CSc 84200

E-Commerce and Computational Economics

Agent-Human Interaction 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

Minimal-Intelligence Agents for Bargaining Behaviors in Market Based Environments

D. Cliff and J. Bruten. Technical Report HPL-97-91, Hewlett Packard Labs, 1997.

Constantinos Djouvas

Agent-Human Interaction in the Continuous Double Auction

Introduction

- Present situation: Simple bidding agents by e-bay and Amazon
- Expected: continued growth in the variety and sophistication of automated economic decision-making technologies.
- Agents:
  - They will have to be economically intelligent, capable of making effective decisions about pricing, purchasing, or bidding.
  - Their economic performance must exceed that of humans on average, otherwise, people will not entrust agents with making economic decisions.
- This paper provides a demonstration of agents competing against humans.

Continuous Double Auction (CDA)

- The dominant market institution for real-world trading of equities, commodities.
- An environment where both humans and agents can participate simultaneously.
- This paper study of economic interactions between agents and humans utilizes a simplified model of a Continuous Double Auction (CDA) market.

Continuous Double Auction (CDA)

- A fixed trading period, during which 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.

CDA model, as adapted in the article

- Multiple units of a single hypothetical commodity can be bought or sold.
- Participants are assigned a fixed role, either a Buyer or Seller.
- There are several periods of trading; at the start of each period participants are given a list of "limit values" (private value) for each unit to be bought or sold.
- The limit values 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 value – trade price) for buyers and (trade price – limit value) for sellers.

CDA model, as adapted in the article

- 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)
  - Demand (the total number of units 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 price p^*.

Experiment

- A hybrid system that combined GEM, a distributed system for experimental economics was developed.
- To ensure that agents and humans could interact seamlessly with one another, humans and agents used the same set of messages to communicate with the GEM auctioneer.
- Agent and 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.

Experiment

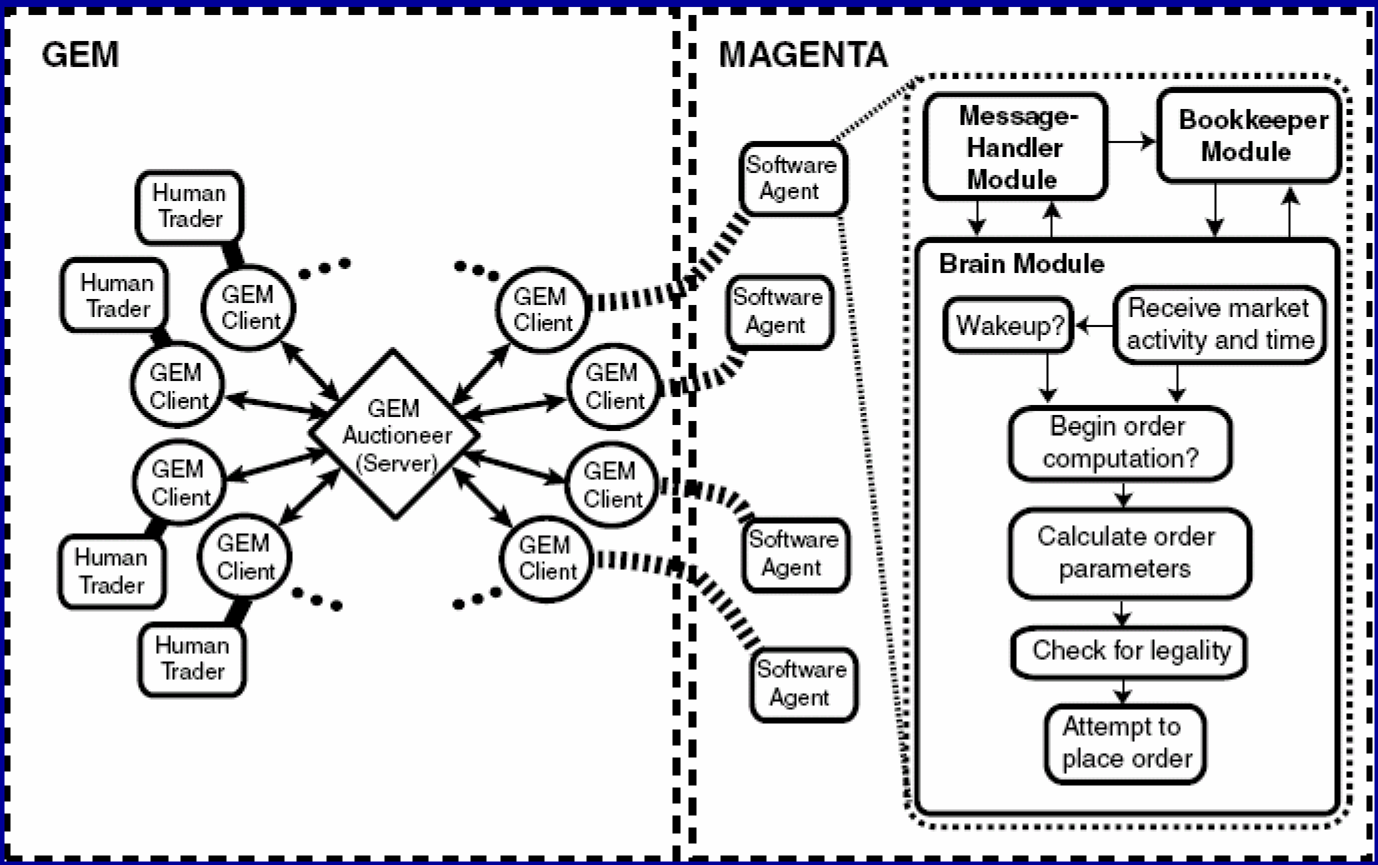

GEM
MAGENTA
Human Trader
Human Trader
Human Trader
Human Trader
GEM Client
GEM Client
GEM Client
GEM Client
GEM Auctioneer (Server)
GEM Client
GEM Client
GEM Client
GEM Client
Software Agent
Software Agent
Software Agent
Software Agent
Message-Handler Module
Bookkeeper Module
Brain Module
Wakeup?
Receive market activity and time
Begin order computation?
Calculate order parameters
Check for legality
Attempt to place order

Agents for the CDA

- The time order was based on a sleep-wake cycle.
- The sleep time was set to a fixed interval of seconds.
- Two types of agents based on timing:
  - “Fast” agents: s=1, wake up on all orders and trades
  - “Slow” agents: s=5, wake up only on trades.

- After waking up, the agent computes an order price using its pricing algorithm.

- Two types of agents based on pricing strategy:
  - Zero-Intelligence-Plus (ZIP) strategy
  - Gjerstad – Dickhaut (GD) strategy.

Gjerstad – Dickhaut (GD) Strategy

- Each agent constructs an order and trade history *H*.
- Based on *H,* a GD buyer or seller forms a subjective "belief" function *f(p)*

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

- *Where:*
  - *AAG(p)* is the number of accepted asks in *H* with price ≥ *p*
  - *BG(p)* is the number of bids in *H* with price ≥ *p*
  - *UAL(p)* is the number of unaccepted asks in *H* with price ≤ *p*

Gjerstad – Dickhaut (GD) Strategy (Version used for simulation)

- Old asks and bids were retained in a queue.
- A vector of limit prices was handled (original algorithm assumed a single trade-able unit).
- Empirically found that the original GD model could behave pathologically for “fast” agents, which placed orders whenever an order or trade had been placed in the market.
  - There are no unsuccessful orders in the history.
  - False assumption that any price would be accepted.
  - Agents place absurdly low bids or high asks.
  - Gradually lowering them until trades began to occur again.
- Reduced this phenomenon using softer form of history truncation.

Zero-Intelligence-Plus (ZIP) strategy

Minimal-Intelligence Agents for Bargaining Behaviors in Market Based Environments

D. Cliff and J. Bruten. Technical Report HPL-97-91, Hewlett Packard Labs, 1997.

Pages 41 – 63

- Explores the minimum degree of agent intelligence required to reach market equilibrium in a simple version of the CDA

Zero-Intelligence-Plus (ZIP) strategy

- Profit margin determines the difference between the traders limit price and shout-price.
- Initially, the only information known to a trader is the limit prices for the units the trader is entitled to sell or buy.
- Traders adjust their profit margins using market price information.

Zero-Intelligence-Plus (ZIP) strategy

- Each ZIP trader alters its profit margin on the basis of four factors:
  1. The trader is active (still capable of making a transaction) or inactive (has sold or bought its full entitlement of units and has dropped out of the market)
- The other three factors concern the last shout:
  2. Its price denoted by q
  3. Whether it was a bid or an offer
  4. Whether it was accepted or rejected
- Shout Price (p) : profit margin (μ) x limit price (λ) (Increase in μ raises p for a seller and lowers p for a buyer)

Zero-Intelligence-Plus (ZIP) strategy

- Buyer: buy from any seller that makes an offer less than the buyers current bid shout price.
- Seller: sells to any buyer making a bid greater than the sellers current offer shout price.

Sellers behavior

If the last shout was accepted at price q

then

any seller s_i for which $p_i \leq q$ should raise its profit margin

if the last shout was a bid

then

any active seller s_i for which $p_i \geq q$ should lower its margin

Else

if the last shout was an offer

then

any active seller s_i for which $p_i \geq q$ should lower its margin

Buyers behavior

If the last shout was accepted at price q
then
any buyer b_i for which $p_i \geq q$ should raise its profit margin
if the last shout was an offer
then
any active buyer b_i for which $p_i \leq q$ should lower its margin
Else
if the last shout was an bid
then
any active buyer b_i for which $p_i \leq q$ should lower its margin

Adaptation

- At a given time t, a ZIP trader i calculates the shout-price $p_i(t)$ for unit j with limit price $\lambda_{i,j}$, using the profit margin $\mu_i(t)$ according to:

$$p_i(t)=\lambda_{i,j}(1+\mu_i(t)) \quad (1)$$

where: for sellers $\mu_i(t) \in [0, \infty); \forall t$

for buyers $\mu_i(t) \in [-1, 0]; \forall t$

- Simple update rule (Widrow-Hoff "delta rule")

$$A(t+1)=A(t)+\Delta(t) \quad (2)$$

where: $A(t)$ is the actual output at time t

$A(t+1)$ is the actual output on the next time step

$\Delta(t)$ is the change in output, where

$$\Delta(t)=\beta(D(t)-A(t)) \quad (3)$$

and β is learning rate coefficient (learning speed)

$D(t)$ is the desired output at time t.

Adaptation

- Rearranging (1) we can get the profit margin μ_i on the transition from time t to $t+1$:

$$\mu_i(t+1)=(p_i(t)+\Delta_i(t))/\lambda_{i,j}-1 \quad (4)$$

where:

$$\Delta_i(t)=\beta_i(r_i(t)-p_i(t)) \quad (5)$$

and $r_i(t)$ is the target price.

- There are many ways in which the target price could be determined.

Adaptation

- In the current ZIP traders the target price is generated using a stochastic function of the shout price $q(t)$

$$r_i(t)=R_i(t)q(t)+A_i(t) \quad (6)$$

where:

  - R_i is a randomly generated coefficient that sets the target price in relation to the price $q(t)$ of the last shout.
  - $A_i(t)$ is a small random absolute price alteration.

- To increase shout price R_i>1.0 *and* A_i>0.0.
- To decrease shout price R_i<1.0 *and* A_i<0.0.

Review

- Shout Price (p) : profit margin (μ) x limit price (λ)
- $\mu_i(t+1)=(p_i(t)+\Delta_i(t))/\lambda_{i,j}-1$
- $\Delta_i(t)=\beta_i(r_i(t)-p_i(t))$
- $r_i(t)=R_i(t)q(t)+A_i(t)$

	Increase	Decrease
R	[1.0,1.05]	[0.95,1.0]
A	[0.0,0.05]	[-0.05,0.0]
β	[0.1,0.5]	

Zero-Intelligence-Plus (ZIP) strategy (Version used for simulation)

- Each agent maintains a vector of limit prices $\vec{p}$.
- If a sufficiently long time has passed without a trade taking place (1.0 seconds), ZIP buyers and sellers adjust p_i in the direction of improving upon the best open competing bid or ask.
- There is a global constraint that each p_i must always correspond to non-negative agent surplus.

Experimental Results

- There were significant interactions and trades between agents and humans, even though the agents were potentially much faster.
- As a group, the agents outperformed the humans in all six experiments, with 20% average more than the total human surplus.
- 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.
  - A consistent edge in agent surplus over human surplus by about 5 – 7% was still found.
- Markets tended to have a lopsided character, in which either buyers consistently exploited sellers or vice versa.

Experimental Results

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

Summary of the six agent-human CDA experiments.

Experimental Results
GD Agents vs. Humans

- The buyers were able to extract more surplus from the market than the sellers as most trades occurred below p^*.
- The agent buyers and the agent sellers extracted more surplus than their human counterparts.
- Most of the lowest-priced trades 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.

Experimental Results

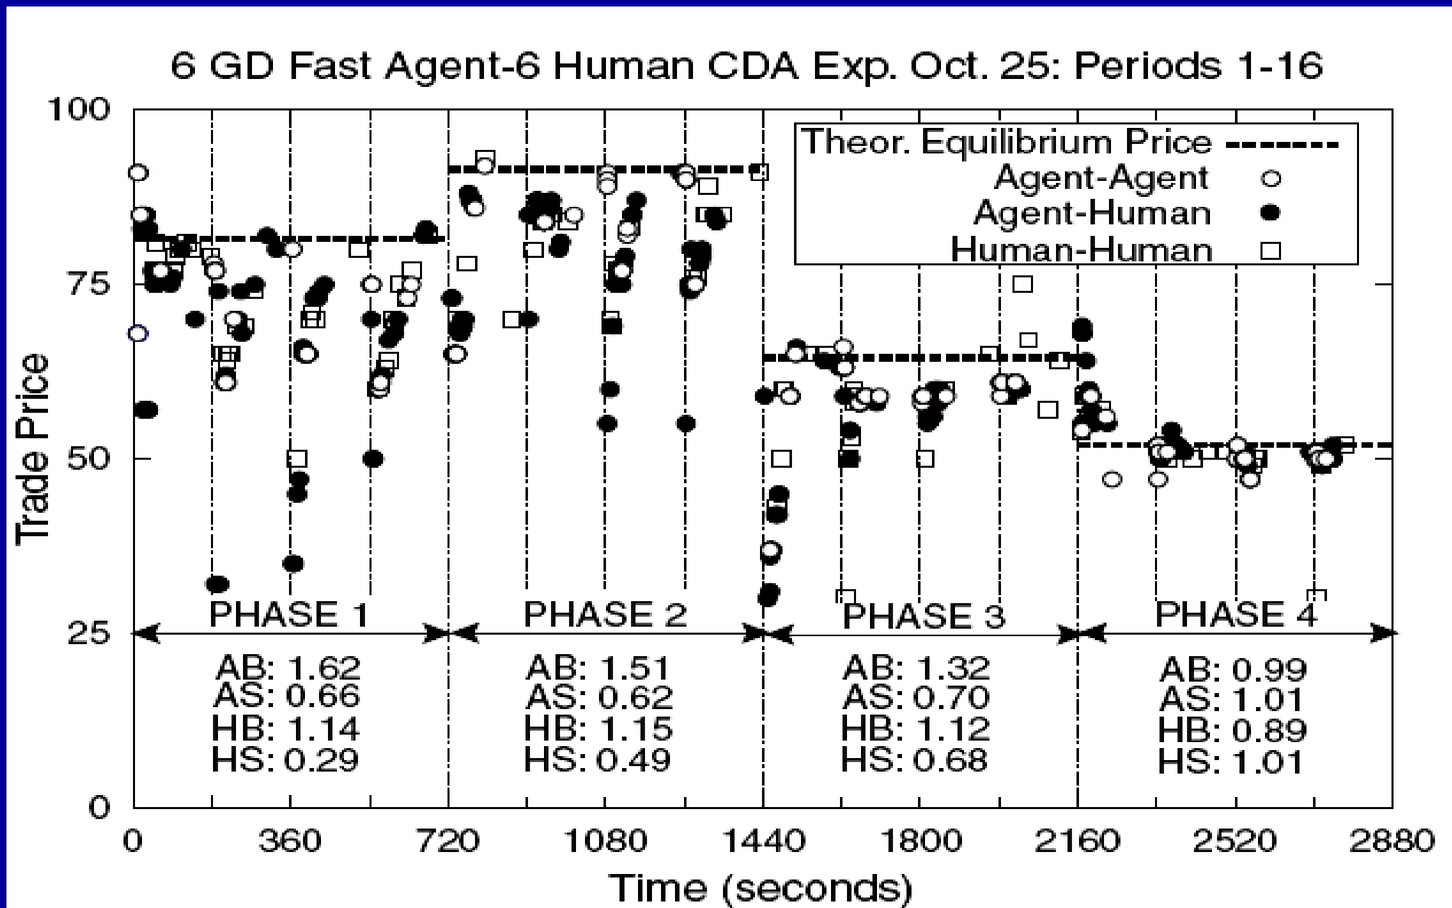

Dashed line represents the equilibrium price p^*

Experimental Results
ZIP Agents vs. Humans

- Buyers extracted much more surplus than sellers.
- In each period, trades typically tended to occur first between agents, then between agents and humans, and finally between humans.
- Although the agents as a group outperformed the humans, agent sellers actually obtained less surplus than human sellers.
  - ‘Fixed-profit-ratio’ strategy by same human sellers.

Conclusions

- In many real marketplaces, agents of sufficient quality might 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.