

Speculation Agents for Dynamic Multi-period Continuous Double Auctions in B2B Exchanges

Li Li

Graduate School of Industrial Administration
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: lil@andrew.cmu.edu

Stephen F. Smith

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: sfs@cs.cmu.edu

Abstract—Business-to-business (B2B) exchanges provide opportunities for companies to streamline their supply chains in dynamic business situations, but they also create new management challenges. Managers are faced with more frequent buying and selling decisions, more available information, and even new problems such as speculation. Due to these new challenges, trading support systems will play an important role in helping companies to achieve maximum profits in B2B exchanges. In this paper, we focus on providing support for two core capabilities, bidding and speculating, in the context of B2B exchanges with Continuous Double Auctions (CDA). Our work extends previous research on bidding agents to more dynamic and realistic B2B exchange situations where both demands and supplies change from period to period. Changes in demand and supply cause price fluctuations and motivate companies to speculate on inventory. The action of speculation is one fundamental aspect of exchanges in general, and is the main way in which B2B exchanges can hedge against costly shortages. We develop a multi-agent system in which a speculation agent makes speculation decisions for a given buyer or seller and interacts with a set of bidding agents. The related algorithms are discussed in detail. We also propose a theoretical model for the economic equilibrium in this dynamic situation. Under various market conditions, the experiments show that our system significantly increases the overall profit of the entire market.

I. INTRODUCTION

According to a recent study by eMarketer [2], despite the economy slow down, worldwide B2B e-commerce will increase to \$823 billion in 2002 from \$474 billion in 2001, and the strong growth will continue through 2004. As a fast growing section of e-commerce, Business-to-business (B2B) exchanges provide opportunities for cost savings through streamlining supply chain processes and opportunities for boosting revenue by bringing better products to market faster. However, B2B exchanges have also created new management challenges. First, in the dynamic business situation within a B2B exchange, decisions such as product buying and selling must be made more frequently and quickly. Furthermore, these decisions become more complicated and involve more information. For instance, rather than producing products in the quantity required by buyers at a fixed price, a seller must make decisions about how many products to offer in the market, at what price, how to respond to buyers' bids and compete with other sellers. To maximize profits, these decisions need to be made based the seller's production condition and the market

condition in the exchange. Finally, managers even have to confront new types of business problems, such as speculation. Price fluctuation is a basic property of B2B exchanges. When the price is at a low level, a manufacturer needs to decide how many components on top of current needs should be purchased to reduce future costs, *i.e.* how to speculate on inventory.

Due to these new challenges, trading support systems will play an important role in helping companies achieve maximum profits in B2B exchanges. As one approach to the design of the bidding functionality in such systems, agent-based bidding systems have gained significant research interest in recent years. In this paper, we will extend previous research and propose an agent-based framework that focuses on not only the bidding but also the speculating capability. The continuous double auction (CDA) is one of the most common exchange institutions and is extensively used in stock exchanges, commodity markets, and B2B exchanges. Therefore, we present our framework in the context of continuous double auctions.

The research on the design of agents for CDAs has its root in early experimental economics work. Smith [7] has conducted experimental double auctions with a small number of human traders and demonstrated a rapid convergence of transaction prices to the competitive equilibrium. This work has set up a framework for experimental analysis of competitive market behavior. Gode and Sunder [4] have replaced the human traders with "zero-intelligence" (ZI) programs that submit random bids and offers and has shown similar price convergence. Cliff and Bruten [1] have further designed "zero-intelligence-plus" (ZIP) agents by employing an elementary form of machine learning in ZI traders and achieved a performance significantly closer to the human data. Preist and van Tol [6] have used different heuristics for determining target profit margins in ZIP agents, which are demonstrated to achieve equilibrium significantly faster and are more robust. Gjerstad and Dickhaut [3] have proposed an agent algorithm, often referred to as the GD algorithm, from a different approach. Following this algorithm, each trader first estimates the probability for a possible bid or offer to be accepted based on recent market trading activities, and then places the bid or offer that maximizes its expected surplus. Tesauro and Das [8] have made several improvements to the GD algorithm. A principal limitation of these works is that they assume that the demand and supply do

not fluctuate over time, or in other words, auctions that are held at different times always assume the same demand and supply. This assumption is not valid in B2B exchanges, where demand constantly changes due to the changing economic condition and other reasons such as the product-life cycle of products.

In this paper, we extend previous works to the situations in B2B exchanges where both demands and supplies change from period to period. Specifically, we develop a multi-agent system in which a speculation agent makes speculation decisions for a given buyer or seller and interacts with a set of bidding agents.

The importance of speculation decisions to companies is the following. First of all, speculation in a dynamic market with price fluctuations is a rational decision to maximize profit. Moreover, speculation widely exists in stock exchanges, commodity markets and B2B exchanges. Therefore, a company must understand how other companies' speculation actions influence the market situation and react to them. Finally, speculation improves the performance of the market in terms of higher market efficiency, lower price volatility, fewer product shortages, and greater ability to meet changing demands.

In the next section, we analyze CDAs in B2B exchanges from the economics point of view, and discuss how to model competitive equilibria in B2B exchanges. We detail the agent framework and related algorithms in §III. Some simulation experiments and results are discussed in §IV and we conclude in §V.

II. AN ANALYSIS OF CDA IN B2B EXCHANGE

In this section, we start with the single period CDA and the related economics theory and then further explore the multi-period situation and model the behavior of B2B exchanges with CDA.

A. CDA in a Single Period

A continuous double auction is a market institution where a group of buyers and a group of sellers simultaneously and asynchronously announce bids and offers. At any time, a seller is free to accept the bid of a buyer, and a buyer is free to accept the offer of a seller. Compared with other styles of auctions, a CDA normally generate competitive outcomes more quickly and reliably.

In each period, sellers sell products they are able to produce in the current period. For each unit of product, there is a minimum price, called the *limit price*, below which the seller will not sell this unit. Let l_i^k be the limit price of seller i for her k th unit. The total supply in this period, noted by S , is the total number of these units. Given a price p , sellers are willing to sell all units with a limit price lower than p , and let function $S(p)$ be the number of these units. Naturally, $S(p)$, called *supply curve*, is non-decreasing in p .

In each period, each buyer demands multiple units of products. For the k th unit of buyer j 's demand, there is a limit price l_j^k , above which buyer j will not buy this unit. Define function $D(p)$ as the number of units with limit price higher than p . $D(p)$ is the total number of units buyers are

willing to buy at price p . $D(p)$ is non-increasing in p , and let $D = D(0)$.

At the price p^* s.t. $S(p^*) = D(p^*)$, all buyers and sellers are able to trade products at the quantity that they are willing to trade, and the market is "cleared." Price p^* is called *equilibrium price* or *competitive equilibrium*. In the CDA, since the limit prices of a trader (a buyer or seller) are her private information, each trader has no knowledge of the equilibrium price, and, therefore, a trade may take place at a price other than p^* . However, previous research has demonstrated that transaction prices generally converge to p^* . In economics, p^* is represented as the function $f(D, S)$, called the *price function*, due to the fact that D and S are the main parameters in curve $S(p)$ and $D(p)$.

B. CDA in Multi-period B2B Exchange

When we consider multiple periods, the situation is more complicated. The demand changes from period to period, while the available supply is constrained by the capacity of sellers. These variations in demand along with constrained supply cause price fluctuations. Price fluctuations motivate companies to accumulate inventory. For instance, when price is low, a buyer may purchase more products than current needs to cut future purchase cost, while a seller may be willing to keep more products on hand and sell at a higher price later. These extra products are called *speculation inventory*.

To further complicate matters, the speculation inventory in turn decreases price fluctuations by increasing demand when prices are low and increasing supply when prices are high. The new price fluctuation would then influence the speculation inventory. This interaction between the speculation inventory and the price fluctuation ends when equilibrium is reached.

To understand CDAs in this situation, we further consider the supply or demand of each seller and buyer in these auctions. At the beginning of period t , seller j sells both the products she will produce in this period and her speculation inventory from period $t - 1$. Limit prices of to-be-produced units are the same as those in the single period case, while limit prices of units from speculation inventory are prices at which they were purchased in period $t - 1$ plus the inventory cost of holding them.

As to the demand of buyer i in period t , we first consider the situation in period $t - 1$. In period $t - 1$, if i 's demand is not fully satisfied, the unsatisfied demand is called buyer i 's shortfall and noted by s_{t-1}^i . In most B2B exchanges, shortfalls at least partially contribute to the demand in the next period. Among a buyer's shortfall, we assume that a constant percentage, noted by α , with the highest limit prices is carried over to the next period. Besides the shortfall, i 's speculation inventory purchased in $t - 1$ is also carried over the period t . Therefore, buyer i 's total demand t equals its original demand, plus αs_{t-1}^i and minus its speculation inventory from $t - 1$.

III. THE AGENT-BASED FRAMEWORK AND RELATED ALGORITHMS

Based on the economic analysis above, we first discuss an agent-based framework for bidding and speculating in CDAs

conducted in a B2B exchange, and then further detail the algorithms for each agent.

A. The Agent-Based Framework

The overall framework is shown in Figure 1. A seller's agent system is composed of one speculation agent (SPA) and multiple selling agents (SAs), while a buyer's system contains a SPA and multiple buying agents (BAs) and SAs. Each BA or SA can only buy or sell one unit of products, say 1000 keyboards, while a SPA may buy multiple units for speculation purposes. Similar to the situation in stock markets, a B2B exchange opens for a certain amount of time, called a trading period, and then closes till the next period. In each period, SAs and BAs continuously decide their offers or bids based on the market situation and internal algorithms; SPAs continuously monitor the market situation and make speculation decisions. Each agent has no information about the type or owner of another agent unless they belong to the same buyer or seller, but all current offers and bids are common knowledge to all agents. Whenever the current highest bid in the market, noted by B_{max} , is equal to or higher than the current lowest offer in the market, noted by O_{min} , a trade occurs, and the BA and SA involved in this transaction are removed from the system.

To understand the interaction of agents inside the system of a buyer or seller, we first discuss different SAs. Besides an already-produced product, a currently available production capacity may be sold to buyers as long as the seller is able to meet the due date. Therefore, SAs may be classified into on-hand (product) SAs and capacity SAs. Moreover, the product or capacity may or may not belong to the speculation inventory, and therefore, a SA may be a speculation SA or a normal SA. In this paper, in order to focus on speculation, we assume that the production lead-time is short enough so that sellers need not produce a product before it is sold in the exchange. Thus, normal (capacity) SA, capacity speculation SA and on-hand speculation SA are three states that a SA may be in; the transfer from one state to another represents the change of what the SA is selling.

In the system of a buyer or seller, the SPA behaves like a coordinator and may control the SAs and BAs when necessary. The dynamics in a seller's system are as the following. At the beginning of a period, a normal SA is created for each unit of production capacity. During the auction, if the SPA decides to beef up its speculation inventory, it may either 1) try to buy a unit in the auction and then create an on-hand speculation SA for this unit, or 2) reserve the capacity for speculation, *i.e.*, convert a normal SA to a capacity speculation SA. Similarly, when the SPA decides to reduce its speculation inventory, it may either convert to a capacity speculation SA back to a normal SA or force an on-hand speculation SA to offer at the current highest bid for an instant sell. A SA is removed once its unit is sold. At the end of the period, all unused normal SAs are removed; all capacity speculation SAs are converted on-hand speculation SAs as the capacities are utilized, and all speculation SAs are carried over to the next period.

In the system of a buyer, an on-hand speculation SA is the only form of SAs. At the beginning of a period, if some SAs are carried over from the last period, these SAs are used to satisfy this buyer's own demands first, starting from the demand with the highest limit price. If the number of SAs is greater than the demand, the leftover SAs stay in the system; otherwise BAs are created for the remaining demand. During the auction, if the SPA decides to beef up its speculation inventory, it tries to buy a unit in the auction and create an on-hand speculation SA for this unit. When the SPA decides to reduce its speculation inventory, it forces an on-hand speculation SA to offer at the current highest bid. At the end of the period, all SAs are carried over and BAs are removed.

The limit price of a capacity speculation SA is the limit price of its original SA, while the limit price of an on-hand speculation SA is its purchase price (or the limit price of the capacity speculation SA that it is converted from) plus the inventory cost of holding this unit since it was acquired (and thus its limit price increases over time).

In the following sections, we first introduce BA and SA algorithms and then detail the algorithms of the SPA, with a discussion of mechanisms for updating the profit margin and estimating the next period price.

Notations used in specifying the bidding agent algorithms include:

- B_m : the current bid of BA m ;
- O_n : the current offer of SA n ;
- τ_m : the target price of agent m ;
- Γ_m : the bid adjustment;
- l_m : the limit price of agent m ;
- γ_m : momentum coefficient;
- β_m : learning rate;
- r_m^1, r_m^2 : parameters in setting target prices;

Notations used in specifying the speculation agent algorithms include:

- λ_i : the required minimum profit margin of speculation;
- \mathcal{B}_i : the set of BAs of buyer i ;
- \mathcal{S}_j : the set of normal SAs of seller j ;
- ψ : the discount factor taking account of inventory costs.

B. Algorithms of Buying Agent and Selling Agent

Different bidding agent algorithms, such as ZI, ZIP and GD algorithms, may be employed in our framework. In our experiments, we use the modified ZIP agent algorithm in [6] due to its simplicity. For the completeness of the paper, we summarize this algorithm below.

The basic logic of this algorithm is the following. A buyer's goal is to buy products at the lowest price possible. It is rational to start bidding at a very low price and gradually increase the bid toward a target price until a trade is realized. This target price is set based on current market situation and is different from the limit price. Following this logic, the algorithm consists of two phrases: 1) determine the target price based on the current market situation and 2) adjust the bid

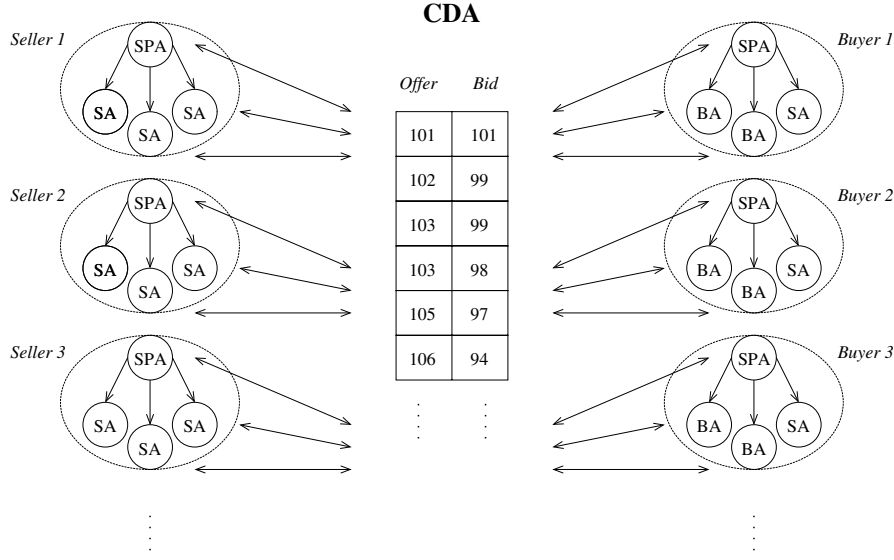


Fig. 1. The Agent-based Framework

towards the target price using learning rules, such as *Widrow-Hoff with momentum* (subject to the constraint that the bid should not exceed the limit price). The seller's situation is similar. Figure 2 details the algorithm. Please refer to [6] and [1] for details, such as how to decide the opening bids and offers.

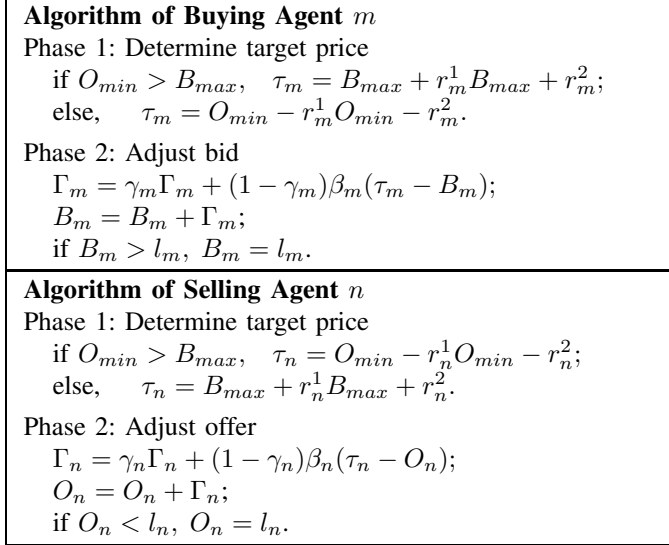


Fig. 2. Bidding Algorithms

C. Speculation Agent Algorithms

A SPA monitors price fluctuations and either 1) buys more speculation inventory when the current price is low and the expected next period price is high, 2) reduces speculation inventory by forcing a speculation SA to offer at B_{max} , or 3) does nothing. Making these decisions needs to follow certain economic rules, but also involves lots of uncertainty.

Therefore, the SPA algorithm is a combination of decision rules and autonomous learning algorithms. We first introduce four decision rules and the decision process. Based on this analysis, the SPA algorithm is proposed, and finally we discuss how to adjust the required profit margin and estimate the next period price.

1) *Decision Rules*: Making speculation decisions is a complicated process involving the information of the overall market as well as the buyer or seller herself. In this process, several decision rules are employed:

- Rule 1: speculation vs. normal trading. A trader considers buying for speculation only when she has no BA m such that $l_m \geq O_{min}$ or no SA n such that $l_n \leq B_{max}$. The rationale is that normal trading involves much less risk than speculation and, therefore, is preferred. Moreover, the speculation risk decreases as more information are available later in the period.
- Rule 2: buy or convert. To obtain speculation inventory, a seller prefers converting a normal SA to capacity speculation SA to buying a unit from the market as long as she has a normal SA with a limit price lower than O_{min} , due to the flexibility to convert the capacity speculation SA back.
- Rule 3: inventory cost. Due to the cost for holding inventory (mainly interest cost for most products), selling one unit of speculation inventory at price P_{t+1} in period $t+1$ is equivalent to selling this unit at ψP_{t+1} in period t , where $\psi < 1$. The inventory cost of carrying this unit for t to $t+1$ is $P_t(1/\psi - 1)$.
- Rule 4: profit margin. Trader i buys for speculation only if the expected profit margin is higher than the required profit margin λ_i .

2) *Decision Process and the SPA Algorithm*: With rules discussed above, buyer i 's SPA continually makes decisions

about buying and selling speculation inventory as follows. First, the SPA must compute its estimation of the next period price, noted by P_{t+1}^i . Then, it observes the market activities and buys or sells speculation inventory in the following two cases respectively:

- 1) Case 1: when $\psi P_{t+1}^i > O_{min}(1 + \lambda_i)$. In this case, holding speculation inventory is profitable, with an expected margin higher than λ_i . Buyer i 's SPA bids for one unit at O_{min} if buyer i has no BA with limit price higher than O_{min} ; otherwise it does not consider buying for speculation according to rule 1.
- 2) Case 2: when $\psi P_{t+1}^i < B_{max}$. In this case, unloading some speculation inventory is a logical decision. Buyer i 's SPA forces a speculation SA to offer at B_{max} ; however, if buyer i has BAs with limit price higher than B_{max} , it's more profitable to sell this speculation unit to one of such BAs.

The decision process of seller j 's SPA is similar but more complicated. With the estimated P_{t+1}^j , the SPA observes market activities and acts in two cases:

- 1) Case 1: when $\psi P_{t+1}^j > O_{min}(1 + \lambda_j)$. If seller j has normal SA with limit price lower than B_{max} , no speculation. Otherwise, convert a normal SA to capacity speculation SA if seller j has a normal SA with a limit price lower than O_{min} ; otherwise the SPA bids for one unit at O_{min} .
- 2) Case 2: when $\psi P_{t+1}^j < B_{max}$. If seller j has on-hand speculation inventory, this on-hand unit will be forced to be offered at price B_{max} on the market; otherwise, the capacity speculation SA with the highest limit price will be converted back to a normal SA.

Based on the analysis above, Figure 3 details the algorithms of the speculation agent.

3) *Adjust Required Profit Margin*: The profit margin λ_i is an important parameter in trader i 's decision making and needs to be adjusted according to the environment. On the one hand, a too large λ_i would waste chances to make profits. On the other hand, the risk of losing money on speculation is higher when λ_i is smaller. The problem of finding the λ_i that maximizes the expected speculation profit under current environment may be solved by various AI techniques. In this paper, we solve it as a gradient search problem. Figure 4 details the procedure to estimate the derivative of the expected speculation profit over λ_i . In this procedure, $P_{t+1}^i(\xi)$ and $O_{min}(\xi)$ are the P_{t+1}^i and O_{min} when ξ is bought, and $P(\xi)$ is the price at which ξ is sold; $|\mathcal{R}|$ gives the number of ξ in the set \mathcal{R} . Whenever a trader sold one unit of speculation inventory, her SPA will compute the derivative and adjust the λ_i by $\Delta\lambda_i$ which equals the derivative times a constant step size.

4) *Estimate The Next Period Price*: Trader i 's speculation decisions depend heavily on P_{t+1}^i , her estimation of the next period price. In a B2B exchange, the overall economic condition decides the market-wide supply and demand, and, therefore, decides the prices. Following the tradition in eco-

Algorithms of Buyer i 's Speculation Agent

- step1: compute i 's expected next period price P_{t+1}^i ;
find $\hat{m} \in \mathcal{B}_i$ with the highest limit price.
- step2: if $l_{\hat{m}} < O_{min}$ and $\psi P_{t+1}^i > O_{min}(1 + \lambda_i)$:
Bid at O_{min} . If success, create an on-hand speculation SA with limit price O_{min} ; otherwise cancel the bid;
- step3: if $\psi P_{t+1}^i < B_{max}$ and i has speculation inventory:
If $l_{\hat{m}} > B_{max}$, sell 1 unit speculation inventory to \hat{m} ;
otherwise, force a speculation SA to offer at B_{max} .

Algorithm of Seller j 's Speculation Agent

- step1: compute j 's expected next period price P_{t+1}^j ;
find $\hat{m} \in \mathcal{S}_j$ with the lowest limit price.
- step2: if $l_{\hat{m}} > B_{max}$ and $\psi P_{t+1}^j > O_{min}(1 + \lambda_j)$:
If $l_{\hat{m}} < O_{min}$, convert \hat{m} to capacity speculation SA.
Otherwise, bid at O_{min} ; if success, create an on-hand speculation SA with limit price O_{min} ; if not, cancel the bid.
- step3: if $\psi P_{t+1}^j < B_{max}$ and j has speculation inventory:
If has an on-hand speculation SA n , let n offer at B_{max} ;
Otherwise, convert the capacity speculation SA with the highest limit price to normal SA.

Fig. 3. Algorithms of Speculation Agents

When a unit ξ of i 's speculation inventory is sold:

- step 1: for this unit ξ , record $P_{t+1}^i(\xi)$, $O_{min}(\xi)$ and $P(\xi)$;
- step 2: initiate the step-size ϵ ; let \mathcal{N} be the N most recently sold ξ ;
- step 3: let \mathcal{R} be the set of ξ such that $\xi \in \mathcal{N}$, and $\lambda_i \leq \psi P_{t+1}^i(\xi)/O_{min}(\xi) - 1 \leq \lambda_i + \epsilon$;
- step 4: if $|\mathcal{R}|$ is less than the minimum requirement, double ϵ and go to step 3;
- step 5: derivative = $\frac{1}{|\mathcal{R}|\epsilon} \sum_{\xi \in \mathcal{R}} (P(\xi) - O_{min}(\xi))$.

Fig. 4. Procedure to estimate the derivative of the expected speculation profit over λ_i .

nomics, the overall economic condition may be modeled as an irreducible Markov process $\{a_t\}$ with stationary transition probabilities $\pi(a_{t+1}|a_t)$. The demand of all buyers, noted by \hat{D}_t , changes from period to period as random variables depending on state a_t .

Deferent techniques may be developed to estimate the next period price based on this economic model. In this paper, we consider three next-period estimation techniques.

- 1) Estimation using historic prices (HP): the weighted average of historic prices provides a good estimation of the next period price. In period k , let $\hat{P}_k^i(a)$ be the weighted average of prices of all past periods whose state is a . Then, in period t , $P_{t+1}^i = E_{a_t}[\hat{P}_t^i(a_{t+1})]$. The method used in our simulation to compute the weighted average price is called Exponential Smoothing, which updates $\hat{P}_k^i(a)$ as follows. In period k , if the state in k

is a , $\hat{P}_k^i(a) = \gamma P_k + (1 - \gamma)\hat{P}_{k-1}^i(a)$, where P_k is the price in period k ; otherwise $\hat{P}_k^i(a) = \hat{P}_{k-1}^i(a)$. γ is the smoothing constant.

- 2) Estimation using a price function (PF): P_{t+1}^i may be estimated by price function $f(D, S)$ as follows: first estimate the shortfall s_t and the speculation inventory I_t based on the current market condition; assume $I_{t+1} = 0$; then $P_{t+1}^i = E_{a_t}[f(\hat{D}_{t+1} + \alpha s_t, C + I_t)]$, where C is the overall production capacity of all suppliers in one period.
- 3) Estimation using an equilibrium function (EF): Li and Smith [5] have shown the existence of rational expectations equilibrium of price and market-wide speculation inventory. P_{t+1}^i may be also estimated by equilibrium functions $\bar{I}_t(I_{t-1}, a_t, D_t)$ and $\bar{P}_t(I_{t-1}, a_t, D_t)$ as follows: compute $\bar{I}_t(I_{t-1}, a_t, D_t)$ and $\bar{P}_t(I_{t-1}, a_t, D_t)$ by the procedure given in [5]; use known data I_{t-1}, a_t , and D_t to estimate I_t by function $\bar{I}_t(I_{t-1}, a_t, D_t)$; then $P_{t+1}^i = E_{a_t}[\bar{P}_{t+1}(I_t, a_{t+1}, D_{t+1})]$.

HP is easy to implement, but an obvious drawback is that the current market situation in period t is not utilized in the estimation. PF is a one-period look-ahead approach and, therefore, is myopic. Theoretically, EF is a better estimator compared with HP and PF. However, EF's performance depends on the market's ability to converge to the theoretical equilibrium. The performance of these three methods are compared through simulations in §IV-C.

IV. EXPERIMENTS

In this section we investigate the performance of the agent-based system discussed in §III under different business situations, such as economic change pattern, market size and industry structure. The setup of experiments is first described. Then we discuss different ways to measure the performance of a B2B exchange system. Finally the experiment results are presented.

A. Experiment Setup

All experiments are conducted on a B2B exchange system with the following features:

- *Economic condition.* 10 sellers and 10 buyers are trading in the exchange. The overall economic condition that decides demand is modeled as a Markov process with 3 states: bad, normal, and good. To simulate the common cyclic patterns of economic condition, we assume the Markov process behaves according to a state transition diagram given in Figure 5. The corresponding transition probabilities $\pi(a_{t+1}|a_t)$ is given in Table I. $x = 0.7$ if not specified. In each period, the demand of a buyer follows a truncated Normal distribution, *i.e.*, the positive portion of $z \times N(0.6, 0.15)$. Let $z = 5, 10, 20$ for states bad, normal and good respectively, and $N(0.6, 0.15)$ represents a Normal distribution with mean 0.6 and standard deviation 0.15. The shortfall carryover rate $\alpha = 0.3$. On the seller side, each seller is able to produce 8 units of products in each period. The discount factor ψ is 0.999.

- *CDA rule and parameters in agent algorithms.* In the CDA, whenever $B_{max} \geq O_{min}$, a trade occurs and the transaction price is $(B_{max} + O_{min})/2$. In our experiments, we use the modified ZIP agents in [6] as the BA and SA algorithms due to its simplicity, and parameters $r_m^1, r_m^2, \gamma_m, \beta_m$ are generated in the same fashion as in [6]. r_m^1 and r_m^2 are independently and identically distributed in the range $[0, 0.2]$; γ_m is uniformly distributed over $[0.2, 0.8]$ and β_m is uniformly distributed over $[0.1, 0.5]$. In each period, we assume that the limit price of each unit a buyer demands is uniformly distributed over $[50, 100]$; the limit price of each unit that a seller is able to produce in the current period is uniformly distributed over $[50, 100]$. When the demand is higher, there are more BAs with high limit prices in the market.
- *Simulation Parameters.* The simulation is carried on for 1000 periods with the first 50 periods as warm-up periods. In each period, there are 3000 moments. In every moment, each agent makes its decision based the current market situation. The parameter γ in HP is 0.3.




Fig. 5. The State Transition Diagram

TABLE I
TRANSITION PROBABILITY $\pi(a_{t+1}|a_t)$

| a_{t+1} | bad | normal | good |
|-----------|---------|---------|---------|
| bad | $1 - x$ | x | 0 |
| normal | $x/2$ | $1 - x$ | $x/2$ |
| good | 0 | x | $1 - x$ |

B. Measuring the Performance of a B2B exchange

Several different parameters may be used to evaluate the performance of a B2B exchange system, including:

- *Overall utility:* when one unit is traded, the buyer's utility equals her limit price minus the sale price; the seller's utility is the sale price minus the seller's limit price and the inventory cost if this unit is a speculation inventory. The total utility of all traders is the most important indicator of the efficiency of auctions at allocating resource.
- *Throughput:* the throughput of the system may be measured as the actual number of products that buyers actually buy for non-speculation reasons in a given time range. Since the throughput of the system decides the total number of final products that manufacturers (buyers)

are able to produce to meet the demand from consumers, it decides the overall profitability of all sellers and buyers.

- Price standard deviation (σ_{P_t}): for manufacturers (buyers), high price standard deviation means high uncertainty in supply and will cause higher costs in related production processes.

C. Simulation Results

1) *A B2B exchange with or without speculation:* We first compare the performance of two systems: a B2B exchange in which agents bid and speculate, against a B2B exchange in which agents only bid. Simulation results are reported in Table II, where NS represents the B2B exchange with no speculation allowed. *Utility* is the overall utility of all traders; \bar{P}_t and σ_{P_t} are the mean and standard deviation of price P_t over 950 periods; Number of all trades in 950 periods is reported in *# of Trades*; *Speculation Utility* is the utility generated by carrying and selling speculation inventory; in each period the speculation inventory is counted, and the sum over 950 periods is given by *Speculation Inventory*; the inventory holding cost is given by *Speculation Inv. Cost*. If one unit of speculation inventory is obtained in period t , the inventory cost of this unit is $P_t(1/\psi - 1)$ times the number of periods this unit has been carried. *Speculation Inv. Cost* gives the sum of inventory costs of all speculation inventory sold in 950 periods.

When agents make speculation decisions, the system performance is significantly improved in terms of overall utility, throughput and standard deviation. For instance, with EF as the next period price estimator, the overall utility rises 8.5%, the throughput increases 16.4% and the standard deviation of prices is reduced by 57%. The gain in the overall utility mainly comes from the speculation utility. Interestingly, this speculation utility is more than the increase of the overall utility. One explanation is that many products that otherwise would be sold instantly are held as speculation inventory instead by sellers. As expected, due to the speculation, trading price is reduced when it is high and boosted when low, and, therefore, the standard deviation σ_{P_t} is significantly reduced.

Another benefit of the speculation is better market liquidity; using EF, the trading volume is tripled. This result is intuitive and the market-maker could profit from it. Simulation results also suggest slight increase in the average trading price.

2) *Performance under different change patterns of economic conditions:* We further investigate the performance of HP, PF, and EF under different economic change patterns. According to the results shown in Table III, both PF and EF perform significantly better than HP in terms of all three performance measurements. The difference of PF and EF's performances is much smaller, but EF demonstrates better ability to estimate next period prices. Overall, compared with the case without speculation, all three next period estimation techniques are able to help improving the overall performance of B2B exchanges. Even when the cyclic pattern in economic changes is weak, *i.e.* when $x = 0.5$, HP, PF and EF increase the overall utility by 4.3%, 5.8% and 7.9% respectively. When the economic changes show stronger cyclic pattern with $x =$

0.9, this performance improvements are boosted to 4.4%, 6.7% and 10.1%.

TABLE III
UNDER DIFFERENT ECONOMIC CONDITIONS

| | Utility | Specu. Ut. | Throughput | σ_{P_t} |
|-----------|----------|------------|------------|----------------|
| $x = 0.5$ | | | | |
| NS | 739267.2 | 0 | 34004 | 11.8 |
| HP | 771432.6 | 130380.5 | 36568 | 7.6 |
| PF | 782340.9 | 143109.1 | 37449 | 7.0 |
| EF | 798082.6 | 195809.4 | 37437 | 5.3 |
| $x = 0.7$ | | | | |
| NS | 744739.7 | 0 | 34245 | 11.4 |
| HP | 777042.2 | 129989.6 | 36792 | 7.3 |
| PF | 794858.6 | 143600.4 | 38121 | 6.6 |
| EF | 807671.2 | 201784.3 | 39876 | 4.9 |
| $x = 0.9$ | | | | |
| NS | 747097.6 | 0 | 34418 | 11.1 |
| HP | 779904.3 | 135790.4 | 36959 | 7.2 |
| PF | 797609.6 | 134553.4 | 37929 | 6.1 |
| EF | 822493.3 | 207449.4 | 39938 | 4.4 |

3) *Performance in markets of different sizes:* In this section, we study the influence of market size on the performance of the B2B exchange system. As shown in Table IV, the performance of HP, PF and EF is robust when the market size changes.

TABLE IV
PERFORMANCE IN MARKETS IN DIFFERENT SIZES

| | Utility | Specu. Ut. | Throughput | σ_{P_t} |
|-----------------------|-----------|------------|------------|----------------|
| 10 buyers, 10 sellers | | | | |
| NS | 744739.7 | 0 | 34245 | 11.4 |
| HP | 777042.2 | 129989.6 | 36792 | 7.3 |
| PF | 794858.6 | 143600.4 | 38121 | 6.6 |
| EF | 807671.2 | 201784.3 | 39876 | 4.9 |
| 20 buyers, 20 sellers | | | | |
| NS | 1515059.7 | 0 | 72682 | 11.4 |
| HP | 1561642.0 | 255714.5 | 74264 | 7.1 |
| PF | 1618485.0 | 282399.9 | 76206 | 6.5 |
| EF | 1624393.4 | 362539.3 | 78614 | 5.0 |

4) *Negotiation power:* In a buyer-seller negotiation (bidding) process, one side, either buyers or sellers, may have the negotiation power over the other side. For instance, in the automobile industry buyers like General Motors have the negotiation power over their suppliers in most cases. This negotiation power depends on many factors such as company sizes and the competition. It is an important aspect of the structure of an industry, and in B2B exchanges, it has direct impact on each player's bidding behavior and the outcome of CDAs.

We will simulate the negotiation power and study its influence on the performance of our agent-based system. Generally, when sellers have the negotiation power, they are reluctant to reduce price for a large quantity, and, therefore, the slope of the supply curve is smaller compared to other cases. To achieve this effect, we reduce the range over which sellers' limit prices are uniformly distributed. In the Table V, for the "Balanced" case, both the buyer's range and the seller's range is [50,100]. In the "Seller Power" and "Strong Seller Power" cases, sellers are willing to sell products at higher prices, and, therefore, the

TABLE II
B2B EXCHANGE WITH AND W/O SPECULATION

| | NS | HP | PF | EF |
|-----------------------|-----------|-----------|-----------|-----------|
| Utility | 744739.67 | 777042.16 | 794858.58 | 807671.24 |
| Throughput | 34245 | 36792 | 38121 | 39876 |
| σ_{P_t} | 11.35 | 7.31 | 6.60 | 4.86 |
| Average P_t | 71.946251 | 73.2 | 74.50 | 74.11 |
| # of Trades | 34245 | 64019 | 90341 | 109523 |
| Speculation Utility | 0 | 129989.55 | 143600.35 | 201784.29 |
| Speculation Inventory | 0 | 3409 | 4732 | 15041 |
| Speculation Inv. Cost | 0 | 413.02 | 585.13 | 1913.28 |

seller's range is changed to [65,100] and [80,100] respectively. In the "Buyer Power" and "Strong Buyer Power" cases, buyers are willing to lower prices, and the buyer's range is changed to [50,85] and [50,70] respectively. The experiments show that, as either the buyers or the sellers gain negotiation power, the overall utility of the system is hampered. Even with using estimator EF, the utility drops 76% both in the Strong Buyer Power case and the Strong Seller Power case. Similarly, the benefit of allowing speculation is also limited due to the buyer or seller power. In the Strong Sell Power and Strong Buyer Power cases, the performance improvements of EF decrease to 7.3% and 5.6% respectively from 8.5%.

TABLE V

THE OVERALL UTILITY UNDER DIFFERENT NEGOTIATION POWERS

| | NS | HP | PF | EF |
|---------------------|----------|----------|----------|----------|
| Strong Seller Power | 179442.0 | 182840.8 | 193275.4 | 192578.4 |
| Seller Power | 435118.1 | 449473.5 | 463738.7 | 472394.3 |
| Balanced | 744739.7 | 777042.2 | 794858.6 | 807671.2 |
| Buyer Power | 434342.6 | 455014.3 | 466047.7 | 477237.5 |
| Strong Buyer Power | 181322.5 | 185345.7 | 186212.9 | 191567.6 |

5) *Performance with different inventory costs:* An important issue to consider in speculation is the inventory cost. When this cost is high, holding speculation inventory will be unprofitable in many situations, and, therefore, the overall speculation in the market will be less active. As shown in Table VI, when the inventory cost increase from 0.1% per period ($\psi = 0.999$) to 1% per period ($\psi = 0.99$) and 10% per period ($\psi = 0.91$), in the EF case, the total amount of speculation inventory decreases from 15041 to 9781 and 776 respectively. At the same time, the performance improvement due to speculation is also hampered. For instance, the Utility improvement decreases from 8.5% to 6.5% and 1.5% respectively.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented an agent-based framework for make bidding and speculating decisions in B2B exchanges. We analyzed the continuous double auction in the B2B exchange context, discussed the decision process of speculation in detail, and developed algorithms to make speculation decisions, adjust required margin, and estimate the next period price. Simulation results suggest a significant performance improvement of the B2B exchange system when traders (represented as agent-based systems) are able to make

TABLE VI
PERFORMANCE WITH DIFFERENT INVENTORY COSTS

| | Utility | Specu. Ut. | Specu. Inv. | Throughput | σ_{P_t} |
|----------------|----------|------------|-------------|------------|----------------|
| $\psi = 0.999$ | | | | | |
| NS | 744739.7 | 0 | 0 | 34245 | 11.4 |
| HP | 777042.2 | 129989.6 | 3409 | 36792 | 7.3 |
| PF | 794858.6 | 143600.4 | 4732 | 38121 | 6.6 |
| EF | 807671.2 | 201784.3 | 15041 | 39876 | 4.9 |
| $\psi = 0.99$ | | | | | |
| NS | 744739.7 | 0 | 0 | 34245 | 11.4 |
| HP | 770359.5 | 107186.0 | 2127 | 36845 | 7.8 |
| PF | 785355.1 | 126960.1 | 4075 | 38077 | 7.0 |
| EF | 793167.7 | 168330.3 | 9781 | 38523 | 6.0 |
| $\psi = 0.91$ | | | | | |
| NS | 744739.7 | 0 | 0 | 34245 | 11.4 |
| HP | 749178.8 | 15247.3 | 403 | 36532 | 10.0 |
| PF | 755208.4 | 20783.4 | 772 | 37663 | 9.3 |
| EF | 756259.9 | 21944.6 | 776 | 38154 | 9.3 |

speculation decisions. We also proposed and discussed several next period estimation techniques. Our work extends the preview research on the design of bidding agents to more dynamic and realistic situations in B2B exchanges.

Although the agent-based framework and related algorithms are presented in the context of continuous double auctions, the framework and ideas may apply to other market institutions. Developing more efficient algorithms to adjust required profit margins and estimate next period prices is one direction of our future research. For instance, an artificial neural network might be used to estimate the next period price.

ACKNOWLEDGMENTS

The research reported in this paper has been funded in part by the Department of Defense Advanced Research Projects Agency under contract F30602-00-2-0503 and by the Robotics Institute at Carnegie Mellon University. The views and conclusions of the paper should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or of the U.S. Government.

REFERENCES

- [1] D. Cliff. and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, Hewlett Packard Labs, 1997.
- [2] eMarketer. The e-commerce trade and b2b exchanges report. Technical report, March 2002.
- [3] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22.

- [4] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137.
- [5] L. Li and S. F. Smith. A quantitative analysis of business-to-business spot markets. Working paper, Carnegie Mellon University, 2002.
- [6] C. Preist and M. van Tol. Adaptive agents in a persistent shout double auction. In *Proceedings of ICE-98*, pages 11–18. ACM.
- [7] V. L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111–137.
- [8] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *ACM Conference on Electronic Commerce (EC'01)*, pages 206–209. ACM.