

A Fuzzy-Logic Based Bidding Strategy for Autonomous Agents in Continuous Double Auctions

Minghua He, Ho-fung Leung, and Nicholas R. Jennings

Abstract—Increasingly, many systems are being conceptualized, designed, and implemented as marketplaces in which autonomous software entities (agents) trade services. These services can be commodities in e-commerce applications or data and knowledge services in information economies. In many of these cases, there are both multiple agents that are looking to procure services and multiple agents that are looking to sell services at any one time. Such marketplaces are termed continuous double auctions (CDAs). Against this background, this paper develops new algorithms that buyer and seller agents can use to participate in CDAs. These algorithms employ heuristic fuzzy rules and fuzzy reasoning mechanisms in order to determine the best bid to make given the state of the marketplace. Moreover, we show how an agent can dynamically adjust its bidding behavior to respond effectively to changes in the supply and demand in the marketplace. We then show, by empirical evaluations, how our agents outperform four of the most prominent algorithms previously developed for CDAs (several of which have been shown to outperform human bidders in experimental studies).

Index Terms—Intelligent agents, service marketplaces, continuous double auction, fuzzy logic, e-commerce.

1 INTRODUCTION

THE advent of global network structures, such as the Internet, has facilitated the development of many large-scale, open distributed systems in a wide range of industrial, commercial, and educational domains. In many cases, these systems can be viewed using a service-oriented metaphor in which various entities offer services to one another in some form of marketplace [1], [2]. For example, in deregulated electricity markets, power generators compete with one another to provide the service of supplying electricity for consumers [3], in digital libraries, information services aim to discover relevant content from providers who offer a variety of document and archive services [4], and in grid computing, high performance applications seek to procure the necessary computational resource services to run [5]. In all of these applications, and in many others besides, the marketplace in which the service producers and the service consumers interact is some form of online auction.¹ The reason for this is that auctions are a very efficient and effective method of allocating goods/services, in dynamic situations, to the entities that value them most highly [6]. While there are many different types of auction [7], the most common forms

are the simple single sided varieties (e.g., English, first-price ascending; Dutch, first-price descending; First-Price-Sealed Bid and Vickrey, second-price sealed-bid) in which there is a single seller and multiple buyers or a single buyer and multiple sellers (a reverse auction). However, in many applications, these simple auction protocols² are inadequate because there are multiple sellers and multiple buyers that want to trade simultaneously. This can occur, for example, because entities may want to resell services they have procured or because delivery of a service requires the provider to procure subsequent component subservices from others. Such auctions are called double auctions [9] and they allow sellers to indicate the services they offer at various prices (called asks) and buyers to indicate the services they desire and the price they are willing to pay (called bids). The most common variety of double auction is the continuous double auction (CDA)³ which permits trade at any time in a

2. In economic and game theory, interactions consist of two components: a protocol and a strategy [8]. The former defines the valid behavior of the agents during the interaction (e.g., who can say what to whom at what time). The later is the method the agents employ to achieve their negotiation objectives *within the specified protocol*. The protocol is set at design time by the marketplace owner and is publicly known to all the participants. The strategy is designed by each individual participant and is private (divulging it may leave them open to exploitation). Moreover, the effectiveness of the strategy is very much determined by the protocol; an optimal strategy for one protocol may well perform very poorly for other protocols.

3. Although CDAs all conform to this basic protocol, there are several ways in which they can differ. These variants relate to factors such as whether the identity of bidders is revealed, whether bids and asks are for single or multiple units, and whether unaccepted offers are queued or replaced by better offers [10]. In our case, we do not reveal the identity of bidders, we allow bids for single units only (one outstanding bid and one outstanding ask), and unaccepted offers are erased as soon as there is a more favorable bid (ask). The most restrictive of these assumptions is the fact that we deal with single unit trades (as do most of the algorithms against which we benchmark our agent's performance). However, even with this restriction in place, the CDA has shown to be a highly efficient protocol [11].

1. In the domain of e-commerce, for example, it is estimated that there are currently more than 2,000 auction sites on the Internet (<http://www.internetauctionlist.com>).

- M. He and N.R. Jennings are with the Department of Electronics and Computer Science, The University of Southampton, Southampton, United Kingdom. E-mail: {mh00r, nrj}@ecs.soton.ac.uk.
- H.-f. Leung is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong SAR, P.R. China. E-mail: lhf@cse.cuhk.edu.hk.

Manuscript received 18 May 2001; revised 27 Nov. 2001; accepted 12 Apr. 2002.

For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number 114173.

trading period (cf. trades only being allowed at discrete time points) and which allows buyers and sellers to continuously update their bids and asks at any time throughout the trading period [12]. CDAs are widely used in the nononline world to trade stocks, agricultural commodities, metals, currencies, and derivative instruments [10].

Against this background, this paper develops and evaluates a new algorithm that autonomous software agents [13] can employ to submit bids and asks in CDAs. We believe this is an important step in the development of large-scale open and distributed systems because the existence of effective strategies means that CDAs can be more readily deployed as the marketplace protocol. Without such strategies, there has been some reluctance to choose CDAs in the online world even though they are the most obvious protocol in many cases. Furthermore, we believe that, in the long term, software agents will be more effective than human bidders in these more complex auction settings. Preliminary evidence for this is contained in [14] which shows that agents, employing the algorithms against which we benchmarked our algorithm, outperformed their human counterparts in CDAs.

In more detail, the bidding algorithms we develop are heuristic methods that exploit fuzzy logic techniques [15], especially fuzzy rules, to undertake their reasoning. The reason for this choice is that, in CDAs, there is no optimal bidding strategy [9]. This is because an agent's decision making about bidding involves uncertainty, multiple factors, and nondeterminism that are affected by the attitudes toward risk of its opponents, the nature of the market supply (demand), and the preferences of the other bidders. Since no agent can have all this information in advance (it is, after all, a competitive environment), the best that can be achieved is a satisfying strategy [16]. We chose to adopt a fuzzy logic-based approach, in particular, because we wish to develop a practicable agent that can cope with the uncertainties in a timely manner and fuzzy techniques have proven to be successful in a wide range of domains with these characteristics (e.g., fuzzy control to drive car-like vehicles [17], making medical diagnosis [18], vehicle dispatching [19] and emergency electric power distribution [20]). See Section 6 for details of the other alternatives we considered.

The specific contributions of this paper are as follows: First, we develop a novel fuzzy logic-based bidding strategy—the FL strategy—for agents that participate in CDAs. Second, we present the design, implementation, and evaluation of this strategy for buyer and seller agents. This strategy is shown, via empirical studies, to outperform the main strategies that have previously been proposed for CDAs. Third, we enhance the basic strategy so that it can adapt its behavior to the supply (demand) in the market (this revised strategy is called the adaptive FL-strategy). We then show how this revised strategy leads to a further improvement in the performance of both the individual agents (buyers and sellers) and of the overall marketplace.

The remainder of this paper is organized as follows: Section 2 formalizes a CDA and outlines the basics of our fuzzy reasoning mechanism. Section 3 presents the

FL-strategy. In Section 4, the behaviors of our FL-agents are analyzed in a range of experiments. Section 5 discusses the adaptive FL-agents and their evaluation. Section 6 discusses the related work. Finally, Section 7 concludes this paper and discusses the future work.

2 PRELIMINARIES

This section outlines the basis of our FL-strategy—presenting a formal account of our CDA protocol and describing the fuzzy reasoning mechanism we employ.

2.1 Continuous Double Auctions

According to the parameterization of CDAs given in [10], we deal with the situation in which there are more than two goods in the market; two-way traders and the numbers of buyers and sellers are greater than three; single indivisible units are to be traded (thus, at any one time, there is one outstanding bid and one outstanding ask); the preferences of the traders are the reservation prices of the goods; and traders have incomplete information of the market. The CDA terminates after a specified period of inactivity.

In more detail, there are agents that are willing to sell goods (s-agents) and agents that are willing to buy goods (b-agents). A given agent can be either a buyer or a seller in a given context. Specifically, an *ask* a is the amount submitted by an s-agent willing to sell a unit of good. The lowest ask in the market is called the *outstanding ask*, denoted a_o . Similarly, a *bid* b is the amount submitted by a b-agent willing to buy a unit of good. The highest bid in the market is called the *outstanding bid*, denoted b_o . A CDA can thus be described as a place where s-agents submit asks to decrease a_o , while b-agents submit bids to increase b_o , until b_o is not less than a_o [11]. At this moment, the s-agent that submits a_o and the b-agent that submits b_o can make a transaction, and the price of the transaction is called the transaction price. Formally, we have:

Definition 1. *The descriptor of a CDA is*

$$P_{CDA} = \langle g, \mathcal{B}, \mathcal{S}, V_b, C_s, \Delta_{price}, t_{round} \rangle,$$

where:

1. g is the good to be auctioned.
2. $\mathcal{B} = \{b_1, \dots, b_n\}$ is the finite set of identifiers of b-agents, where n is the number of b-agents.
3. $\mathcal{S} = \{s_1, \dots, s_m\}$ is the finite set of identifiers of s-agents, where m is the number of s-agents.
4. $V_b = (\vec{V}_1, \dots, \vec{V}_n)$, where $\vec{V}_i(v_{i1}, v_{i2}, \dots, v_{in_i})$ is a vector of unit valuations of b-agent b_i . Here, n_i is the number of units of g that b_i requires, and v_{ij} is the valuation value for the j th unit acquired.
5. $C_s = (\vec{C}_1, \dots, \vec{C}_m)$, where $\vec{C}_i(c_{i1}, \dots, c_{im_i})$ is a vector of unit costs of s-agent s_i . Here, m_i is the number of units that s_i wants to sell, and c_{ij} is the cost of the j th unit.
6. Δ_{price} is the minimum price step required in the auction. That is, a b-agent (s-agent) must increase (decrease) its bid (ask) at $n \times \Delta_{price}$, where n is a nonnegative integer.

7. t_{round} is used for defining the condition for terminating the CDA; that is, if there are no new asks or bids during a time period t_{round} , the CDA terminates.⁴

Definition 2. A *round* in a CDA is the time period between two successive deals or the period from the beginning of the CDA to the time when the first deal takes place. If a round is the r th ($r \in \mathbb{N}^+$) round of the CDA, then r is called the *round number*. A CDA usually consists of multiple rounds.

Definition 3. For a CDA that has lasted r ($r > 0$) rounds, let p_i ($1 \leq i \leq r$) denote the price of the i th transaction. A *history* H_l in a CDA is the set of transaction prices during the last l rounds,

$$H_l = \{p_{r-l+1}, \dots, p_i, \dots, p_r\},$$

where p_i ($r-l+1 \leq i \leq r$) is the transaction price of round i , and l ($l \leq r$) is called the *history length*.⁵

The following is the formal definition of the valid behaviors of agents during a CDA.

Definition 4. A CDA protocol with the descriptor P_{CDA} consists of the following steps:

1. $r=0$.
2. A new round of the CDA starts, $r = r + 1$, $a_o = \infty$, and $b_o = 0$.
3. Several situations might arise during a round:
 - a. When an s -agent submits an ask a ,
 - i. if $a \geq a_o$ then a is an invalid ask;
 - ii. if $b_o < a < a_o$, then a_o is updated to a ;
 - iii. if $a \leq b_o$, then this s -agent makes a deal at b_o ; goto 2.
 - b. When a b -agent submits a bid of b ,
 - i. if $b \leq b_o$, then b is an invalid bid;
 - ii. if $b_o < b < a_o$, then b_o is updated to b ;
 - iii. if $b \geq a_o$, then this b -agent makes a deal at a_o ; goto 2.
4. Step 3 repeats until no new bids (asks) are submitted during a time period t_{round} .

As can be seen, the outstanding ask and outstanding bid define the *bid-ask spread* $[b_o, a_o]$ [11] and only bids and asks that fall within this region are considered valid.

2.2 Fuzzy Reasoning Mechanisms

The fuzzy reasoning inference mechanism employed in this paper is based on the Sugeno controllers [21], [22]. Consider the following block of fuzzy IF-THEN rules:

4. Note that, by this definition, we exclude from this paper CDAs that last infinite periods of time (such as stock markets). To model this, t_{round} can be set to infinity.

5. Through experiments where both the history length (l) and the value (cost) of the goods that the agents trade varied, the performance of the agents with different history lengths was investigated. The results showed that the behavior of FL-agents with a long history length ($l > 20$) was similar to or worse than that of an agent with a history length ranging from 3 to 20. This result shows that agents with short or intermediate history lengths can react more rapidly to changes in a CDA market. When the history length varied from 3 to 20, we found that 10 was a reasonable history length where almost all the agents achieve their highest profit. Thus, this is the value selected for all the experiments in the rest of this paper.

\mathcal{R}_1 : if x is A_1 and y is B_1 then $z_1 = c_1$
also
 \mathcal{R}_2 : if x is A_2 and y is B_2 then $z_2 = c_2$
also
 \vdots
also
 \mathcal{R}_n : if x is A_n and y is B_n then $z_n = c_n$
fact : x is x_0 and y is y_0
consequence : z_0 ,

where A_1, \dots, A_n and B_1, \dots, B_n are fuzzy sets, and z_1, \dots, z_n are real numbers. The firing level α_i of the rules \mathcal{R}_i is computed by the *Min* operator. That is,

$$\alpha_i = \min\{A_i(x_0), B_i(y_0)\}, \quad (1)$$

where $A_i(x)$ and $B_i(y)$ are the membership functions of the corresponding fuzzy sets A_i and B_i , respectively. If the output of the individual rule is denoted as z_i , then according to the Sugeno controller definition, the crisp control action of the rule base is obtained by:

$$z_0 = \frac{\sum_{i=1}^n \alpha_i z_i}{\sum_{i=1}^n \alpha_i}. \quad (2)$$

The extension principle [15] is one of the main means of fuzzifying a formula with crisply defined numbers. In particular, we extend (2) to the situation where these real numbers z_i ($1 \leq i \leq n$) are changed to triangular fuzzy numbers. We made this change because in developing our rules, we felt unable to estimate the action using a single real value chosen from within a predefined range. Rather, we found it easier to estimate a parameter with fuzzy values and this led us to use triangular fuzzy numbers [23]. Also, by the extension principle, arithmetic operations on trapezoidal fuzzy numbers have already been obtained [24], [25] and fuzzy triangular numbers are special cases of fuzzy trapezoidal numbers [26]. Thus, the arithmetic operations on fuzzy triangle numbers can be obtained from the arithmetic operations on fuzzy trapezoidal numbers. Given all this, in our inference mechanism, the output of each rule is a triangular fuzzy number defined with the following triple:

$$\tilde{a} = (m, \theta, \chi),$$

where m is called the center, and θ and χ are called the left and right spreads, respectively [27] (Fig. 1). For two triangular fuzzy numbers $\tilde{a}_1 = (m_1, \theta_1, \chi_1)$, and $\tilde{a}_2 = (m_2, \theta_2, \chi_2)$ ($\tilde{a}_1, \tilde{a}_2 > 0$) and $k \in \mathbb{R}$, the following formulae hold [24], [25]:

$$\begin{aligned} \tilde{a}_1 + \tilde{a}_2 &= (m_1 + m_2, \theta_1 + \theta_2, \chi_1 + \chi_2), \\ \tilde{a}_1 - \tilde{a}_2 &= (m_1 - m_2, \theta_1 + \chi_2, \theta_2 + \chi_1), \\ \tilde{a}_1 \times \tilde{a}_2 &= (m_1 m_2, m_1 \theta_2 + m_2 \theta_1 - \theta_1 \theta_2, m_1 \chi_2 + m_2 \chi_1 + \chi_1 \chi_2), \\ k \times \tilde{a}_1 &= (k m_1, k \theta_1, k \chi_1). \end{aligned}$$

From the above formulae, (2) can be extended to the following in the situation where $\tilde{z}_i = (m_i, \theta_i, \chi_i)$ ($1 \leq i \leq n$):




Fig. 1. Triangular fuzzy number $\tilde{a} = (m, \theta, \chi)$, where m is the center, θ is left spread, and χ is right spread.

$$\begin{aligned} \tilde{z}_0 &= \frac{\sum_{i=1}^n (\alpha_i \times \tilde{z}_i)}{\sum_{i=1}^n \alpha_i} = \frac{\sum_{i=1}^n (\alpha_i \times (m_i, \theta_i, \chi_i))}{\sum_{i=1}^n \alpha_i} \\ &= \left(\frac{\sum_{i=1}^n (\alpha_i \times m_i)}{\sum_{i=1}^n \alpha_i}, \frac{\sum_{i=1}^n (\alpha_i \times \theta_i)}{\sum_{i=1}^n \alpha_i}, \frac{\sum_{i=1}^n (\alpha_i \times \chi_i)}{\sum_{i=1}^n \alpha_i} \right). \end{aligned} \quad (3)$$

Thus, the reasoning mechanism becomes:

\mathcal{R}_1 : if x is A_1 and y is B_1 then \tilde{z}_1 is \tilde{c}_1
also
 \mathcal{R}_2 : if x is A_2 and y is B_2 then \tilde{z}_2 is \tilde{c}_2
also
 \vdots
also
 \mathcal{R}_n : if x is A_n and y is B_n then \tilde{z}_n is \tilde{c}_n
fact : x is x_0 and y is y_0
consequence : \tilde{z}_0

Having defined the protocol and the reasoning mechanism, we can now turn to the FL-strategy itself.

3 THE FL-STRATEGY

Building on the foundations of the previous section, this section describes our FL-strategy and demonstrates how it works in an exemplar scenario.

3.1 Basic Notation and Concepts

In order to detail the FL-strategy, we first need to introduce a number of underpinning notations and concepts.

Definition 5. A *situation* s^* during the course of a CDA is a 6-tuple,

$$s^* = \langle r, \mathcal{B}, \mathcal{S}, a_o, b_o, H_l \rangle,$$

where r is the current round number, \mathcal{B} and \mathcal{S} are the sets of b -agents and s -agents; a_o and b_o are the outstanding ask and the outstanding bid, respectively, and H_l is the history of the last l rounds.⁶

Definition 6. Given a situation s^* , the *valid bids set* (D_b) is the set of the valid bids that a b -agent could submit:

$$D_b = \{b \mid b_o < b \leq \min(a_o, v_{ij})\}, \quad (4)$$

where b is the price at which a b -agent submits a bid and v_{ij} is the valuation of the j th unit of the good by buyer i .

Definition 7. Given a situation s^* , the *valid asks set* (D_s) is the set of valid asks that an s -agent could submit:

$$D_s = \{a \mid \max(b_o, c_{ij}) \leq a < a_o\}, \quad (5)$$

6. Recall that l is the remembered history length of an agent and, thus, l is not necessarily equal to $r - 1$.

where a is the price at which an s -agent submits an ask and c_{ij} is the cost of the j th unit of the good for seller i .

The prices of previous transactions are stored as history and may be referred to by the agents in the subsequent rounds. Generally speaking, CDA markets produce very efficient allocations and prices [28], and the transaction prices often converge to a competitive equilibrium price⁷ while the CDA is in progress. Thus, the transaction prices in a CDA provide an important point for reference. To reflect this fact, we define the **reference price** P_R in the situation s^* as the median of the ordered price history.⁸ A reference price, as its name suggests, provides a reference point that an agent can use to guide its subsequent bidding behavior. Formally, we have:

Definition 8. Let r be the current round number ($r > 0$).

Suppose the price history is a series of prices

$$H_l = \{p_{r-l}, \dots, p_i, \dots, p_{r-1}\},$$

where p_i ($r - l \leq i \leq r - 1$) is the price in round i . Let their ordered series be denoted as

$$p_{(1)} \leq \dots \leq p_{(i)} \leq \dots \leq p_{(l)}. \quad (6)$$

Then, the reference price, P_R , is given by

$$P_R = p_{(\lfloor \frac{l+1}{2} \rfloor)}. \quad (7)$$

To summarize, when an agent submits its next ask (bid), it will consider the outstanding ask, the outstanding bid, the cost (valuation) of the current unit of good, and the reference price. The way in which these values are used is described in the next section.

3.2 Fuzzy Reasoning in the FL-Strategy

The FL-strategy is based on a number of heuristic rules and the fuzzy reasoning mechanism outlined in Section 2.2. The relation of P_R , a_o , and b_o during a round in a CDA falls into one of the cases below:

1. $P_R \leq b_o < a_o$,
2. $b_o < a_o \leq P_R$, and
3. $b_o \leq P_R \leq a_o$.

In the first two cases, we use some heuristic rules (given below); the bidding issue in the third case, which is more complicated, is handled through the fuzzy reasoning mechanism on a rule base (described at the end of this section). Fig. 2 describes all the fuzzy sets used in the heuristic rules. The heuristic rules applied in the first two cases for s -agents are:

- When $P_R \leq b_o < a_o$, the heuristic rule is:

7. The equilibrium price is determined by the intersection of the supply and demand curves of the market, and it is the point where the quantity supplied is equal to the quantity demanded [29].

8. Originally, both the mean and the median of the ordered price history were used; however, experimental results showed that the median is more effective in providing a reference price. This is because the mean price can be overly influenced by a too high (low) price offered by an irrational agent. In contrast, the median of the ordered price history is less susceptible to such bias.

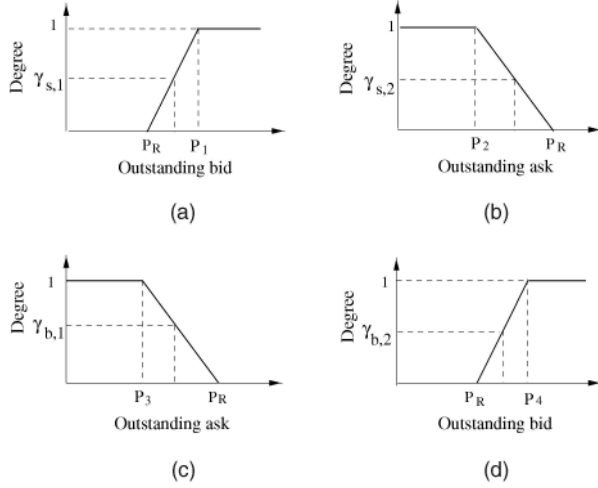


Fig. 2. Fuzzy sets in heuristic rules. (a) Outstanding bid is much_bigger than P_R . (b) Outstanding ask is much_smaller than P_R . (c) Outstanding ask is much_smaller than P_R . (d) Outstanding bid is much_bigger than P_R .

(SR₁) IF b_o is *much_bigger* than P_R
 THEN accept b_o
 ELSE ask is $(a_o - \beta_{s,1}, \theta, \chi)$.

- When $b_o < a_o \leq P_R$, the heuristic rule is:

(SR₂) IF a_o is *much_smaller* than P_R
 THEN no new ask
 ELSE ask is $(a_o - \beta_{s,2}, \theta, \chi)$.

Intuitively, SR_1 states that when the outstanding bid b_o is *much_bigger* than the reference price P_R , it is already very profitable for an s-agent to accept the current outstanding bid. The relation " b_o is *much_bigger* than P_R " can be expressed as fuzzy set A_1 . Let the threshold be $\gamma_{s,1}$, that is, if $A_1(b_o) \geq \gamma_{s,1}$, the s-agent will accept b_o . At this point, a transaction takes place between the s-agent and the b-agent which submits the outstanding bid. Otherwise, the s-agent will decrease the outstanding ask a_o to a fuzzy number $(a_o - \beta_{s,1}, \theta, \chi)$ (see Section 2.2), where $a_o - \beta_{s,1}$ is the center of the new ask, and θ and χ are the left and right spread. $\beta_{s,1}$ shows how much the agent would like to decrease its ask and this is decided by the agents' attitude to risk (to be discussed in Section 4.2). SR_2 is applied when a_o is much smaller than P_R . At this moment, an s-agent is in an unfavorable position and it should be reluctant to decrease a_o . Thus, the s-agent only decreases a_o by a small step. The relationship " a_o is *much_smaller* than P_R " is expressed as a fuzzy set A_2 . Let $\gamma_{s,2}$ be the threshold, that is, if $A_2(a_o) \geq \gamma_{s,2}$, the agent believes the current ask is much smaller than P_R . In this case, the s-agent will not submit a new ask.

Similar heuristic rules also apply to b-agents:

- When $b_o < a_o \leq P_R$, the heuristic rule is:

(BR₁) IF a_o is *much_smaller* than P_R
 THEN accept a_o
 ELSE bid is $(b_o + \beta_{b,1}, \theta, \chi)$.

- When $P_R \leq b_o < a_o$, the heuristic rule is:

TABLE 1
 Fuzzy Rule Base for s-agents

IF	$(b_o$ is <i>far_from</i> or <i>medium_to</i> P_R) and $(a_o$ is <i>far_from</i> P_R)
THEN	ask is $(a_o - \lambda_{s,1}, \theta, \chi)$.
IF	$(b_o$ is <i>far_from</i> or <i>medium_to</i> P_R) and $(a_o$ is <i>medium_to</i> P_R)
THEN	ask is $(a_o - \lambda_{s,2}, \theta, \chi)$.
IF	$(b_o$ is <i>far_from</i> or <i>medium_to</i> P_R) and $(a_o$ is <i>close_to</i> P_R)
THEN	ask is $(a_o - \lambda_{s,3}, \theta, \chi)$.
IF	b_o is <i>close_to</i> P_R
THEN	ask is $(P_R + \lambda_{s,4}, \theta, \chi)$.

(BR₂) IF b_o is *much_bigger* than P_R
 THEN no new bid
 ELSE bid is $(b_o + \beta_{b,2}, \theta, \chi)$.

The relationship " a_o is *much_smaller* than P_R " can be expressed as a fuzzy set A_3 . Let $\gamma_{b,1}$ be the threshold, that is, if $A_3(a_o) \geq \gamma_{b,1}$, a_o is regarded as being much smaller than P_R , and a b-agent will accept a_o ; otherwise, a b-agent will increase b_o to a fuzzy number $(b_o + \beta_{b,1}, \theta, \chi)$. The fuzzy set A_4 defines the relationship " b_o is *much_bigger* than P_R ." Let $\gamma_{b,2}$ be the threshold for this rule, that is, if $A_4(b_o) \geq \gamma_{b,2}$, a b-agent will not submit a new bid because b_o is already high enough and no profit can be made according to its preference; otherwise, it will increase b_o to a fuzzy number $(b_o + \beta_{b,2}, \theta, \chi)$. In the above, P_1 , P_2 , P_3 , and P_4 are the parameters of the fuzzy sets (see Fig. 2) and they are decided by human intuition and experience according to the range of the cost and valuation of the goods. The fuzzy number produced by these heuristic rules is dealt with in the same way as the fuzzy number produced by the reasoning mechanism (which we will discuss below).

Now, for the third case ($b_o \leq P_R \leq a_o$), the fuzzy reasoning on a rule base is required. First, the rule bases for the s-agents and b-agents are presented in Tables 1 and 2, respectively. Again, the fuzzy numbers are all triangular fuzzy numbers as described in Section 2.2; the distance between a_o (or b_o) and P_R is expressed using the fuzzy linguistic terms: *far_from*, *medium_to*, and *close_to*, which are defined in Fig. 3, and *or* corresponds to operator *Max*. $\lambda_{s,1}, \dots, \lambda_{s,4}$ and $\lambda_{b,1}, \dots, \lambda_{b,4}$ are parameters decided by the risk attitude of the agent (see Section 4.2). Based on these rule bases, we can perform inference through the fuzzy reasoning mechanism presented in Section 2.2. The overall output of our fuzzy reasoning is a fuzzy number, i.e., a set of asks (bids) with membership degrees. For example, \tilde{z} may equal $(2.0, 0.02, 0.04)$, where 2.0 is the center, 0.02 is the left spread, and 0.04 is the right spread, and its membership degree might be given by:

TABLE 2
Fuzzy Rule Base for b-agents

IF	$(a_o \text{ is } \textit{far_from} \text{ or } \textit{medium_to } P_R) \text{ and } (b_o \text{ is } \textit{far_from } P_R)$
THEN	bid is $(b_o + \lambda_{b,1}, \theta, \chi)$.
IF	$(a_o \text{ is } \textit{far_from} \text{ or } \textit{medium_to } P_R) \text{ and } (b_o \text{ is } \textit{medium_to } P_R)$
THEN	bid is $(b_o + \lambda_{b,2}, \theta, \chi)$.
IF	$(a_o \text{ is } \textit{far_from} \text{ or } \textit{medium_to } P_R) \text{ and } (b_o \text{ is } \textit{close_to } P_R)$
THEN	bid is $(b_o + \lambda_{b,3}, \theta, \chi)$.
IF	$a_o \text{ is } \textit{close_to } P_R$
THEN	bid is $(P_R - \lambda_{b,4}, \theta, \chi)$.

$$\tilde{z}(x) = \begin{cases} 50x - 99 & \text{if } 1.98 \leq x \leq 2.0, \\ -25x + 51 & \text{if } 2.0 < x \leq 2.04. \end{cases}$$

Now, the decision sets DS_s (acceptable asks for s-agents) and DS_b (acceptable bids for b-agents) can be determined. Suppose $z_s = (m_s, \theta_s, \chi_s)$ is the output fuzzy number of the fuzzy reasoning or the heuristic rules for an s-agent, $z_b = (m_b, \theta_b, \chi_b)$ is the output fuzzy number of the fuzzy reasoning or the heuristic rules for a b-agent, and the parameter π_s , for the s-agent, and π_b , for the b-agent, are the thresholds to decide to which degree the ask (bid) could be accepted. Again, π_s and π_b can be decided by the risk attitudes of the agents. The asks that the s-agent could submit are in the decision set:

$$DS_s = \{a \mid a \in D_s \cap \{a \mid z_s(a) \geq \pi_s\}\}, \quad (8)$$

where D_s is the valid asks set (see Definition 7). Similarly, the bids that the b-agent could submit are in the decision set:

$$DS_b = \{b \mid b \in D_b \cap \{b \mid z_b(b) \geq \pi_b\}\}, \quad (9)$$

where D_b is the valid bids set (see Definition 6).

Finally, the agent can decide whether to accept an ask (bid) or submit an ask (bid), or submit nothing. For an FL-agent, if the decision set, DS_s (DS_b), is empty, it shows that there is no acceptable asks (bids) at which this agent can make any profit, thus it will not submit an ask or a bid. Otherwise, the ask (bid) to be submitted is decided by the following formulae:

- for FL s-agents:

$$ask = \begin{cases} b_o & \text{if } b_o \in DS_s, \\ \arg \max_{a \in DS_s} \{z_s(a)\} & \text{otherwise;} \end{cases} \quad (10)$$

- for FL b-agents:

$$bid = \begin{cases} a_o & \text{if } a_o \in DS_b, \\ \arg \max_{b \in DS_b} \{z_b(b)\} & \text{otherwise.} \end{cases} \quad (11)$$




Fig. 3. Fuzzy sets used in fuzzy reasoning.

For an FL s-agent (b-agent), if the outstanding bid (ask) falls into DS_s (DS_b), it is a sign that b_o (a_o) is acceptable. The FL s-agent (b-agent) will submit b_o (a_o) in order to make a transaction at b_o (a_o). Otherwise, it will select the ask (bid) which corresponds to the maximum similarity degree among those asks (bids) constrained by DS_s (DS_b).

This completes the description of our FL-strategy for both buyer and seller agents in a CDA. We now illustrate its use in an exemplar scenario.

3.3 The FL-Strategy in Operation

Assume there are three valuation vectors for b-agents b_1 , b_2 , and b_3 :

$$\vec{V}_1 = \{3.3, 2.7, 2.4\},$$

$$\vec{V}_2 = \{2.8, 2.5, 2.2\},$$

$$\vec{V}_3 = \{2.7, 2.4, 2.1\},$$

and three cost vectors for s-agents s_1 , s_2 , and s_3 :

$$\vec{C}_1 = \{1.6, 2.2, 2.4\},$$

$$\vec{C}_2 = \{1.75, 2.0, 2.3\},$$

$$\vec{C}_3 = \{1.6, 1.9, 2.1\}.$$

Furthermore, suppose the CDA market is as follows (see Definition 1):

$$P_{CDA} = \langle g, \{b_1, b_2, b_3\}, \{s_1, s_2, s_3\}, (\vec{V}_1, \vec{V}_2, \vec{V}_3), (\vec{C}_1, \vec{C}_2, \vec{C}_3), 0.01, 30 \rangle.$$

In this market, there are three b-agents, each with valuations for three units, and three s-agents each with costs for three units. Consider the following situation (see Definition 5):

$$s^* = \langle 6, \{b_1, b_2, b_3\}, \{s_1, s_2, s_3\}, a_o, b_o, H_l \rangle.$$

The fuzzy sets employed in the FL-strategy are shown in Figs. 2 and 3. Based on the ranges of asks and bids, a difference below 0.01 in the ask (bid) value is assumed to be indifferent to the users here. Thus, we choose 0.01 as the price step, i.e., $\Delta_{price} = 0.01$. Also, for simplicity, the thresholds for all the fuzzy sets used in the rules are set to 0.5, i.e., $\gamma_{s,1} = \gamma_{s,2} = \gamma_{b,1} = \gamma_{b,2} = 0.5$. For all the fuzzy numbers involved, suppose their left spread $\theta = 0.02$ and their right spread $\chi = 0.02$, which ensures a reasonable degree of flexibility in this context.

Example 1. This example shows how to use the heuristic rules to submit a bid. For ease of explanation, let the history length $l = 3$ (however, it does not affect the rationale of the strategy). Let $H_l = \{2.3, 2.2, 2.1\}$, $b_o = 2.4$, $a_o = 2.5$, $P_1 = 2.5$, and an s-agent be about to submit its next ask. First, from s^* , $r = 6$. Then, by Definition 8, we have

$$P_R = p_{(\lfloor \frac{l+1}{2} \rfloor)} = p_{(\lfloor \frac{3+1}{2} \rfloor)} = p_{(2)} = 2.2.$$

From Fig. 2a, we can find that

$$A_1(b_o) = A_1(2.4) = 0.667 > \gamma_{s,1} = 0.5.$$

That is, $b_o = 2.4$ is considered to be much bigger than $P_R = 2.2$. Therefore, according to rule SR_{11} , the agent will accept b_o , i.e., $ask = b_o = 2.4$.

Example 2. This example shows how to use the fuzzy reasoning mechanism to submit a bid. For ease of explanation, we set the history length to be 5. Let $H_l = \{2.0, 2.4, 2.3, 2.2, 2.1\}$ ($l = 5$), $a_o = 2.85$, $b_o = 1.2$, and the FL b-agent b_1 (with valuation \vec{V}_1) be about to submit a new bid for its second unit of good, that is, the valuation of the second unit of good is $v_{12} = 2.7$. By Definition 8,

$$P_R = p_{(\lfloor \frac{l+1}{2} \rfloor)} = p_{(\lfloor \frac{5+1}{2} \rfloor)} = p_{(3)} = 2.2.$$

Since $b_o \leq P_R \leq a_o$, the fuzzy reasoning on the rule base is employed. Let $\lambda_{b,1} = 0.05$, $\lambda_{b,2} = 0.04$, $\lambda_{b,3} = 0.01$, and $\lambda_{b,4} = 0.02$. From Fig. 3, we can find $close_to(a_o) = 0.7$, $medium_to(a_o) = 0.5$, $far_from(a_o) = 0$, $close_to(b_o) = 0$, $medium_to(b_o) = 1$, and $far_from(b_o) = 0$. By (1), the four rules' firing levels in Table 2 are:

$$\begin{aligned} \alpha_1 &= \min\{\max\{far_from(a_o), medium_to(a_o)\}, \\ &\quad far_from(b_o)\} = \min\{\max\{0, 0.5\}, 0\} = 0, \\ \alpha_2 &= \min\{\max\{far_from(a_o), medium_to(a_o)\}, \\ &\quad medium_to(b_o)\} = \min\{\max\{0, 0.5\}, 1\} = 0.5, \\ \alpha_3 &= \min\{\max\{far_from(a_o), medium_to(a_o)\}, \\ &\quad close_to(b_o)\} = \min\{\max\{0, 0.5\}, 0\} = 0, \\ \alpha_4 &= close_to(a_o) = 0.7. \end{aligned}$$

Thus, according to Table 2, the four rules' outputs are:

$$\begin{aligned} \tilde{z}_1 &= (b_o + \lambda_{b,1}, \theta, \chi) = (1.20 + 0.05, 0.02, 0.02) \\ &= (1.25, 0.02, 0.02), \\ \tilde{z}_2 &= (b_o + \lambda_{b,2}, \theta, \chi) = (1.20 + 0.04, 0.02, 0.02) \\ &= (1.24, 0.02, 0.02), \\ \tilde{z}_3 &= (b_o + \lambda_{b,3}, \theta, \chi) = (1.20 + 0.01, 0.02, 0.02) \\ &= (1.21, 0.02, 0.02), \\ \tilde{z}_4 &= (P_R - \lambda_{b,4}, \theta, \chi) = (2.2 - 0.02, 0.02, 0.02) \\ &= (2.18, 0.02, 0.02). \end{aligned}$$

Finally, by (3), the overall output fuzzy number is calculated as follows:

$$\begin{aligned} \tilde{z}_0 &= \frac{\alpha_1 \times \tilde{z}_1 + \alpha_2 \times \tilde{z}_2 + \alpha_3 \times \tilde{z}_3 + \alpha_4 \times \tilde{z}_4}{\alpha_1 + \alpha_2 + \alpha_3 + \alpha_4} \\ &= \frac{0 \times \tilde{z}_1 + 0.5 \times \tilde{z}_2 + 0 \times \tilde{z}_3 + 0.7 \times \tilde{z}_4}{0 + 0.5 + 0 + 0.7} \\ &= \frac{0.5 \times \tilde{z}_2 + 0.7 \times \tilde{z}_4}{1.2} \\ &= \frac{0.5 \times (1.24, 0.02, 0.02)}{1.2} + \frac{0.7 \times (2.18, 0.02, 0.02)}{1.2} \\ &= \frac{(0.5 \times 1.24, 0.5 \times 0.02, 0.5 \times 0.02)}{1.2} \\ &\quad + \frac{(0.7 \times 2.18, 0.7 \times 0.02, 0.7 \times 0.02)}{1.2} \\ &= \frac{(0.62, 0.01, 0.01) + (1.526, 0.014, 0.014)}{1.2} \\ &= (1.79, 0.02, 0.02). \end{aligned}$$

Then, by Definition 6, the valid bids set is:

$$\begin{aligned} D_b &= \{b \mid b_o < b \leq \min(a_o, v_{ij})\} \\ &= \{b \mid 1.2 < b \leq \min(2.85, 2.7)\} \\ &= \{b \mid 1.2 < b \leq 2.7\}. \end{aligned}$$

And then, by (9), the bids that this b-agent can accept are in the decision set:

$$DS_b = \{b \mid b \in D_b \cap \{b \mid z_b(b) \geq \pi_b\}\} = \{1.78, 1.79, 1.80\},$$

where $\pi_b = 0.5$. Finally, by (11), we have

$$\arg \max_{b \in DS_b} \{z_b(1.78), z_b(1.79), z_b(1.80)\} = 1.79.$$

Thus, the bid that the b-agent will submit is $bid = 1.79$.

4 EVALUATION OF FL-AGENTS

This section investigates, in an empirical fashion, the influence of the key parameters of the FL-strategy, the selection of these parameters, and the comparison of the FL-strategy with a number of other prominent strategies that have been proposed in the literature.

4.1 The Experimental Setting

This section describes the settings for the experiments conducted in the rest of this paper. First, the time period that an agent can allow to elapse before sending a message about asks or bids is specified as an exponentially distributed random variable. This is chosen because: 1) each agent's timing decision is independent of domain characteristics such as costs or valuations, and 2) exponential distribution is often a good approximation of the actual distribution [30]. Second, to measure how well an agent performs in a CDA, we evaluate its profit (the monetary gain for the agent). For an s-agent, the gain on its i th unit sold is the difference between the price, p_i , received from a b-agent for that unit, and the cost, c_i , at which the unit is produced, i.e., $p_i - c_i$. If the s-agent sells m units at prices p_1, \dots, p_m , then its profit is $\sum_{1 \leq i \leq m} (p_i - c_i)$. Similarly, for a b-agent, if this agent trades n units of goods, its profit is $\sum_{1 \leq i \leq n} (v_i - p_i)$, where v_i is the valuation value for the i th unit and p_i is the price of buying the i th unit of good. For the




Fig. 4. Utility functions of agents with different attitudes.

rest of this paper, an agent's profit is calculated as the sum of the profit in 1,000 simulations.⁹

Based on the above settings, each experiment is composed of many sessions and then each session consists of 1,000 runs.¹⁰ In each run of the session, an s-agent is endowed with a number of units of goods whose costs are independently drawn randomly from a uniform distribution with support [1.00, 3.00]. A b-agent is endowed with a number of units of goods whose valuations are independently drawn from a uniform distribution with support [2.00, 4.00]. These values were chosen because the cost values are generally smaller than the valuation values [31]. Thus, this is consistent with reality. The supply of the market is calculated by the total number of units of goods that all the s-agents want to sell; and the demand is calculated by the total number of units of goods that all the b-agents desire to buy. For example, if there are five s-agents and five b-agents in the market, each s-agent is endowed with five units of goods, and each b-agent is endowed with six units of goods to buy. The supply is 25 and the demand is 30.

4.2 Agents with Attitudes Toward Risk

This section first defines the different attitudes toward risk that an agent can adopt and then analyzes the influences of these attitudes through experiments based on the settings described above. Due to the complexity and uncertainty of the CDA bidding problem, it is not possible to analytically determine the optimal configuration of parameter values for a given context [10]. The best that can be achieved is to know the likely range of parameters such that the agent will perform effectively. To this end, the concept of attitude toward risk is introduced. Individual attitudes to risk can be characterized according to how an agent approaches a fair gamble [32]; they can be: risk-neutral, risk-averse, or risk-seeking. Take the utility functions of an s-agent as an example. In Fig. 4, the price of the outstanding bid (b_o) appears on the horizontal axis and the utility generated by accepting the current b_o is shown on the vertical axis. For the same value of p , agents with different risk attitudes have different utilities; that is, $U^{(A)}(p) \geq U^{(N)}(p) \geq U^{(S)}(p)$. The agent with the utility function $U^{(A)}$ represents the risk-averse agent which takes minimal risks with its actions. Suppose the cost of the current unit of the good is c , as a result, it is unwilling to sacrifice a sure profit of $(b_o - c)$, although there may be a greater chance of gaining more

profit. In short, risk-averse agents reject fair gambles. In contrast, there are agents that actually prefer fair gambles to sure results. These agents are called risk-seeking and are represented by the utility function $U^{(S)}$. The agents with the attitude between these two extremes are called risk-neutral agents and their utility function is always represented as a straight line. This kind of agent will be indifferent if the sure result and the gamble have the same expected utilities.

Thus, in the FL-strategy, given the same fuzzy sets, different parameters will correspond to different agent attitudes.

Definition 9. Given the same situation, suppose the utilities of an ask a that two s-agents s and s' submit are $U_s(a)$ and $U_{s'}(a)$, respectively. For all $a \in D_s$, $s \succeq_a^s s'$ if and only if $U_s(a) \geq U_{s'}(a)$. If $s \succeq_a^s s'$, then agent s is said to be more averse toward risk than agent s' .

The following corollary is a straightforward result of Definition 9.

Corollary 1. Given the same situation, suppose the utilities of submitting an ask a for three agents $s^{(A)}$, $s^{(N)}$, and $s^{(S)}$ are $U_s^{(A)}(a)$, $U_s^{(N)}(a)$, and $U_s^{(S)}(a)$, respectively. For all $a \in D_s$, $s^{(A)} \succeq_a^s s^{(N)} \succeq_a^s s^{(S)}$ if and only if

$$U_s^{(A)}(a) \geq U_s^{(N)}(a) \geq U_s^{(S)}(a).$$

Corollary 2. For three s-agents $s^{(A)}$, $s^{(N)}$, and $s^{(S)}$, represented by

$$(\beta_{s,1}^{(A)}, \beta_{s,2}^{(A)}, \gamma_{s,1}^{(A)}, \gamma_{s,2}^{(A)}, \lambda_{s,1}^{(A)}, \dots, \lambda_{s,4}^{(A)}),$$

$$(\beta_{s,1}^{(N)}, \beta_{s,2}^{(N)}, \gamma_{s,1}^{(N)}, \gamma_{s,2}^{(N)}, \lambda_{s,1}^{(N)}, \dots, \lambda_{s,4}^{(N)}),$$

and $(\beta_{s,1}^{(S)}, \beta_{s,2}^{(S)}, \gamma_{s,1}^{(S)}, \gamma_{s,2}^{(S)}, \lambda_{s,1}^{(S)}, \dots, \lambda_{s,4}^{(S)})$, respectively. If all the following conditions hold:

1. $\beta_{s,i}^{(A)} > \beta_{s,i}^{(N)} > \beta_{s,i}^{(S)}$ (for $i = 1$ and 2)
2. $\gamma_{s,1}^{(A)} < \gamma_{s,1}^{(N)} < \gamma_{s,1}^{(S)}$,
3. $\gamma_{s,2}^{(A)} > \gamma_{s,2}^{(N)} > \gamma_{s,2}^{(S)}$,
4. $\lambda_{s,j}^{(A)} > \lambda_{s,j}^{(N)} > \lambda_{s,j}^{(S)}$ (for $j=1, 2$, and 3) and

$$\lambda_{s,4}^{(A)} < \lambda_{s,4}^{(N)} < \lambda_{s,4}^{(S)},$$

then, $s^{(A)} \succeq_a^s s^{(N)} \succeq_a^s s^{(S)}$.

Proof. For an s-agent, for all $a \in D_s$, $U_s(a)$ is a nondecreasing function. That is, the bigger the ask, the more utility the agent obtains. Let the ask submitted by each s-agent be $a^{(A)}$, $a^{(N)}$, and $a^{(S)}$. From conditions 1 to 4, we can always get $a^{(A)} < a^{(N)} < a^{(S)}$. That is, $s^{(A)}$ always submits a lower ask compared with $s^{(N)}$ and $s^{(S)}$, and that $s^{(N)}$ is always lower than $s^{(S)}$. Thus, $U_s^{(A)}(a) \geq U_s^{(N)}(a) \geq U_s^{(S)}(a)$. Based on Corollary 1, we have: $s^{(A)} \succeq_a^s s^{(N)} \succeq_a^s s^{(S)}$. \square

Similarly, \succeq_a^b can be defined as follows:

Definition 10. Given the same situation, suppose the utilities of a bid b for two b-agents b and b' submit are $U_b(b)$ and $U_{b'}(b)$, respectively. For all $b \in D_b$, $b \succeq_a^b b'$ if and only if $U_b(b) \geq U_{b'}(b)$. If $b \succeq_a^b b'$, then agent b is said to be more averse toward risk than agent b' .

9. This number is chosen because it is sufficient to produce statistically significant results. By a t -test, the p value of 0.0007 is reported from the sample of 900 runs and that of 1,000 runs. Thus, the profit variance for the two samples are virtually the same and the results are therefore statistically significant.

10. From the beginning of the CDA to its termination is called a run. One thousand runs with the same s-agents and b-agents make up a session.

The following corollary is a straightforward result of Definition 10.

Corollary 3. *Given the same situation, suppose the utilities of submitting a bid b for the b-agents $b^{(A)}$, $b^{(N)}$, and $b^{(S)}$ are $U_b^{(A)}(b)$, $U_b^{(N)}(b)$, and $U_b^{(S)}(b)$ respectively. For all $b \in D_b$,*

$$b^{(A)} \succeq_a b^{(N)} \succeq_a b^{(S)}$$

if and only if $U_b^{(A)}(b) \geq U_b^{(N)}(b) \geq U_b^{(S)}(b)$.

Corollary 4. *For three agents $b^{(A)}$, $b^{(N)}$, and $b^{(S)}$ represented by*

$$\begin{aligned} &(\beta_{b,1}^{(A)}, \beta_{b,2}^{(A)}, \gamma_{b,1}^{(A)}, \gamma_{b,2}^{(A)}, \lambda_{b,1}^{(A)}, \dots, \lambda_{b,4}^{(A)}), \\ &(\beta_{b,1}^{(N)}, \beta_{b,2}^{(N)}, \gamma_{b,1}^{(N)}, \gamma_{b,2}^{(N)}, \lambda_{b,1}^{(N)}, \dots, \lambda_{b,4}^{(N)}), \end{aligned}$$

and $(\beta_{b,1}^{(S)}, \beta_{b,2}^{(S)}, \gamma_{b,1}^{(S)}, \gamma_{b,2}^{(S)}, \lambda_{b,1}^{(S)}, \dots, \lambda_{b,4}^{(S)})$, respectively. If all the following conditions hold:

1. $\beta_{b,i}^{(A)} > \beta_{b,i}^{(N)} > \beta_{b,i}^{(S)}$ (for $i = 1$ and 2)
2. $\gamma_{b,1}^{(A)} < \gamma_{b,1}^{(N)} < \gamma_{b,1}^{(S)}$
3. $\gamma_{b,2}^{(A)} > \gamma_{b,2}^{(N)} > \gamma_{b,2}^{(S)}$
4. $\lambda_{b,j}^{(A)} > \lambda_{b,j}^{(N)} > \lambda_{b,j}^{(S)}$ (for $j = 1, 2,$ and 3) and $\lambda_{b,4}^{(A)} < \lambda_{b,4}^{(N)} < \lambda_{b,4}^{(S)}$

then, $b^{(A)} \succeq_a b^{(N)} \succeq_a b^{(S)}$.

The proof for Corollary 4 is not presented because of space limitations, although it is similar to that of the above.

Now, given the fact that different parameters correspond to different attitudes toward risk, the key question is how to choose the appropriate risk attitudes of agents given a particular environment? The rest of this section is devoted to answering this question. In particular, the influence of the relation of supply and demand quantity (a key environmental factor) is considered.

Conjecture. The relation of supply and demand quantity influences the performance of agents with different attitudes. If the supply (demand) quantity is greater than the demand (supply) quantity, an s-agent (b-agent) with an averse attitude toward risk can make more profit.

To test our conjecture, a series of six experiments were conducted. Beside the three aforementioned kinds of agents (risk-averse, risk-seeking, and risk-neutral) two extra kinds of agents are considered: agents between the neutral and averse attitude (weakly averse), and agents between the neutral and risk attitude (weakly risky). These are added in order to make the trend of influence of the risk attitudes to the market supply (demand) more clear. In each session, only one agent uses the FL-strategy, and from session to session, the attitude of the FL-agent varies from risk-averse to risk-seeking. All the other agents utilize one of our benchmark strategies.¹¹

Fig. 5 shows the profit of FL-agents in different sessions. Figs. 5a and 5b show the profit of agents when supply is equal to demand. In these cases, the left-hand bars in each group are always taller than the other bars in the same group. These bars represent the profits of the averse agents.

11. In fact, they use the ZI-strategy which will be described in Section 4.3 (this is the earliest and simplest of our benchmark strategies). Thus, if there is one FL s-agent in the market, there are four s-agents and five b-agents using the ZI-strategy; if there is one FL b-agent, there are five s-agents and four b-agents using the ZI-strategy.

A risk-averse agent is easily satisfied, so it can make transactions quickly and with a high volume. As a result, its profit is high. Thus, in this environmental setting, an averse agent can make more profit, whether it is an s-agent or a b-agent. In Fig. 5c, supply is greater than demand, thus an s-agent is in an unfavorable position. The chances of selling a good are small because there are very few b-agents. So, an averse agent makes more profit. This explanation also holds for the b-agents in Fig. 5f. Fig. 5d shows the behavior of agents when supply is less than demand. Here, the attitude with which an s-agent performs best varies with the change of the difference between supply and demand quantity. At the beginning, when the difference is small, the weakly averse agents always make more profit. However, with the increase in supply, the trend is that an s-agent with a risk-seeking attitude can make more profit. This is caused by the fact that the market is not very competitive when the difference between supply and demand is small. This also explains the bars for b-agents in the situation when supply is less than demand (Fig. 5e).

In summary, this experiment clearly shows how to choose the suitable attitude for an agent in different situations *if we have knowledge about the real-time supply and demand*. As indicated by our conjecture, an s-agent should be averse when supply is greater than or equal to demand, and the attitude should change from risk-averse to risk-seeking with the increase of demand when the supply is less than demand. For a b-agent, it should be averse when supply is less than or equal to demand, and the attitude should change from risk-averse to risk-seeking with the increase of supply when the supply is greater than demand.

4.3 Benchmarking the FL-Strategy

Having determined the best parameter configuration with respect to risk in a given environment, this section compares our strategy (with the risk attitude tailored to the environment) with a number of others that have been proposed in the literature. These other strategies represent the most widely cited strategies for agents participating in CDAs. In more detail, the benchmark strategies are:

- **Zero Intelligence (ZI) strategy** [33]. A ZI b-agent submits a bid drawn randomly between outstanding bid (b_o) and the valuation of its current unit. Similarly, a ZI s-agent submits an ask drawn randomly between the cost of its current unit and outstanding ask (a_o). This strategy is an extension of the “budget-constraint zero intelligence” trader in the economics literature¹² [33]. In [33], the lower limit that a b-agent submits as a bid is 0 and the upper limit that an s-agent submits as an ask is 1. We believe our extension is appropriate because we are dealing with single unit trades and because the ask (bid) bounding conditions increase the possibility of matching an outstanding bid (ask).
- **Fixed Mark-up (FM) strategy**. An FM b-agent (s-agent) submits the outstanding bid (ask) plus (minus) some predefined mark-up (a specialization of [35]). This is a simple strategy because the agent

12. Actually, the ZI-strategy is equivalent to the ZI-C strategy proposed by Gode and Sunder [34].

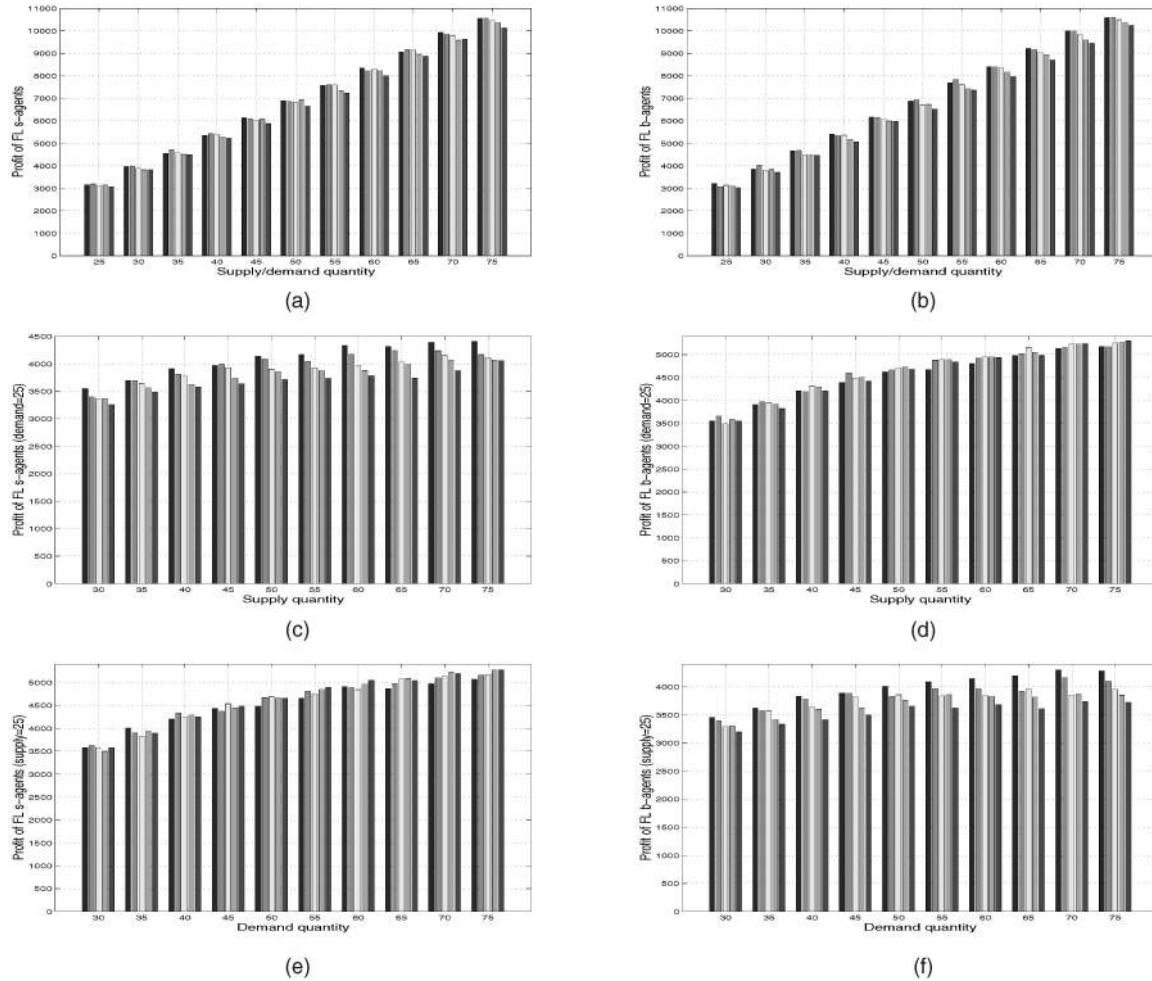


Fig. 5. Performance of FL-agents with different risk attitudes. In each group of bars, the five bars represent, from left to right, the agents with different attitudes: averse, weakly averse, neutral, weakly risky and risky. The horizontal axis shows the supply (demand) quantity of the session. The vertical axis represents the profit of the various FL-agents in the session.

does not need to model other agents and it tries to reduce the ask-bid spread until its cost or valuation is met.

- **Chris Preist (CP) strategy.** The CP-strategy consists of a small number of heuristics and a learning rule [36]. The heuristics first determine the target profit margin based on the current outstanding bid (ask) and an independent random variable distributed in the range $[0, 0.2]$. Then, given the target, a CP agent does not jump straight to that value, but moves toward the target at a learning rate which determines the speed of the adjustment.
- **Gjerstad-Dickhaut (GD) strategy.** The GD-strategy [11] is a more sophisticated strategy. A GD agent records all the asks (bids) made in the history H occurring in the last several transactions. From the history, an agent can compute the probability of a bid or ask being accepted. For example, for a buyer,

$$\hat{q}(b) = \frac{TBL(b) + AL(b)}{TBL(b) + AL(b) + RBG(b)}, \quad (12)$$

where $\hat{q}(b)$ is the probability of b being accepted, $TBL(b)$ is the number of accepted bids not greater than b in H , $AL(b)$ is the number of asks not greater

than b in H , and $RBG(b)$ is the number of rejected bids not less than b in H . Then, cubic spline interpolation is used to compute the probability of a given bid being accepted given the history. A GD b-agent submits a bid, b , which maximizes $\pi_b(v - b)$, where π_b is the belief function of a bid that is accepted, and v is the valuation of the good. Similarly, a GD s-agent submits an ask a which maximizes $\pi_s(a - c)$, where π_s is the belief function of an ask that is accepted, and c is the cost of the good to sell.

To evaluate the behavior of each agent, we compare their profits in three situations:

1. supply equals demand (Figs. 6a and 6b),
2. supply is less than demand (Figs. 6c and 6d) and
3. supply is greater than demand (Figs. 6e and 6f).

In each subfigure of Fig. 6, the horizontal axis shows the supply (demand) quantity and the vertical axis represents the profit of agents using various strategies. There are five curves in each subfigure and each one represents the profit of one kind of strategy. Given the same supply (demand), the bigger the profit, the better the strategy.

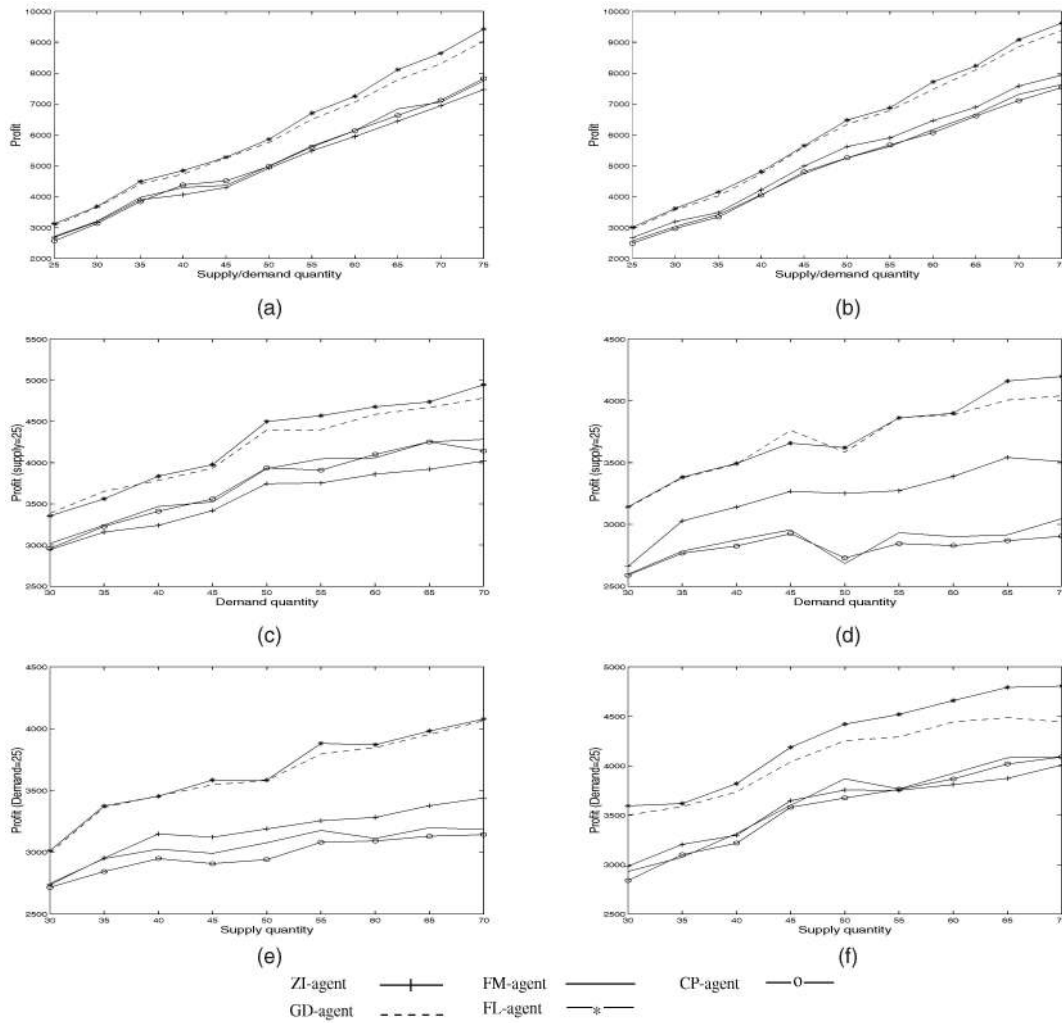


Fig. 6. Performance of agents with different strategies. The horizontal axis represents the demand (or supply) configuration of the corresponding market and the vertical axis represents the total profit of the corresponding agent using a specific strategy in one session. There are 5 s-agents and 5 b-agents in each experiment. In (a) and (b), supply is equal to demand, the unit of good that each agent is endowed with to buy (sell) increases from session to session in the range [5, 15], the total supply (demand) is shown in the horizontal axis. In (c) and (d), supply is less than demand, the unit of good that each s-agent is endowed with is fixed to 5, and the unit of good for b-agents increases from session to session in the range [6, 14]. In (e) and (f), supply is greater than demand, the unit of good that each b-agent is endowed with is fixed to 5, and the unit of good for s-agents increases from session to session in the range [6, 14].

From Fig. 6, we can see that the FL-agents often obtain higher profits than all the other corresponding agents. The exception is that FL-agents perform slightly worse than the GD-strategy when 1) for the b-agent, supply = 25 and demand = 40, 45, and 50, respectively, and 2) for the s-agent, supply = 25 and demand = 30 and 35, respectively. The reason for this inferiority is that our agent cannot adjust its risk attitude during the course of the CDA and the value that is used is based on the experimental results in Section 4.2 which is for the general case. However, in most situations, FL-agents can outperform agents using other strategies. We attribute this success to two factors. First, when making a decision in a given situation, an FL-agent considers the outstanding ask, outstanding bid, and reference price (inferred from history) in deciding its next ask (bid). We believe reference price is a very important factor in bidding, and our strategy is the only one to exploit this information. Second, the FL-strategy can dynamically

vary the rate of increase (decrease) in bid (ask) according to the prevailing context. Sometimes, for example, an FL-agent can jump from a very low price to a transaction price. This is markedly different from the other strategies which only increase (decrease) their bids gradually.

The performance of the other strategies is statistically worse than our FL-agents. GD-agents behave worse than the FL-strategy although they do maintain a history. However, these agents ignore the outstanding ask (bid), which we believe is one of the most important factors in deciding an agent's next bid (ask). The other three strategies ignore the transaction history. ZI-agents have no knowledge about the auction, they submit their bids (asks) randomly. However, ZI-agents can sometimes deal at a very low bid or high ask. FM-agents and CP-agents can only increase (decrease) their bids (asks) in a fixed step or small varied steps without caring about the outstanding bid (ask). Thus, they miss out on some deals which they should have made.




Fig. 7. Risk Attitude of an agent.

5 THE ADAPTIVE FL-AGENTS

In the above experiments, the risk attitude of the FL-agent is selected manually based on design time knowledge of the relation between supply and demand. However, in many environments, this information is a priori unknown. Also, in an open CDA, the number of agents can be changing continuously as new agents enter the market and the existing ones drop out. Further, the parameters that are suitable in one CDA market may not behave well in others because success is inextricably linked to the strategies of the competitor agents. For all these reasons, we believe it is desirable for an FL-agent to have the ability to automatically adapt itself to its market context. Thus, this section reports on a number of extensions in this direction that we made to the basic FL-strategy.

5.1 Learning Principle for FL-agents

As discussed in Section 3.1, each FL-agent has a reference price (P_R) to decide whether it sells (buys) a good at a profitable price. Given this price, an agent can submit an ask (bid) based on its risk attitude (parameters). However, different attitudes can lead to different asks (bids) (Section 4.2). Furthermore, even the same asks (bids) have different effects in different environments. Thus, an agent needs another measure to decide whether it submits too high a bid or too low an ask. To this end, an agent can observe how frequently it can make transactions. If an s-agent (b-agent) waits too long to conduct a deal, it shows that it should be more averse in the next round if it is to make more transactions. On the other hand, if an s-agent (b-agent) can transact very frequently, it is a sign that its bids (asks) are too high (low). Thus, during the next round of the CDA, the agent should change its attitude in the direction of risk-seeking (hoping it can still make a transaction while increasing its profit). We call this kind of hill-climbing behavior the *adaptive FL-agent* (denoted A-FL-agent).

Suppose agent i 's attitude is expressed by $A_{attitude}^{(i)}$ which corresponds to a value in $[-1, +1]$ as shown in Fig. 7. Each value of $A_{attitude}^{(i)}$ corresponds to a group of parameters which define its attitudes toward risk that satisfy Corollary 2 or Corollary 4. Formally, the learning principles for A-FL-agents can be expressed as the rules in Table 3 (where δ is the minimum step and r ($r > 0$) is the learning rate). The terms "waits_long" and "transacts_frequently" are expressed as two fuzzy sets shown in Fig. 8.

In this context, the learning rate r determines the speed with which the adjustment takes place. Some agents may adapt themselves slowly but steadily, while others may change their attitudes quickly. Thus, we compare three different representative adjustment methods:

1. an agent increases (decreases) at the constant rate minimum step δ , that is, $r = 1$;
2. $r = m\delta$, where $m > 1$; that is, the agent increases (decreases) at a bigger step than δ ;

TABLE 3
Learning Rules for A-FL-Agents

IF	agent i waits_long to transact
THEN	$A_{attitude}^{(i)} = A_{attitude}^{(i)} - r\delta$
IF	agent i transacts_frequently
THEN	$A_{attitude}^{(i)} = A_{attitude}^{(i)} + r\delta$

$A_{attitude}^{(i)}$ denotes the attitude of agent i , r is the learning rate, and δ is the minimum step.

3. r is an independent random number uniformly distributed in the range $[1, \tau]$, where τ is the maximum adjustment number.

In these experiments, we assume there is no abrupt increase or decrease in the supply and demand quantity. That is, over any period, the CDA market in each session is relatively stable (i.e., there is a fixed supply and demand quantity). Further, to compare the performance of different learning rates, we compare the three adaptive FL-agents and the FL-agents with the parameters shown by the selection principle in Section 4.2. The experiments are conducted in different situations (see Fig. 9).

Generally, the agents whose learning rate is 1 (the agent which increases (decreases) the attitude value at a small and constant rate) perform best. This is because this kind of agent can fine tune its parameters which avoids overresponse to supply (demand) changes in the market. Thus, we choose this kind of learning rate to adjust the attitude for the adaptive FL-agents. Also, the adaptive FL-agents with a small learning rate do better than other FL-agents. From Fig. 9, it can be seen that the A-FL_1 agents obtain a higher profit than the FL-agents. This means that even without the knowledge of supply and demand, the adaptive agent can effectively tailor its strategy to its prevailing circumstances. While this result is promising, there is the caveat assumption that there is no abrupt change in supply or demand. In such circumstances, a learning rate that takes small steps may not be able to respond quickly enough to be effective. However, the best means of dealing with abrupt change is left for future work.

5.2 Comparison with the Other Strategies

In this experiment, we compare the adaptive FL-agents with the four benchmark strategies of the previous section (ZI, FM, CP, and GD).¹³ Fig. 10 clearly shows that the adaptive strategy is effective. A-FL-agents behave better, sometimes much better, than all the other strategies. This is because an A-FL-agent can tailor its bidding behavior to the prevailing market context.

Besides the profit of the agents, we also investigated the transaction price distribution of each agent. We did this because this metric is a good indication of how consistently an agent performs in a CDA. Table 4 shows the transaction price distributions in 1,000 runs of two agents¹⁴ when supply is equal to demand (Table 4a),

13. The set-up of this experiment is the same as that of Section 4.3.

14. Due to the limitation of space, we only show the transaction price distribution of two agents. In total, in the experiment shown in Fig. 10, there are 28 configurations of different market situations. Two agents (one s-agent and one b-agent) from six sessions are shown in Table 4. To distinguish these two agents, we use *strategy_s/b_1* and *strategy_s/b_2*, respectively. For example, *A-FL_b_1* means b-agent 1 utilizes A-FL strategy.




Fig. 8. Two fuzzy sets for transaction rate. This rate is calculated by the number of transactions made by an agent divided by the total transaction numbers in the market after the latest change of the agent's attitude towards risk.

supply is less than demand (Table 4b), and supply is greater than demand (Table 4c). In Table 4, P_0 is the average equilibrium price of 1,000 runs, obtained from the supply and demand curves. Accordingly, Q_0 is the average quantities at the equilibrium prices. \bar{P} is the average price for each agent and the Change Rate (CR) shows the percentage of $(\bar{P} - P_0)$ to the equilibrium price P_0 . This is a key measure of how well the agent behaves and is calculated in the following way:

$$CR(\bar{P}, P_0) = \frac{\bar{P} - P_0}{P_0} \times 100\%.$$

Generally, for a b-agent, the lower the average price, the better the strategy; for an s-agent, the higher the average price, the better the strategy. In Table 4, the average prices of A-FL b-agents are always the lowest among all the b-agents. This means that A-FL b-agents always pay low prices to acquire goods. Also, the average prices of A-FL s-agents are always the highest among all the s-agents. This means that A-FL s-agents always sell goods at high prices. An agent's Change Rate also gives an indication of how high or low the average transaction price is compared with the equilibrium price of the CDA. For an s-agent, the higher the CR, the higher the price at which it sells its goods; for a b-agent, the lower the CR, the lower the price at which it buys the goods. Our A-FL s-agents always get the maximum CR value and A-FL b-agents always obtain the minimum CR value. This means that our adaptive FL-agents buy goods at the lowest average price and sell goods at the highest average price among all the agents using various strategies. Thus, the adaptive FL-agents outperform all the other strategies.

5.3 Collective Behavior of A-FL-Agents

Since the A-FL strategy is effective in making good profits in a CDA, we expect many A-FL-agents may appear in a given CDA market. Thus, we need to test the efficiency of a CDA market that is populated with multiple A-FL-agents. In particular, we would like to investigate how the performance of an A-FL-agent changes as the percentage of A-FL-agents in the population increases, and to what extent the efficiency of the CDA market is affected by this change in population.

We test the collective behavior of A-FL-agents in situations where the quantity of demand is

1. greater than,
2. equal to, and
3. less than the quantity of supply, respectively.

For each situation, the experiment is composed of multiple sessions. Fig. 11 shows the results of the profits of A-FL-agents in different sessions under different situations.¹⁵ The horizontal axes represent the session numbers and the vertical axes represent the sum of the profits in 1,000 runs of each agent. The curves show each agent's profit in different sessions.

In different sessions, we only change the strategy of one agent. Actually, we increase the number of agents that employ the A-FL strategy by one in each session. Take Fig. 11a as an example. In this situation, the demand is greater than supply, thus the competition among A-FL b-agents is highlighted. There are 10 b-agents and eight s-agents and each agent has three units of goods to buy (sell). In session 1, only b-agent b_1 uses the A-FL strategy; in session 2, we change the strategy of b-agent b_2 to the A-FL strategy, while fixing all the other parameters; in session 3, the strategy that b_3 uses is changed; and, finally, in session 10, all the b-agents use the A-FL strategy.

As shown in Figs. 11a, 11b, and 11c, the profit of A-FL b-agent b_1 decreases initially and then increases steadily. The trend of the change of the profit of agent b_2 and agent b_3 is similar. This phenomenon can be explained as follows: For agent b_1 , with an increasing number of A-FL-agents, more A-FL-agents compete with b_1 , and, thus, the profit of b_1 decreases. When the number of A-FL-agents is sufficiently large,¹⁶ A-FL-agents make the transaction prices of the market low. The same explanation holds for A-FL s-agents. Thus, the reference prices of each agent decreases from session to session and all the A-FL-agents make greater profits.

Since A-FL-agents make more profits in the long term, another question arises. How will the CDA market as a whole be affected with an increasing population of A-FL-agents? We evaluate the profit obtained by all agents in the market divided by the surplus when agents trade their goods at the equilibrium price to determine the efficiency of the market.¹⁷ Tables 5, 6, and 7 summarize this efficiency data of the collective profit of all the agents and the market efficiency from experiments of different sessions. As can be seen, for CDA markets with varying numbers of A-FL b-agents,¹⁸ both the total profit and efficiency of the market increase initially and then decrease little by little. For a CDA market with varying numbers of A-FL s-agents, the trend is obvious. The market becomes less efficient due to the increase in strategic reasoning of the A-FL agents. However, we can see that, from the session when no agent uses the

15. In order to investigate the performance of A-FL-agents, the profit of non-A-FL agents is not shown in Fig. 11. These agents use various randomly selected strategies from our set of benchmarks.

16. In the experiment, this number is 50 percent of the total number of b-agents or s-agents.

17. The efficiency of the market is obtained by the following formula (based on the intuitions described in [11]): $Efficiency = \frac{\sum_i p_a^{(i)}}{\sum_i p_e^{(i)}}$, where $p_a^{(i)}$ is the actual profit of the agent i in a session (1,000 runs); $p_e^{(i)}$ is the profit of the agent if the agent i trades its goods according to the equilibrium price given their costs and valuations of the goods in a session.

18. Note that the total number of agents in each set-up is fixed, that is, in Table 5, there are 8 b-agents and 8 s-agents and, in Table 6, there are 8 b-agents and 10 s-agents, and in Table 7, there are 10 b-agents and 8 s-agents.

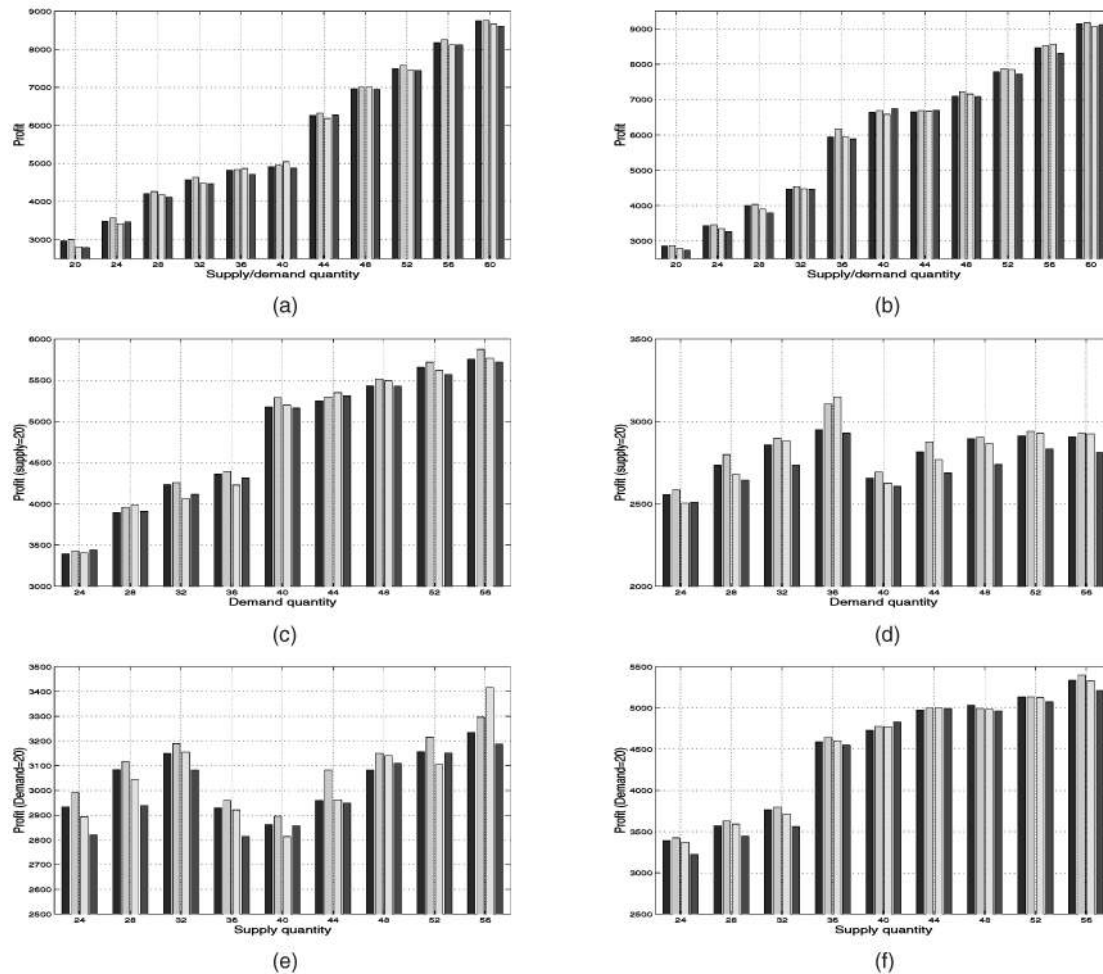


Fig. 9. Performance of A-FL-agents with different adaptive speeds. Initially, all the adaptive agents are risk-neutral. In each image, the bars in each group represent the agents (from left to right): FL-agents (without adaptivity), A-FL_1 agents ($r = 1$), A-FL_2 agents ($r = 5$), and A-FL_3 agents (r is drawn randomly from $[1, 10]$).

A-FL strategy to the session when all the b-agents and s-agents use the A-FL strategy, the efficiency of the market does not decrease significantly. That is, even in the worst case, the efficiency is still reasonably high. The market with all A-FL-agents is also investigated. In this market, all the agents use the A-FL strategy, and the efficiency is 85.45 percent for a market with 5 s-agents and 5 b-agents. This figure is still reasonable with respect to experiments shown in Tables 5, 6, and 7. Thus, we can conclude that widespread adoption of the A-FL strategy does not lead to a significant deterioration in the effectiveness of the overall market.

6 RELATED WORK

There are a number of strands of work that are related to what we have described in this paper. First, the work on bidding strategies for various forms of auctions. Second, the work on using fuzzy techniques to manage an agent's interactions. Finally, alternatives to fuzzy reasoning for coping with the uncertainties in bidding. Each of these areas will now be dealt with in turn.

Noncooperative game theory is an important tool for analyzing strategic interactions between agents [37]. However, one of its weaknesses is that the theory is only

suitable for highly stylized, simple settings [38], thus a clear game-theoretic solution to the CDA problem is not possible. The Recursive Modeling Method [39] has been proposed as an approach for an agent to reason about other agents and generate an appropriate strategy for negotiation. However, in most practical cases, the agent can only build finite nesting models due to the limitation of acquiring knowledge. Thus, with this approach, not all the information in the recursive model may be relevant to the agent and it is possible that little or no information may be available for the agent to use. Park et al. [40] propose the adaptive agent bidding strategy (called the p-strategy) based on stochastic modeling for a CDA. The idea of the p-strategy is to model the auction process using a Markov Chain (MC). However, in many cases, it is hard to obtain the probability values required for the MC model, such as the transition probabilities and the probabilities of success and failure for particular trading actions. Moreover, the computation involved in this approach is large. Badea [41] applied Inductive Logic Programming (ILP) to induce trading rules for a CDA. He first identified buy (sell) opportunities from historical market data. Then, these buy (sell) opportunities are input as examples to an ILP learner to produce understandable

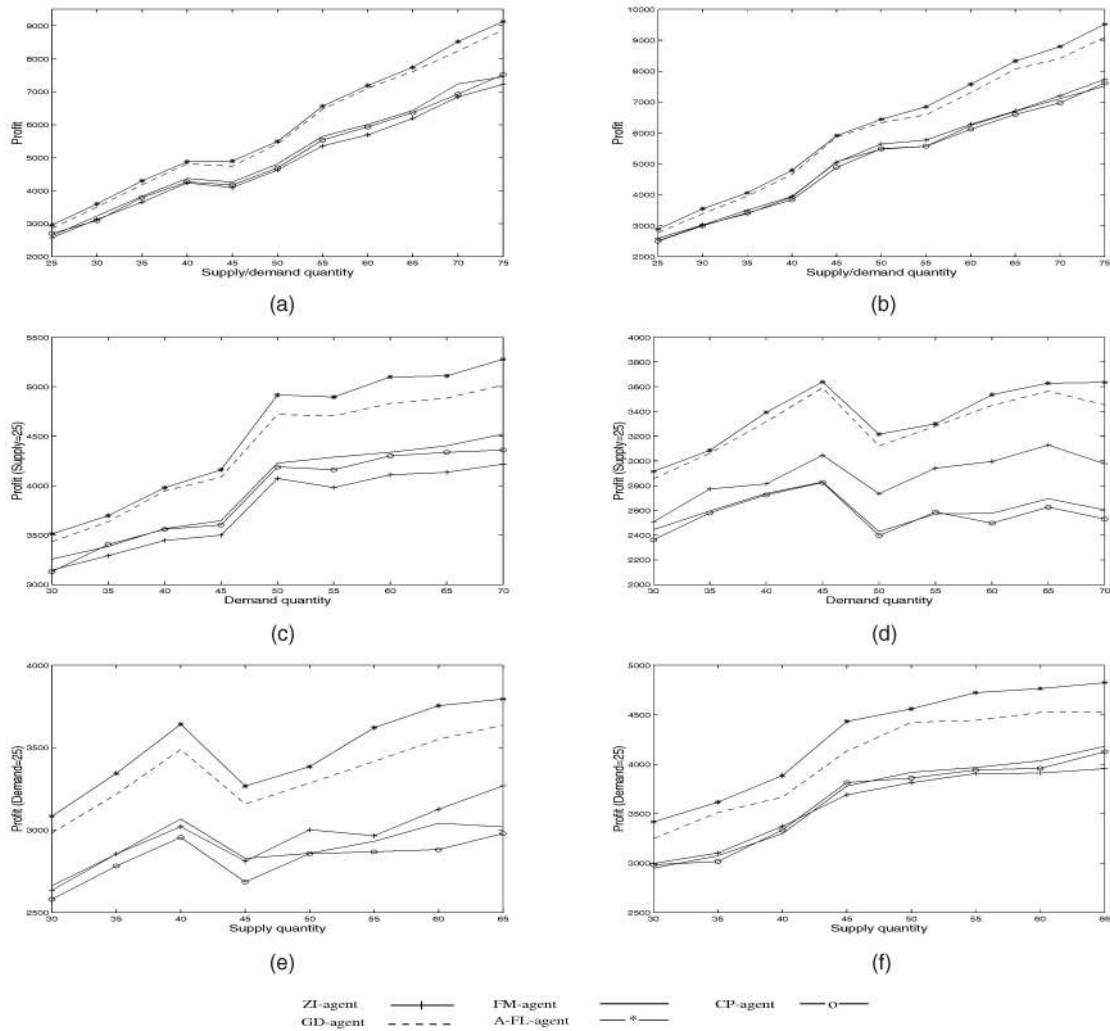


Fig. 10. Competition of A-FL-agents with other strategy agents.

rules. However, this learning strategy relies heavily on historical data which is often not available in the contexts we consider.

Turning now to the use of fuzzy techniques to manage an agent's interactions. Faratin et al. [42] used fuzzy similarity to compute trade offs among multiple attributes during bilateral negotiations. An agent first generates some (all) of the potential contracts for which it receives a score θ . Then, the agent finds the contract on the indifference curve for θ which has the maximum similarity degree to the last proposal from the negotiation opponent. Here, fuzzy techniques are used to deal with a bilateral negotiation and the algorithm aims to find a win-win (cooperative) solution for both parties. However, in our work, there are multiple parties in a CDA and they are essentially competitive. Kowalczyk and Bui [43], [44] modeled the multiissue negotiation process as a fuzzy constraint satisfaction problem (FCSP). Their approach performs negotiation on individual solutions one at a time. During negotiation, an agent evaluates the offers, relaxes the preferences and constraints, and makes counteroffers to find an agreement for both parties. The issues negotiated over actually correspond to the constrained variables and the preferences, constraints, and each party's objectives are

expressed as fuzzy constraints over these issues. Using this method, the FCSP is to find a solution that maximizes the satisfaction of all constraints of the parties. In contrast with our work, they use fuzzy constraints in a bilateral negotiation on multiple issues, while our work aims to find a solution for an agent in a CDA and we use fuzzy sets and fuzzy rules to generate a bid (ask). Luo et al. [45], [46] developed a fuzzy constraint based framework for bilateral multiissue negotiations in competitive trading environments. The framework is expressed via two knowledge models, one for the seller agent and one for the buyer agent. The seller agent's domain knowledge consists of its multi-dimensional representation of the products or services it offers. The buyer agent's domain knowledge consists of the buyer's requirement/preference model (a prioritized fuzzy constraint problem) and buyer's profile model (fuzzy truth propositions). The buyer and seller agents exchange offers and counteroffers with additional constraints revealed or existing constraints being relaxed. Finally, a solution is found if there is one. Compared with our work, they solve the multiissue bilateral negotiation using fuzzy truth propositions and prioritized fuzzy constraint techniques, while we use fuzzy sets and fuzzy rules to deal with making bidding decisions.

TABLE 4
Transaction Price Distributions of Each Agent

$P_0 = 2.50, Q_0 = 37$			$P_0 = 2.50, Q_0 = 55$		
Agent	\bar{P}	CR	Agent	\bar{P}	CR
<i>ZI_s_1</i>	2.46	-1.54%	<i>ZI_s_2</i>	2.39	-4.39%
<i>FM_s_1</i>	2.49	-0.55%	<i>FM_s_2</i>	2.43	-2.65%
<i>CP_s_1</i>	2.47	-1.22%	<i>CP_s_2</i>	2.42	-3.03%
<i>GD_s_1</i>	2.58	3.33%	<i>GD_s_2</i>	2.53	1.31%
A-FL_s_1	2.59	3.47%	A-FL_s_2	2.55	2.09%
<i>ZI_b_1</i>	2.58	3.36%	<i>ZI_b_2</i>	2.54	1.69%
<i>FM_b_1</i>	2.59	3.66%	<i>FM_b_2</i>	2.55	2.08%
<i>CP_b_1</i>	2.59	3.77%	<i>CP_b_2</i>	2.56	2.37%
<i>GD_b_1</i>	2.46	-1.46%	<i>GD_b_2</i>	2.42	-3.11%
A-FL_b_1	2.44	-2.38%	A-FL_b_2	2.39	-4.53%

(a)

$P_0 = 2.64, Q_0 = 25$			$P_0 = 2.92, Q_0 = 29$		
Agent	\bar{P}	CR	Agent	\bar{P}	CR
<i>ZI_s_1</i>	2.60	-1.63%	<i>ZI_s_2</i>	2.81	-3.91%
<i>FM_s_1</i>	2.62	-0.60%	<i>FM_s_2</i>	2.85	-2.24%
<i>CP_s_1</i>	2.62	-0.94%	<i>CP_s_2</i>	2.84	-2.75%
<i>GD_s_1</i>	2.70	2.42%	<i>GD_s_2</i>	2.94	0.85%
A-FL_s_1	2.72	2.99%	A-FL_s_2	3.00	2.77%
<i>ZI_b_1</i>	2.69	1.88%	<i>ZI_b_2</i>	2.92	0%
<i>FM_b_1</i>	2.72	2.89%	<i>FM_b_2</i>	2.98	1.92%
<i>CP_b_1</i>	2.73	3.22%	<i>CP_b_2</i>	3.00	2.64%
<i>GD_b_1</i>	2.60	-1.56%	<i>GD_b_2</i>	2.82	-3.51%
A-FL_b_1	2.58	-2.19%	A-FL_b_2	2.81	-3.77%

(b)

$P_0 = 2.36, Q_0 = 25$			$P_0 = 2.12, Q_0 = 29$		
Agent	\bar{P}	CR	Agent	\bar{P}	CR
<i>ZI_s_1</i>	2.35	-0.57%	<i>ZI_s_2</i>	2.15	1.28%
<i>FM_s_1</i>	2.36	0%	<i>FM_s_2</i>	2.13	0.67%
<i>CP_s_1</i>	2.33	-1.12%	<i>CP_s_2</i>	2.12	-0.11%
<i>GD_s_1</i>	2.43	2.93%	<i>GD_s_2</i>	2.22	4.79%
A-FL_s_1	2.46	4.06%	A-FL_s_2	2.24	5.79%
<i>ZI_b_1</i>	2.44	3.43%	<i>ZI_b_2</i>	2.26	6.56%
<i>FM_b_1</i>	2.44	3.43%	<i>FM_b_2</i>	2.23	5.37%
<i>CP_b_1</i>	2.45	3.63%	<i>CP_b_2</i>	2.25	6.02%
<i>GD_b_1</i>	2.36	-0.14%	<i>GD_b_2</i>	2.15	1.49%
A-FL_b_1	2.32	-1.50%	A-FL_b_2	2.12	-0.23%

(c)

- (a) Transaction prices of each agent when supply = demand.
 (b) Transaction prices of each agent when supply < demand.
 (c) Transaction prices of each agent when supply > demand.

Finally, we consider the alternatives to fuzzy reasoning for handling with uncertainty in agent interactions (see [47] for a comprehensive survey about handling uncertainty in agent systems). As stated, we chose fuzzy logic based methods because they have proven to be a practicable solution in solving decision making problems under uncertainty (e.g., [17], [18], [19], [20]). Fuzzy rules are the most visible manifestation of this approach and have been successfully used in industrial applications, manufacturing, process control, automotive control, and financial trading [48]. There are, however, alternative techniques for handling with uncertainties. For example, the possibility based approach [49], [50] has been used to perform multiagent reasoning under uncertainty for bilateral negotiation. In this work, uncertainties due to the lack of knowledge about other agents' behaviors are modeled by possibility distributions. Based on information from a case base of previous negotiation behaviors, the possibility distributions are generated by choosing the most similar situation to the

current context and the most similar price from the case base. Since this approach relies on a case base, it is unclear what would happen if no similar situations were available. Moreover, even if a similar case exists, it is possible that the strategy used successfully in that situation does not work in the current environment due to the variety of competitors. The Bayesian learning method [51] has also been used to explicitly model multiissue negotiation in a sequential decision making model. In this work, a Bayesian network is used to update the knowledge and belief each agent has about the environment and other agents, and offers and counteroffers between agents during bilateral negotiations are generated based on Bayesian probabilities. However, this method is inappropriate in our context because assigning prior probabilities of a bid (ask) being accepted is practically impossible given the dynamism and uncertainty of the CDA context.

7 CONCLUSIONS AND FUTURE WORK

This paper developed new algorithms that guide an agent's buying and selling behavior in a CDA. The FL-strategy uses heuristic fuzzy rules and a fuzzy reasoning mechanism to decide what bids or asks to place. We then extended this strategy so that the agent could adapt its bidding behavior to its prevailing market context. In both cases, we benchmarked the performance of our algorithm against the most prominent alternatives available in the literature. This evaluation showed the superior performance of our method. This result is especially promising since the benchmark strategies have been shown to outperform human bidders in experimental settings. Speaking more generally, we also believe that the development of efficient and practicable algorithms for bidding behavior increase the opportunities of using CDAs as the auction protocol for service marketplaces. We, therefore, view our contribution as an important step in this direction.

In more detail, the experiments in Section 4.2 show how to select the appropriate risk attitude for an agent in different situations. The result is consistent with our conjecture: If supply (demand) quantity is greater than demand (supply) quantity, an s-agent (b-agent) with an averse attitude toward risk can make more profit. Based on this selection principle, the experiment in Section 4.3 shows that the FL-strategy outperforms some of the most commonly used bidding strategies in a range of situations. Since agents often have no prior knowledge of the relation between supply and demand, it is not always possible to tell in advance what kind of attitude an agent should have. Thus, adaptive FL-agents are introduced which can tailor their strategy to the supply (demand) of the market. Through the experiments in Section 5.1, we find that the learning rate which is adjusted in small steps behaves best in an environment in which the supply and demand do not change abruptly. The experiments in Section 5.2 show that A-FL-agents always outperform other benchmarking strategy agents in various situations. The transaction price distribution of agents using different strategies shows that an A-FL-agent always sells (buys) goods at higher (lower) prices than agents using other strategies. Finally, in Section 5.3, we investigate to what extent the behavior of A-FL-agents and the efficiency of the CDA market are affected by the

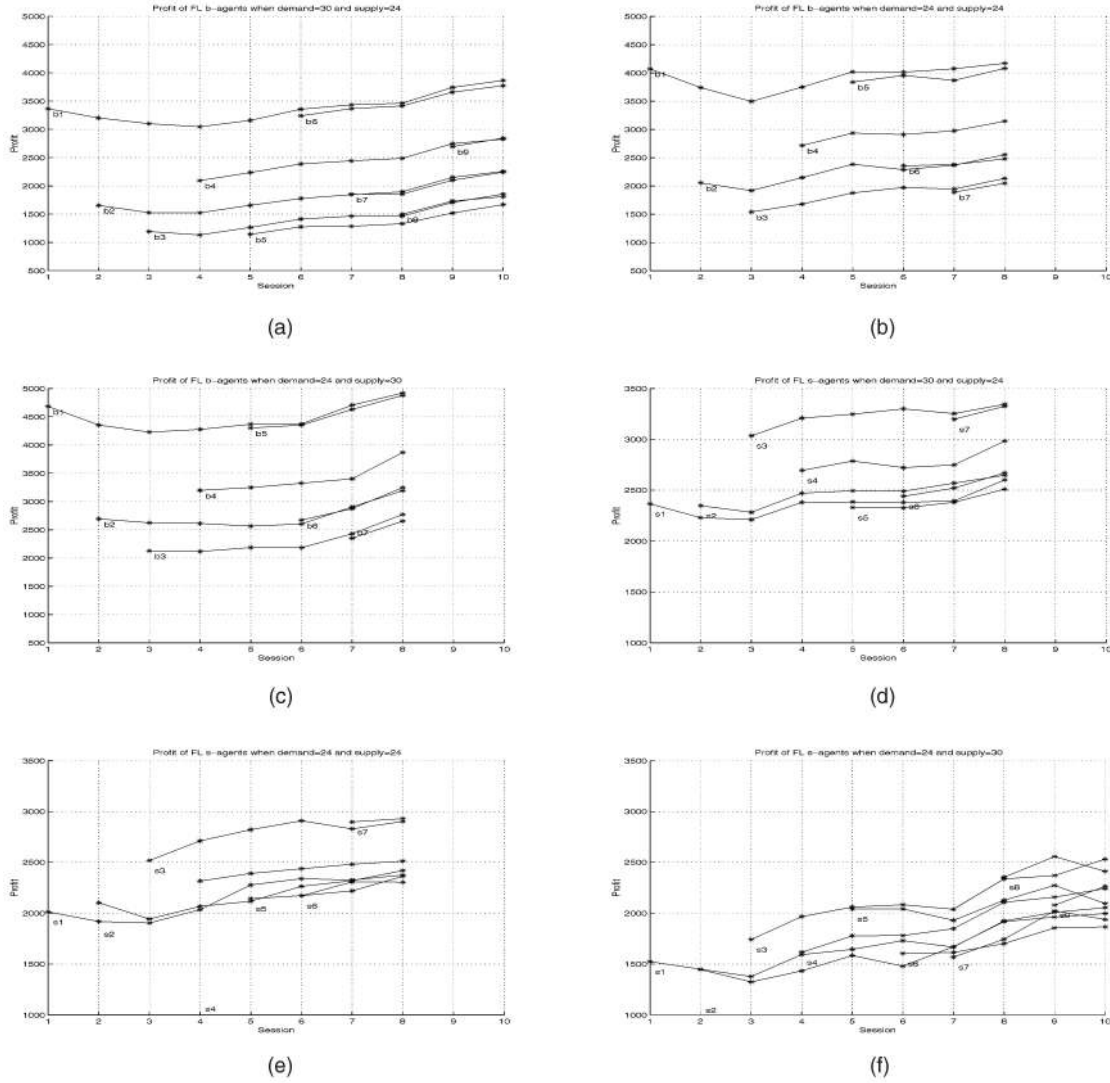


Fig. 11. Profits for A-FL-agents in different sessions. (a) Profits for A-FL b-agents when demand > supply. (number of b-agents is 10 and number of s-agents is 8). (b) Profits for A-FL b-agents when demand = supply. (number of b-agents is 8 and number of s-agents is 8). (c) Profits for A-FL b-agents when demand < supply. (number of b-agents is 8 and number of s-agents is 10). (d) Profits for A-FL s-agents when demand > supply. (number of b-agents is 10 and number of s-agents is 8). (e) Profits for A-FL s-agents when demand = supply. (number of b-agents is 8 and number of s-agents is 8). (f) Profits for A-FL s-agents when demand < supply. (number of b-agents is 8 and number of s-agents is 10).

increasing use of A-FL-agents. This investigation reveals that the profit of an individual A-FL-agent decreases at first and then increases steadily. We also show that with an increase in the number of A-FL-agents, the efficiency of the market is not significantly affected.

As well as being effective, we believe the FL strategy is practical for building autonomous agents for CDAs. The strategy we employ is intuitive and its embodiment in fuzzy rules means that it should be readily comprehensible to the agent’s owner (as have other similar applications of fuzzy rules [52], [53]). Moreover, the information required by the strategy can be readily obtained by monitoring market activities, such as the outstanding ask, the outstanding bid, and the accepted bids or asks in past transactions. In particular, this procedure does not require any information of the cost or valuation of other agents (cf. some of the approaches discussed in Section 3).

For the future, there are four main extensions required for our model. First, we need to further investigate the

setting of the learning rate for markets in which there can be abrupt changes in the supply and demand. Here, our hypothesis is that agents with a learning rate, which can rapidly respond to market changes, should perform best.

TABLE 5
Efficiency Statistics when Demand = Supply

Number of A-FL b-agent	Profit	Efficiency	Number of A-FL s-agent	Profit	Efficiency
0	29,463	96.92%	0	29,463	96.92%
1	29,267	96.27%	1	29,280	96.32%
2	29,560	97.24%	2	29,155	95.91%
3	29,575	97.29%	3	29,283	96.33%
4	29,446	96.86%	4	29,136	95.84%
5	29,103	95.73%	5	28,781	94.67%
6	29,086	95.68%	6	28,332	93.20%
7	29,028	95.49%	7	27,510	90.49%
8	28,449	93.58%	8	26,045	85.67%

TABLE 6
Efficiency Statistics when Demand < Supply

Number of A-FL b-agent	Profit	Efficiency	Number of A-FL s-agent	Profit	Efficiency
0	31,972	94.04%	0	31,972	94.04%
1	32,592	95.86%	1	31,776	93.46%
2	32,971	96.97%	2	31,751	93.39%
3	33,071	97.11%	3	32,140	94.53%
4	32,994	97.04%	4	32,020	94.18%
5	32,909	96.79%	5	32,201	94.71%
6	32,899	96.76%	6	32,081	94.36%
7	32,941	96.89%	7	31,971	94.03%
8	32,860	96.65%	8	30,843	90.71%
			9	29,454	86.63%
			10	28,267	83.14%

TABLE 7
Efficiency Statistics when Demand > Supply

Number of A-FL b-agent	Profit	Efficiency	Number of A-FL s-agent	Profit	Efficiency
0	31,444	86.86%	0	31,444	86.86%
1	31,504	87.03%	1	31,345	86.59%
2	31,480	86.96%	2	31,334	86.56%
3	31,448	86.87%	3	31,285	86.42%
4	31,448	86.87%	4	31,195	86.17%
5	31,316	86.51%	5	31,045	85.76%
6	31,029	85.72%	6	30,808	85.10%
7	30,998	85.63%	7	30,487	84.22%
8	30,961	85.53%	8	29,041	80.22%
9	30,598	84.52%			
10	30,332	83.79%			

Second, since the FL-strategy is only tested in a single unit protocol, the rules and some parameters of our reasoning model may need to be modified to deal with the multiple unit case. Third, we wish to experiment with the use of other fuzzy logic controllers, such as the conventional Mamdani controller [22]. The reason for this is that it may well be the case that another type of controller could improve the performance of our algorithms still further. Fourth, we need to develop tools that will assist users in adjusting the parameters of the fuzzy rules according to the their preferences. These are needed to make the reasoning model more widely accessible and usable.

ACKNOWLEDGMENTS

The authors would like to express their thanks to the anonymous reviewers for their insightful and helpful comments. The authors would also like to thank Xudong Luo (The University of Southampton), Chris Preist and Javier Esplugas Cuadrado (Hewlett-Packard Lab), and Philippe De Wilde (Imperial College) for their useful comments and suggestions. Thanks also to Cora B. Excelente-Toledo (The University of Southampton) for her help. The work described in this paper was partially supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (RGC Ref. No. CUHK4346/02E).

REFERENCES

[1] N. R. Jennings, "An Agent-Based Approach for Building Complex Software Systems," *Comm. ACM*, vol. 44, no. 4, pp. 35-41, 2001.

[2] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific Am.*, May 2001.

[3] F. Ygge, "Market-Oriented Programming and Its Application to Power Load Management," PhD thesis, Dept. of Computer Science, Lund Univ., Sweden, 1998.

[4] L. Crum, "University of Michigan Digital Library Project," *Comm. ACM*, vol. 38, no. 4, pp. 63-64, 1995.

[5] *The Grid: Blueprint for a New Computing Infrastructure*, I. Foster and C. Kesselman, eds. Morgan Kaufmann, 1998.

[6] P.R. Wurman, "Dynamic Pricing in the Virtual Marketplace," *IEEE Internet Computing*, vol. 5, pp. 36-42, Mar./Apr. 2001.

[7] P. Wurman, M. Wellman, and W. Walsh, "The Michigan Internet Auctionbot: A Configurable Auction Server for Human and Software Agents," *Proc. Second Int'l Conf. Autonomous Agents*, pp. 301-308, 1998.

[8] K. Binmore, *Fun and Games: A Text on Game Theory*. Lexington: D.C. Heath, 1992.

[9] D. Friedman and J. Rust, *The Double Auction Market: Institutions, Theories and Evidence*. Addison-Wesley, 1992.

[10] D. Friedman, "The Double Auction Market Institution: A Survey," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 3-25, Addison-Wesley, 1993.

[11] S. Gjerstad and J. Dickhaut, "Price Formation in Double Auction," *Games and Economic Behavior*, vol. 22, pp. 1-29, 1998.

[12] J.H. Kagel and W. Vogt, "Buyer's Bid Double Auctions: Preliminary Experimental Results," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 285-305, Addison-Wesley, 1991.

[13] M. Wooldridge and N.R. Jennings, "Intelligent Agents: Theory and Practice," *The Knowledge Eng. Rev.*, vol. 10, no. 2, pp. 115-152, 1995.

[14] R. Das, J. Hanson, J. Kephart, and G. Tesauro, "Agent-Human Interactions in the Continuous Double Auction," *Proc. 17th Int'l Joint Conf. Artificial Intelligence*, B. Nebel, ed., vol. 2, pp. 1169-1176, Aug. 2001.

[15] L.A. Zadeh, "Fuzzy Sets," *Information and Control*, vol. 8, pp. 338-353, 1965.

[16] H.A. Simon, *The Sciences of the Artificial*. Cambridge, Mass.: The MIT Press, 1997.

[17] T. Fraichard and P. Garnier, "Fuzzy Control to Drive Car-Like Vehicles," *Robotics and Autonomous Systems*, vol. 34, no. 1, pp. 1-22, 2001.

[18] J.F.F. Yao and J.S. Yao, "Fuzzy Decision Making for Medical Diagnosis Based on Fuzzy Number and Compositional Rule of Inference," *Fuzzy Sets and Systems*, vol. 120, no. 2, pp. 351-366, 2001.

[19] K.K. Tan and K.Z. Tang, "Vehicle Dispatching System Based on Taguchi-Tuned Fuzzy Rules," *European J. Operational Research*, vol. 128, no. 3, pp. 545-557, 2001.

[20] S. Mohammadi, I. Hassanzadeh, R. Mathur, and K. Patil, "A New Fuzzy Decision-Making Procedure Applied to Emergency Electric Power Distribution Scheduling," *Eng. Applications of Artificial Intelligence*, vol. 13, no. 6, pp. 731-740, 2000.

[21] M. Sugeno, "An Introductory Survey of Fuzzy Control," *Information Sciences*, vol. 36, pp. 59-83, 1985.

[22] H.-J. Zimmermann, *Fuzzy Set Theory and Its Applications*, chapter 11, pp. 203-240, Kluwer Academic Publishers, 1996.

[23] D. Dubois and H. Prade, "Operations on Fuzzy Numbers," *Int'l J. Systems Sciences*, vol. 9, pp. 613-626, 1978.

[24] P.P. Bonissone and K.S. Decker, "Selecting Uncertainty Calculi and Granularity: An Experiment in Trading-off Precision and Complexity," *Uncertainty in Artificial Intelligence*, L.N. Kanal and J.F. Lemmer, eds., pp. 217-247, Elsevier Science Publishers B.V., 1986.

[25] X. Luo, Y. Qiu, and J. Cai, "PILCF: A Hybrid Uncertain Reasoning Model," *Proc. Eighth Chinese Conf. Artificial Intelligence*, pp. 70-79, 1994.

[26] H. Bandemer and S. Gottwald, *Fuzzy Sets, Fuzzy Logic, Fuzzy Methods with Applications*. Wiley, 1995.

[27] R.R. Yager and D. Filev, *Essentials of Fuzzy Modeling and Control*. 1994.

[28] D. Easley and J. Ledyard, "Theories of Price Formation and Exchange in Double Oral Auctions," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 63-97, Addison-Wesley, 1993.

[29] J.M. Perloff, *Microeconomics*, chapter 2, pp. 29-32, Addison Wesley, 1998.

- [30] S.M. Ross, *Introduction to Probability Models*, chapter 5, pp. 201-247, Academic Press, 1989.
- [31] T.N. Cason and D. Friedman, "An Empirical Analysis of Price Formation in Double Auction Markets," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 253-283, Addison-Wesley, 1991.
- [32] A.R. Schotter, *Microeconomics: A Modern Approach*, chapter 14, pp. 442-445, 1994.
- [33] D. Gode and S. Sunder, "Lower Bounds for Efficiency of Surplus Extraction in Double Auctions," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 199-219, Addison-Wesley, 1993.
- [34] D.K. Gode and S. Sunder, "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *J. Political Economy*, vol. 101, no. 1, pp. 119-137, 1993.
- [35] C. Preist and M. van Tol, "Adaptive Agents in a Persistent Shout Double Auction," *Proc. First Int'l Conf. Information and Computation Economics*, pp. 11-18, 1998.
- [36] C. Preist, "Commodity Trading Using an Agent-Based Iterated Double Auction," *Proc. Third Int'l Conf. Autonomous Agents*, pp. 131-138, 1999.
- [37] D. Kreps, *Game Theory and Economic Modeling*. Oxford Univ. Press, 1990.
- [38] N.R. Jennings, P. Faratin, A.R. Lomuscio et al., "Automated Negotiation: Prospects, Methods and Challenges," *Int'l J. Group Decision and Negotiation*, vol. 10, no. 2, pp. 199-215, 2001.
- [39] J. Vidal and E. Durfee, "The Impact of Nested Agent Models in an Information Economy," *Proc. Second Int'l Conf. Multi-Agent Systems*, pp. 377-384, 1996.
- [40] S. Park, E. Durfee, and W. Birmingham, "An Adaptive Agent Bidding Strategy Based on Stochastic Modeling," *Proc. Third Int'l Conf. Autonomous Agents*, pp. 147-153, 1999.
- [41] L. Badae, "Learning Trading Rules with Inductive Logic Programming," *Machine Learning: ECML 2000*, R. Lopez de Mantaras and E. Plaza, eds., pp. 39-46, 2000.
- [42] P. Faratin, C. Sierra, and N.R. Jennings, "Using Similarity Criteria to Make Negotiation Trade-Offs," *Proc. Fourth Int'l Conf. Multi-Agent Systems*, pp. 119-126, 2000.
- [43] R. Kowalczyk and V. Bui, "On Fuzzy E-Negotiation Agents: Autonomous Negotiation with Incomplete and Imprecise Information," *Proc. DEXA e-Negotiation Workshop*, 2000.
- [44] R. Kowalczyk, "On Negotiation as a Distributed Fuzzy Constraint Satisfaction Problem," *Proc. DEXA e-Negotiation Workshop*, pp. 631-637, 2000.
- [45] X. Luo, N.R. Jennings, N. Shadbolt, H.F. Leung, and J.H.M. Lee, "A Fuzzy Constraint Based Model for Bilateral, Multi-Issue Negotiations in Semi-Competitive Environments," *Artificial Intelligence J.*, vol. 148, nos. 1-2, 2003.
- [46] X. Luo, "A Study of Information Sharing of Heterogenous Reasoning Models in Multi-Agent Environment," PhD thesis, The Univ. New England, Australia, 1998.
- [47] X. Luo, C. Zhang, and H.F. Leung, "Information Sharing between Heterogeneous Uncertain Reasoning Models in a Multi-Agent Environment: A Case Study," *Int'l J. Approximate Reasoning*, vol. 27, no. 1, pp. 27-59, 2001.
- [48] J. Yen, "Fuzzy Logic—A Modern Perspective," *IEEE Trans. Knowledge and Data Eng.*, vol. 11, no. 1, pp. 153-165, Jan./Feb. 1999.
- [49] E. Gimenez-Funes, L. Godo, J. Rodriguez-Aguilar et al., "Designing Bidding Strategies for Trading Agents in Electronic Auctions," *Proc. Third Int'l Conf. Multi-Agent Systems*, pp. 136-143, 1998.
- [50] N. Matos and C. Sierra, "Evolutionary Computing and Negotiating Agents," *Agent-Mediated Electronic Commerce*, vol. 1571, pp. 126-150, 1998.
- [51] D. Zeng and K. Sycara, "Bayesian Learning in Negotiation," *Int'l J. Human-Computer Studies*, vol. 48, pp. 125-141, 1998.
- [52] Z.A. Sosnowski, "Comments on Fuzzyspell: A Large-Scale Expert Systems Shell Using Fuzzy Logic for Uncertainty Reasoning," *IEEE Trans. Fuzzy Systems*, vol. 8, no. 6, pp. 817-820, 2000. Authors' reply, p. 821, 2000.
- [53] Y. Yam and L.T. Koczy, "Representing Membership Functions as Points in High-Dimensional Spaces for Fuzzy Interpolation and Extrapolation," *IEEE Trans. Fuzzy Systems*, vol. 8, no. 6, pp. 761-772, 2000.



Minghua He received the master's degree in computer science from Southwest China Normal University in China in 1999. She is a PhD candidate in the Department of Electronics and Computer Science at the University of Southampton, United Kingdom. Her research interests center around the study of artificial intelligence, especially the area of agent-based automated negotiation for E-commerce.



Ho-fung Leung received the BSc (Hon.) and MPhil degrees in computer science from The Chinese University of Hong Kong. He received the PhD degree from the Imperial College of Science, Technology, and Medicine in 1992. He is an associate professor in the Department of Computer Science and Engineering at The Chinese University of Hong Kong. Dr. Leung joined the Department of Computer Science of The Chinese University of Hong Kong in 1992.

His current research interests include constraint satisfaction, multiagent systems, and logic programming. He is a participating member of the Centre for the Advancement of E-Commerce Technologies (AECT) of the Engineering Faculty.



Nicholas R. Jennings is a professor of computer science in the Department of Electronics and Computer Science at Southampton University, where he carries out basic and applied research in agent-based computing. He is head of the Intelligence, Agents, Multimedia Group and is also the chief scientific officer for lost wax. Professor Jennings helped pioneer the use of agent-based techniques for real-world applications and he has also made theoretical contributions in the areas of automated negotiation and auctions, cooperative problem solving, and socially rational decision making. He has published more than 150 articles on various facets of agent-based computing, holds two patents (three more pending), written one monograph, and coedited five books. He is in the top 250 most cited computer scientists (according to the citeseer digital library). He was the recipient of the Computers and Thought Award (the premier award for a young AI scientist) in 1999 for his contributions to practical agent architectures and applications of multiagent systems (this is the first time in the Award's 30 year history that it has been given to someone based in Europe) and the recipient of an IEE Achievement Medal in 2000 for his work on agent-based computing.

► For more information on this or any computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.