

Each player was given a list of 8-14 limit prices for the units to be traded, arranged in order from most to least valuable (i.e., the buyer values decreased and the seller costs increased). Roughly half of the players’ units were tradeable for positive surplus at equilibrium. The limit prices were generated from a base set of three linear schedules in which each successive unit increased in cost or decreased in value at a constant rate. These rates varied in the three schedules; however, the total theoretical surplus was designed to be about the same in each. Each human had an agent counterpart with the same role and the same limit prices, and hence the same theoretical surplus. The total theoretical surplus was designed to split about evenly between buyers and sellers.

An experiment consisted of 15-16 trading periods of 3 minutes each. Every 4-5 periods, each player’s limit prices were changed by rotating the three limit price schedules (e.g., seller A received seller B’s previous schedule, B received C’s, and C received A’s) and adding or subtracting a constant value to all limit prices, so as to change the equilibrium price.

Our target configuration for the experiments consisted of 6 agents and 6 humans², both split evenly between buyers and sellers. In each experiment, all agents used the same bidding strategy—either ZIP or GD—and were all either the “Fast” or “Slow” variant. Human subjects in four experiments were undergraduates from local colleges; in two others, they were employees of IBM Research. Before the start of each experiment, subjects received instruction on the auction rules and the profit objective, and practiced using the GUI. No discussion of bidding strategies was given. Subjects were told they would receive cash payments proportionate to profits earned in the auction; the conversion factor was set so that the expected payouts were $\sim \$50 - 60$ per player.

²In some experiments, one or more of the scheduled subjects failed to appear, resulting in an asymmetric market with more agents than humans.

6 Experimental results

Table 1 summarizes the results of the six agent-human CDA experiments. Several noteworthy findings were obtained. First, there were significant interactions and trades between agents and humans, even though the agents were potentially much faster. Roughly 30% or more of all trades were between agents and humans. This is a reasonable fraction of the naive expectation of 50% if any trade partner (agent or human) is equally likely, and shows that the laboratory markets did genuinely test agent-human interactivity, as opposed to merely creating two non-interacting sub-markets operating on different time scales.

Second, when considered as a group, the agents outperformed the humans in all six experiments: the total surplus obtained by agents was on average $\sim 20\%$ more than the total human surplus. This was true for both fast and slow agent populations, showing that speed was not the sole factor accounting for the agents’ edge in performance. In terms of average efficiency, agents in aggregate tended to achieve greater than 100% efficiency, which necessarily implies that they were exploiting human errors or weaknesses. Humans, on the other hand, tended to score in the range of $\sim 0.92 - 0.96$, and on two occasions they did much worse. To check our market design, we ran a baseline experiment in which all 12 traders were human, measuring an efficiency of 0.96, which is consistent with what is typically found in all-human CDA experiments. The fact that humans play better against other humans than they do against agents corroborates the evidence that, as a group, the agents are stronger players.

Third, as in prior all-human CDAs, human performance tended to improve during the course of an experiment, as the subjects became more familiar with the GUI and the market behavior, and got a better idea of how to execute a good bidding strategy. Nevertheless, we still found a consistent edge in agent surplus over human surplus by about 5 – 7% in the final periods of each experiment.

Finally, although it is not documented in Table 1, our agent-human markets tended to have a lopsided character in which either buyers consistently exploited sellers, or vice versa. The previously-described scalloping behavior was observed to be more pronounced and longer-lasting than in prior all-human or all-agent CDAs. This will be discussed in detail in Sections 6.1 and 6.2, which examine two specific experiments with different agent bidding strategies (Fast GD and Slow ZIP), yielding distinctly different market dynamics.

6.1 Fast GD Agents vs. Humans

Figure 2 shows the trading activity in experiment Oct25, which was conducted over 16 periods divided equally among four phases with shifts in limit prices. In each period, the time series of trades tends to show scalloping, as trade prices converge towards the equilibrium price p^* . Such scalloping was not observed in agent-only CDA experiments with GD Fast (or Slow) agents. In this experiment, the buyers were able to extract more surplus from the market than the sellers as most trades occurred below p^* , a fact that is also reflected in the average efficiency measures. However, the differences in surpluses (and efficiencies) between the two sides of the market shrank over time, due to improving overall market efficiency,

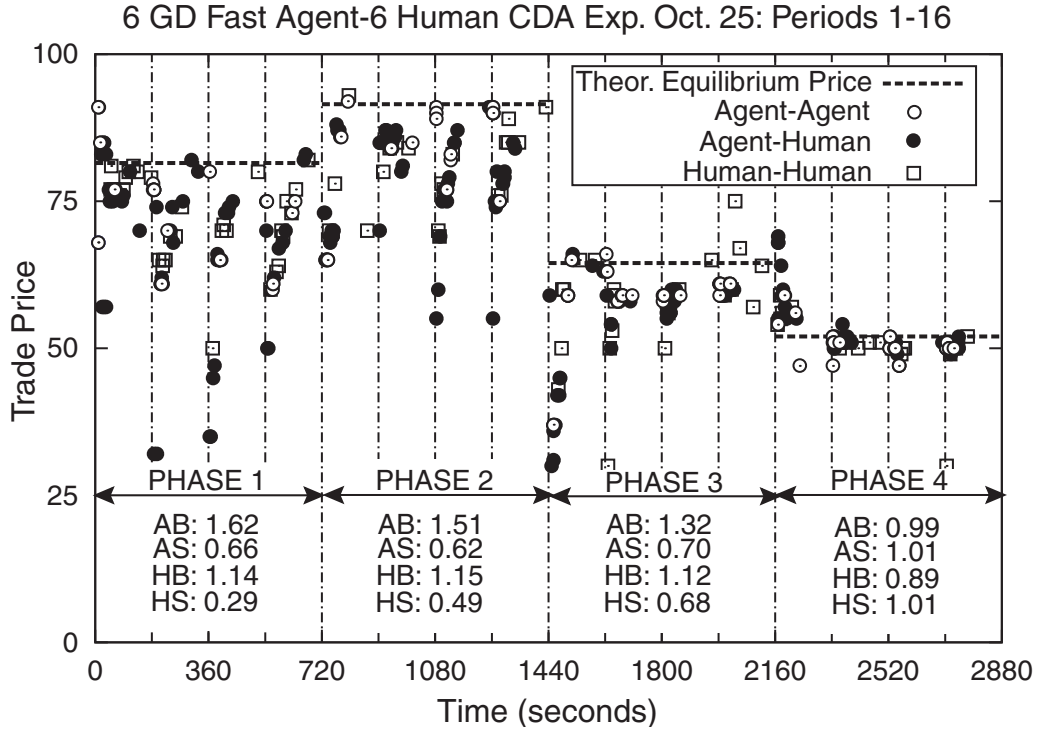


Figure 2: Trade price vs. time for experiment Oct25 with 6 GD fast agents and 6 humans. The vertical dashed lines indicate the start of a new trading period. The 16 trading periods were divided into four phases, each with its own set of limit prices. In each phase, p^* is shown by the horizontal dashed lines. Trades between two agents are shown with open circles, between two humans with open squares, and between an agent and a human with solid circles. The labels AB, AS, HB, and HS refer to the average efficiency of Agent Buyers, Agent Sellers, Human Buyers, and Human Sellers, respectively.

which we attribute to human improvement during the course of the experiment. Note that in all four phases, the agent buyers and the agent sellers are able to extract more surplus than their human counterparts.

Figure 2 also shows that most of the lowest-priced trades, occurring well below p^* , were between agents and humans. An inspection of experimental records reveals that these trades were mostly between human sellers and agent buyers. Apparently the human sellers were consistently offering excessively low asks, and the agent buyers were able to pounce on such mistakes more quickly than their human counterparts.

6.2 Slow ZIP Agents vs. Humans

Figure 3 shows the trading activity in the third phase of experiment Oct24a with 6 ZIP slow agents and 6 humans. The figure shows several interesting features. First, pronounced and repeated scalloping in trade prices is evident, with buyers extracting much more surplus than sellers. This is surprising since such large scalloping is not seen in markets with only ZIP agents. Second, in each period, trades typically tended to occur first between agents, then between agents and humans, and finally between humans. Third, although the agents as a group outperformed the humans, agent sellers actually ob-

tained less surplus than human sellers.

Subsequent interviews with human subjects helped to explain the behavior and ultimately reveal a weakness in the ZIP strategy. It turned out that two of the human sellers in this experiment consistently followed a ‘fixed-profit-ratio’ strategy, that is, their asking price for each unit was a fixed percentage greater than its cost. These sellers repeatedly submitted asks at prices much lower than p^* and most often, agent buyers quickly accepted such offers. Having traded their units at extremely low prices, the ZIP buyers ignored subsequent trades at higher prices, and maintained very low bid prices at the start of the next period. The resulting bid-ask spread at the start of each period was centered well below p^* , and once the human sellers made a few low-priced sales, the ZIP sellers, which were waking up on trades, began quickly dumping their inventory at very low prices. Hence we attribute the unusual market behavior, and the performance ranking of the agent and human traders, to a set of odd interactions between the ZIP strategy and a specific non-optimal human strategy.

To test our hypothesis, we performed separate all-agent experiments with a fixed-profit-ratio agent in a population of only ZIP agents. The resulting dynamics exhibited large scallops in trade prices similar to those in Figure 3. We also

Experiment				Agent			Human		
ID	# Periods	# Trades	Interaction	Strategy	Surplus	Efficiency	# Traders	Surplus	Efficiency
Oct17	15	412	0.38	GD Fast	11058	1.016	5	6991	0.927
Oct18	15	504	0.29	ZIP Fast	11069	1.028	6	7023	0.652
Oct23	16	320	0.33	GD Fast	10495	0.999	3	4582	0.965
Oct24a	16	455	0.48	ZIP Slow	10696	1.032	6	9490	0.916
Oct24b	9	261	0.42	GD Fast	6808	1.026	6	6353	0.958
Oct25	16	433	0.49	GD Fast	12159	1.052	6	9708	0.840

Table 1: Summary of the six agent-human CDA experiments. For each experiment, the table presents: the number of trading periods, the total number of successful trades, the fraction of trades between agents and humans, the bidding strategy employed by all six agents, the number of human traders, and the aggregate agent and human performances in terms of total surplus accumulated over the entire experiment and the average efficiency.

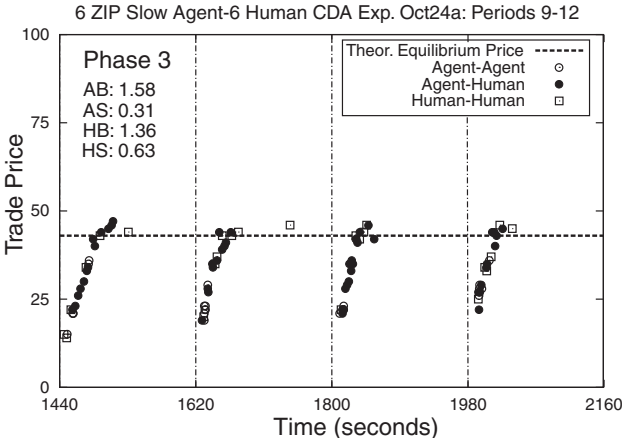


Figure 3: Trade price vs time for experiment Oct24a with 6 ZIP slow agents and 6 humans. Trading activity in periods 9-12 (out of 16 total) is shown. Other details are similar to those in Figure 2.

implemented a modification to the ZIP strategy that is more reluctant to lower its profit margin based on where trades occurred in the last period. Preliminary results show that the modified ZIP agents retain high efficiency and are not easily misled by the fixed-profit-ratio agent.

7 Conclusion

Given the simplicity of our agent bidding strategies, we are encouraged by the results of these first-ever tests against human subjects. (We remark that, while substantial agent-human interactions have already occurred in financial markets, such interactions have an unknown and uncontrolled nature.) Our agents make relatively simple time-independent price calculations, based on established algorithms, and their timing decisions are equally simple, yet they are able to outperform non-expert human subjects by a clear margin. It would be interesting to test our bidding agents against a higher caliber of human opposition, e.g., professional equities or commodities traders. We suspect that such opponents would uncover weaknesses in the agent strategies, and that this would eventually lead to significant algorithmic improve-

ments in the strategies. We are optimistic that CDA strategies can be improved to the point where they outperform all human opposition by making better price inferences based on market history, and by taking time remaining into account in pricing and timing decisions.

Several aspects of the market behavior in our experiments differed from prior studies of all-agent or all-human traders. Convergence to equilibrium in our experiments was generally slower than in prior studies, and in two experiments, there was no evidence of convergence by the end of the experiment. We observed scalloped price trajectories that were more pronounced and longer lasting than seen previously. Also, our markets tended to be much more lopsided (either buyers systematically exploiting sellers, or vice versa) than in earlier studies. Such novel market phenomena merit further investigation: they might be due to specifics of our market design, or they may be more general outcomes of agent-human interactions. Hence it will be important to conduct further agent-human experiments in other types of markets, possibly including greater complexity and real-world detail. Candidates include combinatorial auctions, TAC-type markets [Wellman *et al.*, 2001], and more realistic models of financial markets. The development and deployment of effective automated trading strategies in such markets would have immense practical importance, and could mark the beginning of a large-scale introduction of economic software agents into the world economy [Kephart *et al.*, 2000].

While our results are preliminary, some aspects of our findings may be indicative of what one can expect to occur more generally as economic software agents are developed for real-world markets. We suspect that, in many real marketplaces, agents of sufficient quality will be developed such that most agents beat most humans. A significant component of their advantage will come from their ability to initiate actions, and to react to market events, much faster than humans. As a result, there will be significant economic incentive for humans to employ agents to act on their behalf. Then the competition between agents and humans will evolve into a competition among agents.

Acknowledgments

The authors are especially indebted to Steven Gjerstad, who made several invaluable contributions to these experiments.

Steven co-authored the paper that first described the GD bidding algorithm, which we adapted and extended to handle different market rules. Steven was also responsible for the design of the Watson Experimental Economics Laboratory (WEEL) where the experiments were conducted; for the design of the double auction market rules; for the design of the GEM system that was used as the electronic auctioneer and the auction clients for human traders; and for recruiting most of the experimental subjects.

We would like to thank several other people for their significant contributions. Jason Shachat helped develop the experimental protocol and supervise some of the experiments. Amit Shah helped implement the messaging link between the GEM and Magenta systems. Oconel Johnson, system administrator for WEEL, helped ensure that the GEM hardware and software were ready for the experiments. Weng-Keen Wong and Jonathan Bredin developed the Magenta message handling components for the agents and a Magenta auctioneer used for stand-alone testing, and they implemented and tested several bidding algorithms. David Levine resolved several bugs and inefficiencies in the GEM-Magenta link, and was an invaluable resource on architectural questions.

We would also like to thank the many students and IBM researchers who participated in the experiments, including some early trial experiments not reported here. Finally, we would like to acknowledge the support of the IBM Institute for Advanced Commerce.

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, Hewlett Packard Labs, 1997.
- [Gjerstad and Dickhaut, 1998] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
- [Kephart *et al.*, 2000] J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32:731–752, 2000.
- [Preist and van Tol, 1998] C. Preist and M. van Tol. Adaptive agents in a persistent shout double auction. In *Proceedings of the First International Conference on Information and Computation Economies*, pages 11–18. ACM Press, 1998.
- [Rust *et al.*, 1992] J. Rust, J. Miller, and R. Palmer. Behavior of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, Redwood City, CA, 1992. Addison-Wesley.
- [Schaeffer, 2000] J. Schaeffer. The games computers (and people) play. In M. Zelkowitz, editor, *Advances in Computers 50*, pages 189–266. Academic Press, 2000.
- [Smith, 1962] V. L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70:111–137, 1962.
- [Smith, 1982] V. L. Smith. Microeconomic systems as an experimental science. *American Economic Review*, 72:923–955, 1982.
- [Wellman *et al.*, 2001] M. Wellman, P. Wurman, K. O’Malley, R. Bangera, S. Lin, D. Reeves, and W. Walsh. Designing the market game for a trading agent competition. *IEEE Internet Computing*, pages 43–51, March/April 2001.