ELSEVIER

Decision Support Systems 33 (2002) 335–347

Decision Support Systems

www.elsevier.com/locate/dsw

Cooperation in multi-agent bidding

D.J. Wu*, Yanjun Sun

LeBow College of Business, Drexel University, 101 North 33rd Street, Academic Building, Philadelphia, PA 19104, USA

Abstract

A fundamental problem in business-to-business exchanges is the efficient design of mechanisms to promote cooperation and coordination of various self-interested agents. We study the behavior of artificial agents in a bidding and contracting framework [Eur. J. Oper. Res. (2002); D.J. Wu, P. Kleindorfer, J.E. Zhang, Integrating Contracting and Spot Procurement with Capacity Options, Working Paper, Department of Operations and Information Management, The Wharton School, University of Pennsylvania, 2001]. In this framework, there is a long-term contract market as well as a backstop spot market. Seller agents bid their contract offers in terms of price and capacity via an electronic bulletin board, while Buyer agents decide how much to contract with Sellers and how much to shop from the spot market. This two-tiered market has been modeled [Eur. J. Oper. Res. (2002); D.J. Wu, P. Kleindorfer, J.E. Zhang, Integrating Contracting and Spot Procurement with Capacity Options, Working Paper, Department of Operations and Information Management, The Wharton School, University of Pennsylvania, 2001] as a von-Stackelberg game with Seller agents as leaders, and the necessary and sufficient conditions for the existence of market equilibrium are given. What happens if the resulting equilibrium is noncooperative and Pareto dominated by some nonequilibrium bidding? What happens if there are multiple equilibria, some Pareto dominated by others, then which will be selected? What happens when there is no equilibrium? The goal of this paper is to study equilibrium and disequilibrium behavior of artificial agents in such systems, and explore the efficient design of mechanisms to promote cooperation and coordination of self-interested artificial agents.

Keywords: Artificial agents; Bidding; Capacity options; (Dis)equilibrium behavior

1. Introduction

A fundamental problem in business-to-business exchanges is the efficient design of mechanisms to promote cooperation and coordination of various self-interested agents [35]. We study the behavior of artificial agents in a bidding and contracting framework in Refs. [40,41], hereafter cited as the WKZ papers. In the WKZ framework, there is assumed to be a contract market (which might be thought of as the "month ahead" market), in which Sellers and Buyers interact through an electronic bulletin board, posting bids and offers until agreement has been reached. Capacity not committed through this contracting market is assumed to be offered on the spot market, but such capacity may go unused because of the risk of not finding customers or transportation capacity at the last minute ("on the day"). Buyers face another type of risk for demand not contracted for in the bilateral market, namely price volatility in the spot market. Such price volatility can be quite severe and has caused Buyers (for example, in the electric power market) to pay close attention to the proper balance in their supply

* Corresponding author. Tel.: +1-215-895-2121; fax: +1-215-895-2891.

E-mail addresses: wudj@drexel.edu (D.J. Wu), ys45@drexel.edu (Y. Sun).

PII: S0167-9236(02)00020-9

portfolio between long-term contracting and spot purchases. This two-tiered market has been modeled in the WKZ papers as a von-Stackelberg game with Seller agents as leaders.

The WKZ papers show the optimal contracting strategies for the Buyers and the optimal bidding and auction strategies for the Sellers. WKZ [41] also characterize existence conditions and the structure of market equilibrium for the associated competitive game among sellers, under the assumption that sellers know buyers' demand functions. However, if these conditions are violated, no equilibrium exists. In practical Internet auctions, the conditions as required in the WKZ theoretical framework could well be violated (examples are given below). What happens in such a case is precisely the focus of this paper.

Here, we are interested in using artificial agents [15,21,24,36] to explore the (dis)equilibrium behavior of such a dynamic bidding system. In doing so, we model each selling artificial agent as a learning automaton [24], where the learning mechanism is characterized by a genetic algorithm [14]. Each buying artificial agent is embedded with an optimal contracting strategy as derived in the WKZ papers; hence, buyer agents do not learn since they are already using the optimal strategies given available bids on the electronic market.

The goals of this study are the following. First, we investigate whether or not artificial agents are able to discover equilibrium strategies when equilibrium exists. Second, we investigate to see if the agents can discover good and effective (e.g., cooperative) bidding strategies when playing repeated nonlinear games when the one-shot equilibrium is noncooperative or when there does not exist any equilibrium. Finally, we explore the emergence of trust by studying the conditions on problems and bidding mechanisms that induce cooperation when the above nonlinear game is repeated over time.

The rest of the paper is organized as follows. Section 2 provides a brief literature review. Section 3 describes our bidding models. Section 4 reports the results of baseline computational experiments on multi-agent bidding as well as broader experiments varying a number of underlying parameters. Section 5 investigates nonmyopic bidding strategies (history-dependent "if-then-else" rules, defined in Section 5) and studies the emergence of cooperation in a trust framework. Section 6 summarizes our findings and discusses future research.

2. Literature review

Three streams of research are relevant to our work: (1) models, analysis and human lab experiments on the design of Internet business-to-business exchanges for capital-intensive industries such as the electric power sector, (2) agent-mediated electronic commerce, and (3) cooperation-based social trust. None of the above three streams of work has been combined, however. In this paper, we integrate these separate streams into a prototype Decision Support System that stems from a real world problem.

2.1. Models and analysis

The WKZ papers set up the theoretical framework for the optimal bidding and contracting for capital-intensive industries, which we summarize below. The framework models the strategic interaction of long-term contracting and spot market transactions between Sellers and Buyers for capital-intensive goods. Sellers and Buyers may either contract for delivery in advance (the "contracting" option) or they may sell and buy some or all of their output/input in a spot market. Contract pricing involves both a reservation fee (s) per unit of capacity and an execution fee (g) per unit of output if capacity is called. The key question addressed is the structure of the optimal portfolios of contracting and spot market transactions for these Sellers and Buyers, and the pricing thereof in market equilibrium when such equilibrium exists. The WKZ papers show that when Sellers properly anticipate demands to their bids, bidding a contract execution fee equal to variable cost (b) dominates all other bidding strategies yielding the same contract output. The optimal capacity reservation fees (s) are determined by Sellers to trade-off the risk of underutilized capacity against unit capacity costs. Buyers' optimal portfolios are shown to follow a merit order (or greedy) shopping rule, under which contracts are signed following an index, which is an increasing function of the Seller's reservation cost and execution cost. The index has the following structure: $s+G(g)$, where $G(g)=E\{\min[P_s,g]\}$. $G(g)$ is called the "effective price" at g. It represents the expected value of the price paid by a Buyer who has purchased a capacity option, and reflects the fact that the Buyer will use the option when $g < P_s$, the spot price, but otherwise will use the spot market rather than exercising the option contract. Shopping in order of the index $s + G(g)$ says that the Buyer at the margin evaluates which Seller to do business with by considering the full price of a unit of output purchased under an option. This full price is clearly the option fee s plus the effective price at g, $G(g)$, when owning the option. The Buyer fills his order book or procurement portfolio up to anticipated willingness-to-pay for additional units of contract output. Beyond this, the Buyer plans on purchasing output from the spot market. This implies that the index of the final contract signed by the Buyer always satisfies $s + G(g) < E\{P_s\}$, the mean of the spot price.

Existence conditions and the structure of market equilibrium are characterized in WKZ [41] for the associated competitive game among Sellers, under the assumption that they know Buyers' demand functions. If these conditions are violated, there does not exist any equilibrium in the WKZ bidding game. In practice, nonequilibrium is a real possibility as we will see below. What happens in disequilibrium is therefore of considerable interest.

We investigate (dis)equilibrium behavior using artificial agents. An alternative and complementary approach would be controlled experiments [19]. Both approaches attempt to discover how the rules of the game, regarding information and decision rights, affect the outcome. An excellent example of the use of controlled experiments in this regard is the work of the Arizona School (e.g., Refs. [25,28]) investigating electric power markets. Such markets are quite complex, owing to the real-time balance requirements of electricity, and experimental studies show the sensitivity of market outcomes to very small changes in the rules governing bidding and strategic interactions. It should be noted, however, that none of this work has dealt with the problem of disequilibrium behavior. One obvious reason is that such behavior typically involves cycling among various focal outcomes (we will see precisely this type of cycling in the problem analyzed below), and it would be rather expensive to track such behavior experimentally. Thus, for the question of interest in this paper, the use of artificial agents seems especially appropriate.

2.2. Artificial agents

The study of artificial agents [17] and their applications in electronic commerce [e.g., DSS special issue on agents, vol. 28, 2000] have been growing rapidly in recent years. In that stream of work, topics pertaining to agent-mediated electronic commerce, such as auction bots [16,26,27,44], exchange agents [29,30], and shopping and pricing agents [11,12], are relevant to the work we report here. What distinguishes the present work from previous work is the interest here in market outcomes under (dis)equilibrium conditions, and this for a market having a structure sufficiently realistic to capture the essential characteristics of many emerging B2B exchanges. For example, the work in Refs. [11,12] is concerned only with the consumer side of e-Markets, and does not base its work on any market-level (e.g., equilibrium) outcomes. Similarly, the work in Refs. [16,29,30,44] is valuable for the study of general Internet-based auctions, and associated infrastructure or platforms, but does not model in any way the supply side of those participating/competing in the auctions. However, it is precisely the interaction of supply technology and capacity conditions with buyer demand structure that is fundamental to business-to-business exchange outcomes. Finally, the work in Refs. [26,27] considered bundled auctions, which are important aspects for many consumer-oriented problems, but are not relevant to the large majority of B2B exchanges which are concerned with well-defined commodities, separately priced, and typically not with bundled goods. None of these, by the way, deals with out of equilibrium behavior or disequilibrium behavior. By contrast, the present paper examines a class of problems of direct relevance to B2B exchanges for capital-intensive production, is based on a solid theoretical structure describing rational market outcomes, and describes and interprets both equilibrium and nonequilibrium market behavior.

Particularly relevant to our work are the applications of agent technology to the electric power sector such as pricing and investment [37,38], agents negotiating for load balancing [4], energy management [1], power load management [47], planning of power transmission expansion [46], and distributed transmission cost allocation [45]. This work covers a number of different problems in electric power, from investment problems, to optimizing power flows and demand-side bidding. None of this work, however, treats the fundamental problems of multiple markets (contracting and spot markets) that are the hallmark of optimal supply management by generators/sellers in this market. This is one of the applications to which the general framework developed in WKZ [40,41], and analyzed here, is applicable.

2.3. Social trust

The issue of social trust [32] is extremely important in electronic communities since it is difficult to figure out who you can trust in the electronic marketplace [31,34]. Surprisingly, the notion of social trust has never been agreed upon among researchers, such as philosophers or economists [5,6,18]. This implies a fundamental ambiguity about any computational approach to contexts involving trust [23], since such contexts require a legitimate model or set of conditions under which trust may be expected to emerge from the interactions of rational agents.

In this paper, we view trust as cooperation via reciprocity, as in the computational economics literature (e.g., Refs. [5,9,13]). Most of this work is based on lab experiments of human beings playing various simple trust games [7–9,13,33], and these experiments underscore the importance of outcome equity, transparency and repetition for the emergence of trust/cooperation. As will be shown, particularly relevant to our work is the prisoner's dilemma game, which is one of the trust games that has been extensively studied using both human agents and artificial agents (e.g., Refs. [2,3,24]). The intellectual links between our work and this work lie in exploring the emergence of social trust in the form of cooperative behavior (in a repeated game setting). Our work reported here describes an original bidding game that arises from our earlier theoretical framework characterizing a class of real world applications, and focuses on both equilibrium and nonequilibrium behavior. (The latter has not been studied in the iterated prisoner's dilemma literature.) In particular, we test reciprocity as a device for contract enforcement and for social trust building in a distributed environment (i.e., with no centralization) when the resulting equilibrium is noncooperative or empty. We emphasize here that we are interested in the emergence of social trust/cooperation; we are not here interested in either trusting technology or in user acceptance of information technology [32,39,48].

3. The bidding game

In this section, we describe our bidding game, discuss some interesting bidding strategies, define the normal form of the bidding game, give two examples (one with equilibrium and the other with out), and justify why an agent-based approach is needed.

3.1. Preliminaries and notation

There is a single Buyer and there are N Sellers. There are two markets, a contract market in which Sellers can precommit capacity via capacity options to the Buyer, and a spot market in which Sellers can sell residual capacity (with some risk) and the Buyer can buy additional output. Each Seller i maximizes its expected profit $E\pi_i$ by bidding a contract price $x_i = s_i + G(g_i)$ anticipating the Buyer's optimal contracting strategy $Q_i(x)$, where the competing bids are $x = (x_1, x_2, \ldots, x_N)$. It has been shown in WKZ [40,41] that the optimal contracting strategy of the Buyer, $Q_i(x)$, follows a merit order, i.e., the Buyer shops along a list from the lowest bid to the highest bid until all the Buyer's demand has been satisfied. Each Seller has a capacity limit K_i (thus $\sum_i K_i$ is the overall capacity of all Sellers), a technology index that is a function of the variable cost b_i, and a capacity cost β_i. Denote $c = (c_1, c_2, \ldots, c_N)$ and $K = (K_1, K_2, \ldots, K_N)$. As is standard in the economics literature, we assume there is a linear contract demand,[1] $D(p) = (a - hp)^+$, where $y^+ = \max\{y, 0\}$. Following WKZ [40,41], we adopt the following bid-tie allocation mechanism: If there is a tie in bids among any subset of Sellers, then the Buyer's total demand for that subset of Sellers is allocated to the Sellers in proportion to their respective bid capacities. (This a commonly used allocation mechanism in practice.)

[1] Note this does not imply any linear assumption of the Buyer's total demand. Indeed, it is kinked (i.e., nonlinear) under any spot market price distribution, as shown in Ref. [40].

Seller i's problem is to maximize expected profit,

$$x_i Q_i(x) + c_i[K_i - Q_i(x)] - [\beta_i + G(b_i)]K_i$$

where the first term is Seller i's revenue from the contract market, the second term is the revenue from the spot market, and the third term represents cost. This formula can be rewritten as

$$(x_i - c_i)Q_i(x) + [c_i - \beta_i - G(b_i)]K_i \tag{1}$$

Notice that the second term in Eq. (1) is constant with respect to the decision vector, x, and therefore, Seller i's problem can be simplified as maximize

$$E\pi_i = (x_i - c_i)Q_i(x) \tag{2}$$

3.2. Interesting bidding behaviors

We now consider some of the interesting and useful bidding strategies that might be taken by the Sellers.

The *Monopoly Bidder* (v_i). This is defined as Seller i's bid if this seller is the only supplier in the contract market, i.e.,

$$v_i = \arg\max_{v_i} [(v_i - c_i)D(v_i)]^+$$
$$= \arg\max_{v_i} [(v_i - c_i)(a - hv_i)]^+ = (a + hc_i)/(2h).$$

The Cooperative Bidder ($u_i = u$). This is defined as all the Sellers forming a cartel or a team or an oligopoly and bidding a price that maximizes the total profit of all Sellers as a whole, i.e., $u_i = u$ is the cooperatively determined (denoted as "c"), uniform price solving the following problem:

$$\operatorname*{Max}_u \sum_i E\pi_i$$

Subject to:

$E\pi_i \geq 0$	for all i (individual rationality constraint)
$0 \leq Q_i(u,u,u) \leq K_i$	for all i (capacity constraint)

The solution to the above problem is:

$$u = \left(a + \sum_i c_i K_i / \sum_i K_i\right)/(1+h) \text{ if } u$$
$$> \max\left\{\max\{c_i\}, \left(\left(a - \sum_i K_i\right)/h\right)^+\right\}$$

else no cooperative solution exists.

The Opportunistic Bidder (x_i). Also called the myopic bidder (denote as "m") or the noncooperative bidder. Given the other players' bids, x_{-i}, Seller i bids to maximize its own profit

$$x_i = \arg\max_{x_i} E\pi_i[x_i \mid x_{-i}]$$

where, for any vector $x = (x_1, x_2, \ldots, x_N)$, $x_{-\mathrm{i}}$ is the vector of i-exclusive components of x, e.g., $x_{-1} = (x_2, \ldots, x_N)$.

The Random Bidder (r_i). Seller i bids randomly (denoted as "r") between the lower bound and the upper bound according to the uniform distribution, i.e., $r_i = \mathrm{Rand}(c_i, \max(u, v_i))$. Note here that given the profit function (2), c_i is the minimum feasible bid (this is obviously the least any rational seller should bid, otherwise the seller can make more profit per unit of capacity by participating on the spot market rather than on the contract market).

3.3. Normal form bidding game

We now formally define the bidding game among the Sellers. As will be seen, these games provide a rich platform for study of both equilibrium and non-equilibrium behavior. The players are Sellers i, $i = 1, \ldots, N$. The strategy space for each Seller i is any integer between $[c_i, \max(u, v_i)]$. The Payoff function is computed via Eq. (2). Note that if there is a bid-tie, the capacity is allocated according to the WKZ bid-tie allocation mechanism. Note also that for this one-shot normal form game, there will always exist at least one mixed strategy Nash equilibrium [10]; however, there may not exist any pure strategy Nash equilibrium. (Here, we are only interested in pure strategy equilibria not mixed strategy equilibria.) We have implemented a straightforward algorithm to check whether this normal form game exists any pure strategy Nash equilibrium, and, if so, our algorithm can find it.[2]

3.4. Numerical examples

We now give two examples that will be further analyzed in the following sections.

Example 1. The first example corresponds to row 6, column 4 of Table 1, which has parameter settings of c=(10,10,18), K=(40,40,30), and the contract demand function is $D(p)=(100-p)^{+}$. The unique Nash equilibrium for this one-shot game is the bid vector of x=(18,18,19). However, for this very same game, there exists a cooperative feasible bid vector, (56, 56,56) or $u=56$, that is not a Nash equilibrium, yet Pareto dominates the noncooperative Nash equilibrium (18,18,19). It is not a Nash equilibrium since, when facing the bid (56,56,56), Seller 3 has an incentive to bid slightly less than the other players, say 55, in order to contract all its capacity with the Buyer (to profit more) rather than splitting the Buyer demand with Sellers 1 and 2 (which would happen under the assumed proportional bid-tie allocation mechanism). This bidding game is similar to the Prisoner's Dilemma. The problem of interest is whether artificial agents learn to cooperate when the game is played repeatedly.

Example 2. The second example corresponds to row 7, column 4 of Table 1, which has the parameters of c=(10,12,14) and K=(40,30,20), with the same demand function as in Example 1. There does not exist any pure strategy Nash equilibrium for the one-shot game. Note that there is nothing odd about the above problem parameters. This and other examples (see Table 1 as well as additional examples in Ref. [41]) suggest the need to understand the "rational" bidding behavior of agents in problems such as these and the consequences of such bidding when there does not exist any equilibrium.

[2] This algorithm basically computes the best-response of each player and checks to see if there exists any outcome that constitutes the best-responses of all players (i.e., the Nash equilibrium). If no such outcome can be found, then there does not exist any pure strategy Nash equilibrium, since there is at least one player who has the incentive to deviate given other players' bids.

Given the complexity of such behavior in a repeated game setting, this seems an ideal problem setting to study via artificial agents. This will be the approach we pursue in what follows. We consider only strategies based on price here (as in WKZ), but the same approach could be used to evaluate capacity based bid strategies (using either the Cournot equilibrium, or the Cournot best-response functions [10], to evaluate profits at a specific vector of capacity bids; see Ref. [42] for some initial results).

In what follows, we first investigate two mechanisms for artificial Seller agents: myopic learning (history-independent, Section 4) and nonmyopic learning (history-dependent "if-then-else" rules, Section 5) as defined below, and finally, we briefly interpret agents' behavior in the context of social trust.

4. Myopic bidding

In myopic bidding, all Seller agents use the same rule: the myopic rule ("m") (e.g., Ref. [20]). Basically, what this rule or heuristic says is the following: at the beginning of the game, randomly choose a bid from the feasible strategy space. Then, for the next period, choose the bid that maximizes the agent's payoff for his/her current time period, assuming the other two agents stick with their previous bids.[3] This myopic rule is formally given as follows:

$$x_i(0) = \mathrm{Rand}(c_i, \max(u, v_i))$$

$$x_i(t) = \arg\max_{x_i(t)} E\pi_i[x_i(t) \mid x_{-i}(t-1)] \qquad (3)$$

Here, we assume each Seller can only remember what happened last time, i.e., memory size is 1 (which is standard in the agent literature), so that current strategies cannot depend on previous strategies used by other players prior to the most recent period. (We leave the impact of memory size for a separate study; for initial results, see Ref. [48].)

We want to study the behavior of artificial agents when playing the dynamic game of the above exam-

[3] Thus, this rule calls for players to play their "best-response" function strategy in the usual game-theoretic sense, as defined above.

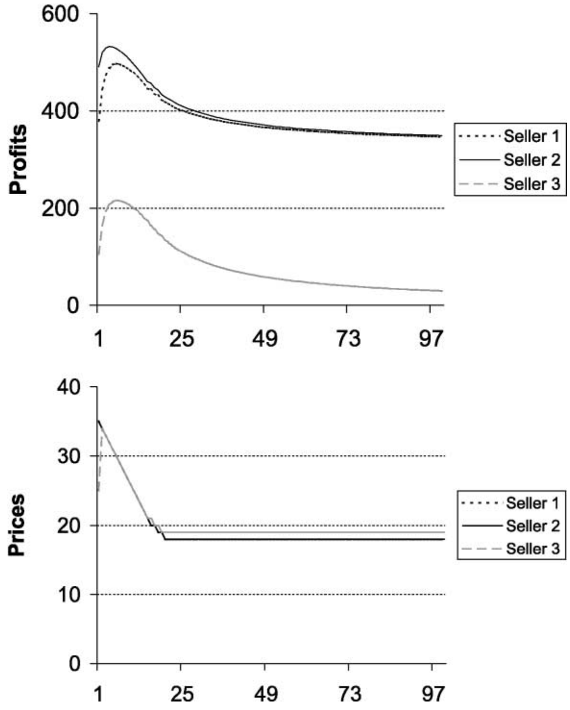

Fig. 1. Dynamic bidding for Example 1: profit and price over time for $c=(10,10,18)$, $K=(40,40,30)$.

ples,[4] e.g., Example 1 where $c=(10,10,18)$ and $K=(40,40,30)$. Under such a myopic (best-response) rule, not surprisingly, artificial agents quickly converge to the noncooperative equilibrium (18,18,19) predicted by our game-theoretic type of analysis. The dynamics of profit and agent price bidding for this simple, benchmark setting is shown in Fig. 1.

In order to test the stability and validity of our findings, we designed a broader experiment to test various competing technologies and contract demand functions. For some of these, equilibrium exists and for some it does not. Table 1 summarizes the experimental design and the results of a game-theoretical analysis of the one-shot three-seller bidding game.[5] When the equilibrium exists, and given that we are using adjustment processes that mimic best-response strategies, we should not be surprised if the myopic adjustment process converges always to the noncooperative equilibrium, indeed, this behavior can be explained by the "backward induction" argument in game theory [10]. When there does not exist any equilibrium, the agents present the expected oscillating behavior (i.e., myopic bidding agents cannot discover an equilibrium in repeated bidding). These are, in fact, what occur, as shown in Table 2, which reports the results using myopic bidding, where the entry "no equilibrium" means no convergence occurs for the indicated bidding game.

A point implicit in Table 2 is that the cooperative outcome (which is different for each cell in this design) did not emerge in any of the test cases. The reason is that the myopic rule used provides incentives for agents to undercut their fellow sellers in order to gain a higher market share and profit through a slight price reduction. The result is that cycling is observed as the artificial agents engage in a "price war" (no cooperation), as plotted in Fig. 2.

Artificial agents engage in "dog-fighting" strategies that take advantage of other players' previous bids, but they never manage to exit from their cycling behavior, and long-term profits are therefore not maximized. None of this is very surprising given the structure of the myopic strategies assumed. However, what is interesting is that using myopic bidding, where there exist multiple Nash equilibria, the agents are able to select the Nash equilibrium that Pareto dominates the others. This is true for all cases with multiple Nash equilibria in our experimental design, as shown in row 2, columns 2, 3 and 4 of Table 2.

5. Nonmyopic bidding

Since myopic bidding does not lead to cooperation, we now investigate whether nonmyopic (or history-dependent) price bidding will do better. If so, we are

[4] In myopic bidding, GA is used to search for the optimal bid by solving Eq. (3). For the simple profit function (2), other search techniques have been used to benchmark with GA (not surprisingly, both found the optimal solutions in these cases). Of course, GA could be used for considerably more complex profit functions than those illustrated here, including profits that are the result of enterprise simulation models, which would preclude the use of many analytic techniques. In any case, the focus in this first example is on system dynamics, rather than on comparing various optimization methods.

[5] Basically, the table shows a two-factor statistical design to cover various cases of the parameter settings, and the results of the game-theoretical analysis are obtained using the algorithm described in footnote 2.

Table 1
Theoretical analysis of the one-shot three-seller bidding game

	$D=(77-p)^+$	$D=(115-p)^+$	$D=(100-p)^+$
$c_i=(13,13,13)$, $K_i=(30,30,30)$	(16,16,16), (15,15,15), (14,14,14)	(30,30,30), (29,29,29)	(21,21,21), (20,20,20)
$c_i=(11,21,21)$, $K_i=(44,23,23)$	(21,22,22)	no equilibrium	(23,24,24)
$c_i=(16,16,25)$, $K_i=(34,34,22)$	no equilibrium	(31,31,31)	no equilibrium
$c_i=(7,12,17)$, $K_i=(45,26,19)$	no equilibrium	no equilibrium	no equilibrium
$c_i=(10,10,18)$, $K_i=(40,40,30)$	no equilibrium	no equilibrium	(18,18,19)
$c_i=(10,12,14)$, $K_i=(40,30,20)$	no equilibrium	no equilibrium	no equilibrium

Table 2
Agents' repeated myopic bidding results. In agent repeated bidding, no cooperation emerged in any of the test cases

	$D=(77-p)^+$	$D=(115-p)^+$	$D=(100-p)^+$
$c_i=(13,13,13)$, $K_i=(30,30,30)$	(16,16,16)	(30,30,30)	(21,21,21)
$c_i=(11,21,21)$, $K_i=(44,23,23)$	(21,22,22)	no equilibrium	(23,24,24)
$c_i=(16,16,25)$, $K_i=(34,34,22)$	no equilibrium	(31,31,31)	no equilibrium
$c_i=(7,12,17)$, $K_i=(45,26,19)$	no equilibrium	no equilibrium	no equilibrium
$c_i=(10,10,18)$, $K_i=(40,40,30)$	no equilibrium	no equilibrium	(18,18,19)
$c_i=(10,12,14)$, $K_i=(40,30,20)$	no equilibrium	no equilibrium	no equilibrium

interested in finding out the conditions under which cooperation and long-run profit maximization occur.

5.1. Rule definitions

The agent strategies used are a population of "if-then-else" rules or "condition/action" pairs that maps the history (previous time period $t-1$) of other players' bidding behavior to the agent's current action (for time period t). Each rule specifies an action (how much to bid for the current time period) to be taken if the condition is satisfied. The rules are coded in the following way.[6] Assume each agent can take an action by either bid randomly ("r"), cooperatively ("c") or myopically ("m") as defined previously. Assume each agent has a memory size of 1. Thus, a nine-bit string completely specifies the action of an agent, with each bit can take three possible values from {r,c,m}. In particular, the string has the structure of "rr, rc, rm, cr, cc, cm, mr, mc, mm", with each bit corresponding to one of the nine possible scenarios from the previous time period. For example, the sixth bit "cm" corresponds to the situation when Seller j bids cooperatively ("c") and Seller k bids myopically ("m") during the previous time period $(t-1)$. The value in this bit specifies the action the agent (Seller i) will take in the current time period (t).[7] Examples of rules using this coding strategy are given below. Each agent-rule (three-value string) in each generation plays a round-

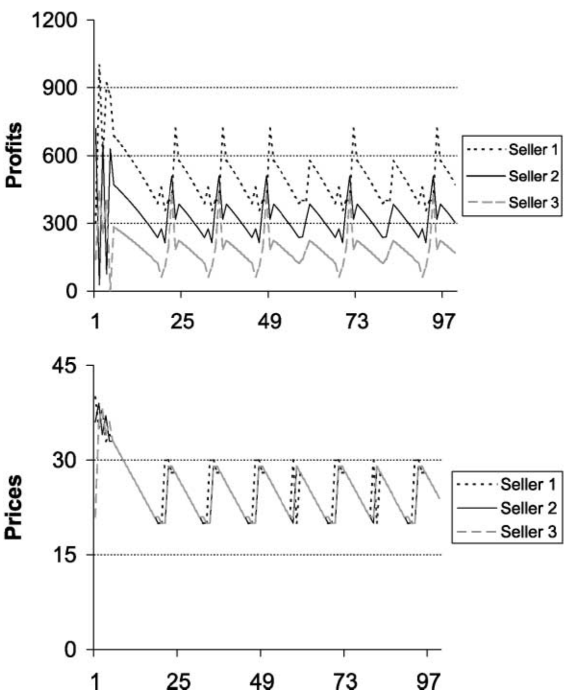

Fig. 2. Dynamic bidding for Example 2: profit and price over time for $c=(10,12,14)$, $K=(40,30,20)$.

[6] This simple rule coding strategy is standard in the genetic algorithms literature when programming agents to play games.

[7] Therefore, the total rule search space for Seller i would be 3^9 for this particular example. The search space would grow exponentially as the size of the memory increases, for example, when the memory size is 2, then the rule space would be 3^{81} using the above coding strategy. It is straightforward to transfer this rule representation to computer implementations such as in binary coding or in gray-coding, obviously the size of the rule space depends on specific computer implementation. We focus on the insights discovered by artificial agents, but not on comparing efficiency of various coding strategies.

robin tournament against each and every other agent-rule with a prespecified population size, say, 70. This means that each rule for each agent has to play against $70 \times 70 = 4900$ combinations of rules of the other two agents, with fitness of the rule for this agent being the total profits achieved over the 4900 games played in this tournament. Standard genetic algorithm operators are used to evolve rules between generations. Each bidding game is repeated 100 times. We are particularly interested in the performance of the following rules. These rules are random (R), cooperative (C), noncooperative (N), and adaptive-learning (L), which are formally defined as the following. Note that the string in quotes, e.g., "r, r, r, r, r, r, r, r, r", is the GA rule representation using the above-described coding strategy. For any Seller i, given the bids of others (Seller j and Seller k):

Random rule: R = "r, r, r, r, r, r, r, r, r". $x_i(t) = \text{Rand}(c_i, \max(u, v_i))$ (uniformly bid an integer number from the feasible strategy space no matter what the others bid);

Cooperative rule: C = "c, c, c, c, c, c, c, c, c". $x_i(t) = u_i = u$ (always bid the cooperative price $u_i = u$, no matter what the others bid. This is also known as the "Nice" rule);

Opportunistic rule: O = "m, m, m, m, m, m, m, m, m". $x_i(t) = \text{argmax}_{x_i(t)} E\pi_i[x_i(t) \mid x_{-i}(t-1)]$ (no matter what the others bid, always bid the best-response strategy that maximizes own profit, assuming the other players stick with their last bids; in the first period, assume other players will bid cooperatively, i.e., $x_{-i}(0) = u$; sometimes this strategy is also called the "Nasty" rule);

Adaptive-learning rule: L = "m, m, m, m, c, m, m, m, m". $x_i(1) = u_i = \mathrm{u}$; if $x_j(t-1) = u$ and $x_k(t-1) = u$, then $x_i(t) = u$ else $x_i(t) = \underset{x_i(t)}{\text{argmax}}\, E\pi_i[x_i(t) \mid x_{-i}(t-1)]$ (bid cooperatively initially; if both the other players bid cooperatively at the last time period, then cooperate; otherwise, bid best-response assuming the other players stick to their last period's bids).

We now report experimental results using non-myopic bidding.

5.2. Experimental results

5.2.1. One agent learning

We now introduce one learning agent, Seller 1, into the game, while the other two agents are using fixed rules from the set of {R,C,O}. Thus, the learning artificial agent could use the adaptive-learning rule (L) as well as the other three strategies "R", "C", and "O". Again these rules are functions of other players' previous bids (one-period memory or history-dependent strategies). The winning rule the GA agent learns is "L" and it is a more sophisticated version of a Tit-for-Tat strategy. Recall the "L" rule says that "if the others cooperated, then cooperate; otherwise, defect and choose a best-response given what they chose in the previous period". This is very easy for the agent to use and it clearly has some characteristics of Tit-for-Tat. However, it is not exactly Tit-for-Tat, as Tit-for-Tat simply mimics other players' strategies, which is not easy for the agent to use in the multi-agent setting since there are multiple opponents here. Whose bids should the agent mimic? For this reason, we labeled "L" as the "adaptive-learning rule".

The experiment shows that the following strategy vector, (L,O,O), exhibits the Nash property, i.e., given the rules used by other players, no player has any incentive to switch to other rules. The result shows that only one agent learning does not lead to cooperation. As in all test cases, the result remains the same as when all players are using myopic bidding (as summarized in Table 2), although the learning agent was able to discover the best-response rule ("L") corresponding to other players' fixed rules ("O").

5.2.2. All agents learning

We now want to study what happens if all three agents (not just one of them) are intelligent, i.e., all three agents can use history-dependent rules such as "R", "C", "O", and "L". The winning rule used by each agent is "L", and the vector (L,L,L) exhibits the Nash property. Co-evolving artificial agents are heading to a Pareto efficient cooperative outcome, (u,u,u), for *all* test cases, as summarized in Table 3. This should not be surprising. If all agents bid u initially, then the above rules guarantee that all agents will continue to bid u in the future! What is interesting is that even if some of the agents defect at times, the dynamics show that learning agents learn to cooperate over time. The results in Table 3 are interesting since the folk theorem [10] in game theory predicts that any outcome in the feasible domain can be reached if the one-shot game is played in an infinite time horizon (there are quite a many of them, and the cooperative

Table 3
With all three agents learning, nonmyopic biddings lead to cooperation, $(x_1,x_2,x_3)=(u,u,u)$

	$D=(77-p)^+$	$D=(115-p)^+$	$D=(100-p)^+$
c_i=(13,13,13), K_i=(30,30,30)	(45,45,45)	(64,64,64)	(56,56,56)
c_i=(11,21,21), K_i=(44,23,23)	(46,46,46)	(65,65,65)	(58,58,58)
c_i=(16,16,25), K_i=(34,34,22)	(47,47,47)	(66,66,66)	(59,59,59)
c_i=(7,12,17), K_i=(45,26,19)	(43,43,43)	(62,62,62)	(55,55,55)
c_i=(10,10,18), K_i=(40,40,30)	(44,44,44)	(63,63,63)	(56,56,56)
c_i=(10,12,14), K_i=(40,30,20)	(44,44,44)	(63,63,63)	(55,55,55)

outcome discovered by artificial agents is one of them). Our results show that learning artificial agents can achieve the Pareto efficient cooperative outcome in a finite time horizon (in the experiment, 100 periods), suggesting a promising alternative to explain the behavior of real agents (human beings) when playing this and other repeated games.

5.3. The emergence of trust

We now further investigate the behavior of a learning artificial agent in the framework of social trust in order to explore conditions that induce cooperation. To this end, we study the impact of the "climate" on this agent's behavior. Here, we define climate as the percentage of adaptive-learning strategies (the "L" rule) used by the other players (excluding the agent itself). We test the behavior of the intelligent agent, say, Seller 1, by allying Seller 2 and Seller 3 in the sense that they

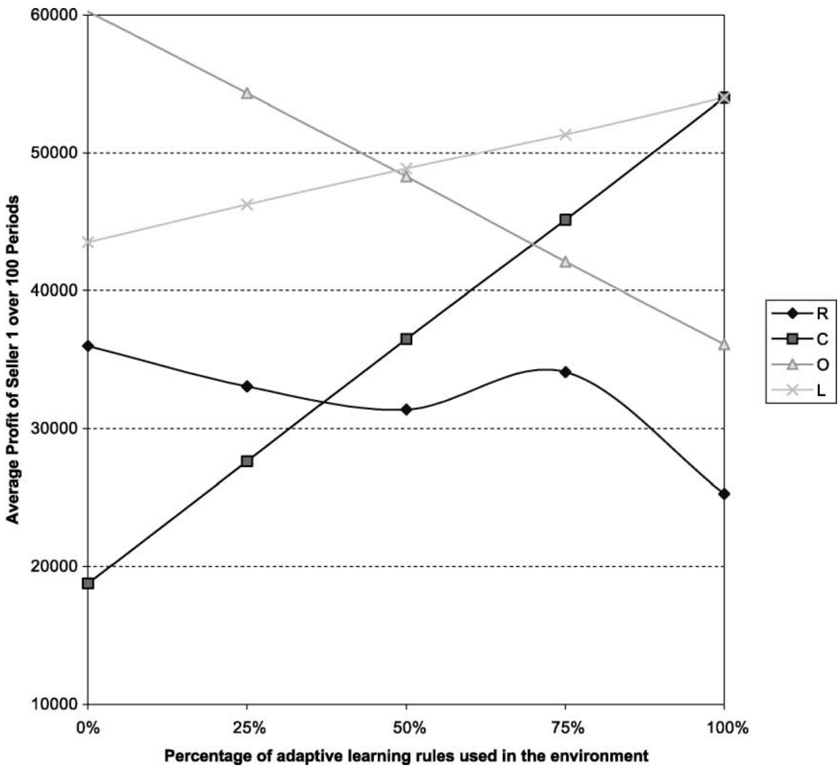

Fig. 3. The emergence of trust: artificial agent Seller 1 learns best rules to use for various climates using Example 1 parameters where $c=(10,10,18)$, $K=(40,40,30)$. When the trustworthiness in the community is low (i.e., when less than 50% of the strategies used by other sellers are adaptive learning rule L), Seller 1 learns to exploit this by using the opportunistic rule; when the trustworthiness is high (i.e., when over 50% of the strategies used by others are adaptive learning rule L), Seller 1 learns to cooperate by switching to the adaptive learning rule. As the degree of trustworthiness increases, it is much better off for an intelligent seller to trust the other players by cooperating (either use the adaptive learning rule L or the cooperative rule C). Similar results have been achieved when using other technological parameters, e.g., Example 2 parameters where c =(10,12,14), $K=(40,30,20)$, or when other sellers are learning artificial agents.

will use identical strategies, fixed over time and drawn from the set {R,C,O,L}, with the same percentage of rule "L", ranging from 0%, 25%, 50%, 75%, to 100%, with the remaining strategies equally split among {R,C,O}. Seller 1's behavior and profits over 100 time periods in different environments is shown in Fig. 3., using parameter settings as in Example 1. Intuitively, the behavior of the artificial intelligent agent suggests the following: When the trustworthiness in the community is low (i.e., when less than 50% of the rules used by other sellers are "L"), use the opportunistic rule "O" to exploit this; when the trustworthiness is high (i.e., when over 50% of the rules used by others are "L"), use rule "L" to cooperate. Shown also in Fig. 3 is that as the degree of trustworthiness increases, it is much better for an intelligent player to trust the other players by cooperating (either use rule "L" or "C"), as the player's profit keeps increasing if using these two cooperative rules suggesting that "trust pays".

Similar results have been achieved when using other technological parameters listed in Table 1 (e.g., Example 2), or when other Sellers are learning artificial agents, e.g., switch the role of Seller 3 with Seller 1, and let Seller 1 and Seller 2 ally with each other, suggesting these results are fairly general. The behavior of the intelligent agent exhibits some degree of "identity", rather than merely having a set of "do or die" strategies. This is interesting and sheds lights on designing "identity-centric" versus "strategy-centric" artificial agents. The former, "identity-centric" artificial agents are of significant importance in strategic contexts [21].

6. Conclusions and future research

We now briefly summarize the findings of the various experiments conducted. First, and most important, we find that artificial agents are viable in an automated marketplace: in the repeated bidding game, they can discover the noncooperative equilibrium if it exists; in cases where there are multiple equilibria (in the one-shot game), the agents discover the Pareto efficient equilibrium; further, they discover the cooperative outcome that Pareto dominated the noncooperative equilibrium. Second, full cooperation is achieved in repeated bidding even when the one-shot equilibrium does not exist. This demonstrates that learning adaptive artificial agents are capable of finding cooperative strategies in a complex dynamic environment. Third, agent learning plays a significant role in inducing agent cooperation. Under myopic bidding, no cooperation results. However, under nonmyopic bidding, the resulting outcomes entail full cooperation and exhibit the Nash property. Finally, we find some preliminary conditions for the emergence of trust: (a) nonmyopic bidding can lead to cooperation; (b) strategies like Tit-for-Tat can induce cooperation; (c) climate has an impact on learning agents' behavior. For example, a "friendly" climate (when the majority are playing nicely using rule "C") does not ensure cooperation, as noted above. In fact, some agents tend to be "opportunistic" when the majority of the group is "nice". This, in turn, disrupts the cooperative behavior or trustworthiness of the whole community.

While these results are consistent with findings in the iterated prisoner's dilemma literature where the equilibrium exists, we believe they open the door to defining computational principles of trust in strategic situations where the equilibrium might or might not exist. Along the same line, using a different learning regime (Q-learning), Kimbrough, Wu and Zhong [22, 39,48] explore coordination and cooperation among "identity-centric" artificial agents in other types of trust games, ultimatum games, and supply chain games. In the long term, we hope a general theory of trust will emerge for both equilibrium and non-equilibrium settings that can be computationally replicated by artificial agents. Such a theory would then have some credible justification for broader use in the application of artificial agents to support efficient e-Business activities.

Acknowledgements

The authors are grateful to Paul R. Kleindorfer for pointing out the research direction explored in this paper as well as for helping us in framing and presenting the results reported here. We thank Tung Bui, Steven O. Kimbrough, Melvin Shakun, and Yao-Hua Tan for many stimulating conversations, and the participants of the FMEC 2000 workshop, HICSS-34 and workshops at the Naval Postgraduate School and the Smeal College of Business of Pennsylvania State University for their comments. We express our special thanks to the anonymous referee of this journal for the critical yet extremely insightful comments that led to the current paper. Thanks to Marjorie Weiler for proofreading this work. Partial results of this paper appeared in Proceedings of HICSS-34 [42,43]. This work was supported in part by a mini-Summer research grant and a summer research fellowship from the Safeguard Scientifics Center for Electronic Commerce Management, Bennett S. LeBow College of Business, Drexel University and an equipment grant from Hewlett-Packard.

References

[1] H. Akkermans, F. Ygge, R. Gustavsson, HOMEBOTS: intelligent decentralized services for energy management, Proceedings of the Fourth International Symposium on the Management of Industrial and Corporate Knowledge (ISMICK), 1996.

[2] R. Axelrod, The Evolution of Cooperation, Basic Books, New York, NY, 1984.

[3] R. Axelrod, W. Hamilton, The evolution of cooperation, Science 211 (27) (March 1981) 1390–1396.

[4] F. Barazier, F. Cornelissen, R. Gustavsson, C.M. Jonker, O.L. Lindeberg, B. Polak, J. Treur, Agent Negotiating for Load Balancing of Electricity Use, Working Paper, Department of Mathematics and Computer Science, Artificial Intelligence Group, Vrije University of Amsterdam. http://www.cs.vu.nl.

[5] J. Berg, J. Dickhaut, K. McCabe, Trust, reciprocity, and social history, Games and Economic Behavior 10 (1994) 122–142.

[6] P. Dasgupta, Trust as a commodity, in: D. Gambetta (Ed.), Trust: Making and Breaking Cooperative Relations, Blackwell, Oxford, 1988.

[7] J. Engle-Warnick, Learning to Trust, Working Paper, Department of Economics, University of Pittsburgh, 2000.

[8] E. Erev, A. Roth, Predicting how people play game: reinforcement learning in experimental games with unique, mixed strategy equilibria, The American Economic Review 88 (1998) 848–881.

[9] E. Fehr, S. Gachter, G. Kirchsteiger, Reciprocity as a contract enforcement device: experimental evidence, Econometrica 65 (1997) 833–860.

[10] D. Fudenberg, J. Tirole, Game Theory, The MIT Press, Cambridge, 1991.

[11] A.R. Greenwald, J.O. Kephart, Shopbots and pricebots, Proceedings of IJCAI'99, 1999.

[12] A.R. Greenwald, J.O. Kephart, Probabilistic pricebots, AGENTS'01, Quebec, Canada, May 2001.

[13] W. Güth, P. Ockenfels, M. Wendel, Cooperation based on trust: an experimental investigation, Journal of Economic Psychology 18 (1997) 15–43.

[14] J. Holland, Artificial adaptive agents in economic theory, The American Economic Review 81 (1975) 365–370.

[15] J. Holland, Adaptation in Natural and Artificial Systems, MIT Press, Cambridge, MA, 1992.

[16] J. Hu, M.P. Wellman, Online learning about other agents in a dynamic multiagent system, Autonomous Agents 98, Minneapolis, MN, ACM Press, New York, NY, 1998.

[17] N.R. Jennings, K. Sycara, M. Wooldridge, A roadmap of agent research and development, Autonomous Agents and Multi-Agent Systems 1 (1998) 275–306.

[18] A.J. Jones, On the concept of trust, Decision Support Systems (this issue).

[19] J.H. Kagel, A.E. Roth (Eds.), The Handbook of Experimental Economics, Princeton University Press, Princeton, NJ, 1995.

[20] M. Kandori, G.J. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, Econometrica 61 (1) (1993) 29–56.

[21] S. Kimbrough, D.J. Wu, Designing Artificial Agents for E-Business, mimeo, Department of Operations and Information Management, The Wharton School, University of Pennsylvania, March 2001.

[22] S. Kimbrough, D.J. Wu, F. Zhong, Computers play the beer game: can artificial agents manage supply chains? Decision Support Systems (this issue).

[23] S. Marsh, Formalizing Trust as a Computational Concept, PhD Thesis, Department of Computing Science and Mathematics, University of Stirling, 1994.

[24] J. Miller, The coevolution of automata in the repeated prisoner's dilemma, Journal of Economic Behavior and Organization 29 (1996) 87–112.

[25] M. Olson, S.J. Rassenti, V. Smith, Market design and motivated human trading behavior in electricity markets, IIE Transactions on Operations Engineering, in press.

[26] D.C. Parkes, *i*Bundle: an efficient ascending price bundle auction, Proc. ACM Conference on Electronic Commerce (EC-99), ACM Press, New York, NY, 1999.

[27] D.C. Parkes, L.H. Ungar, Iterative combinatorial auctions: theory and practice, Proc. 17th National Conference on Artificial Intelligence (AAAI-00), AAAI Press, Mewlo Park, California, 2000.

[28] S.J. Rassenti, V.L. Smith, B.J. Wilson, Using experiments to inform the privatization/deregulation movement in electricity, Forthcoming in The Cato Journal 21 (3) (Winter, 2002).

[29] T. Sandholm, V. Ferrandon, Safe exchange planner, Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS), 2000.

[30] T. Sandholm, V. Lesser, Equilibrium analysis of the possibilities of unenforced exchanges in multiagent systems, IJCAI, (1995) 694–701.

[31] M. Schillo, P. Funk, Who can you trust: dealing with deception, Proceedings of the Workshop Deception, Fraud and Trust in Agent Societies at the Autonomous Agents Conference, ACM Press, New York, NY, 1999.

[32] S. Shapiro, The social control of impersonal trust, American Journal of Sociology 93 (3) (1987) 623–658.

[33] S. Tadelis, What is a name? Reputation as a tradeable asset, working paper, department of economics, Stanford University, 1997.

[34] Y.H. Tan, W. Thoen, Formal aspects of a generic model of trust for electronic commerce, Decision Support Systems (this issue).

[35] N. Vulkan, N.R. Jennings, Efficient mechanisms for the supply of services in multi-agent environments, Decision Support Systems 28 (2000) 5–19.

[36] M.P. Wellman, A market-oriented programming environment and its application to distributed multicommodity flow problems, Journal of Artificial Intelligence Research 1 (1993) 1–22.

[37] D.J. Wu, Discovering near-optimal pricing strategies for the deregulated electronic power marketplace using genetic algorithms, Decision Support Systems 27 (1/2) (1999) 25–45.

[38] D.J. Wu, Artificial agents for discovering business strategies for network industries, International Journal of Electronic Commerce 5 (1) (Fall 2000) 9–36.

[39] D.J. Wu, S. Kimbrough, F. Zhong, Artificial agents play the Mad Mex Trust Game: a computational approach, HICSS-35, 2002.

[40] D.J. Wu, P. Kleindorfer, J.E. Zhang, Optimal bidding and contracting strategies for capital-intensive goods, European Journal of Operational Research 137 (3) (March 2002), 657–676.

[41] D.J. Wu, P. Kleindorfer, J.E. Zhang, Integrating Contracting and Spot Procurement with Capacity Options, Working Paper, Department of Operations and Information Management, The Wharton School, University of Pennsylvania, 2001.

[42] D.J. Wu, Y. Sun, Multi-agent bidding and contracting for non-storable goods, HICSS-34, 2001.

[43] D.J. Wu, Y. Sun, The emergence of trust in multi-agent bidding: a computational approach, HICSS-34, 2001.

[44] P.R. Wurman, M.P. Wellman, W.E. Walsh, The Michigan Internet AuctionBot: a configurable auction server for human and software agents, Autonomous Agents'98, Minneapolis, MN, 1998.

[45] Y. Yan, J. Yen, T. Bui, A multi-agent based negotiation support system for distributed transmission cost allocation, HICSS-33, 2000.

[46] J. Yen, Y. Yan, J. Contreras, P.C. Ma, F. Wu, Multi-agent approach to the planning of power transmission expansion, Decision Support Systems 28 (2000) 279–290.

[47] F. Ygge, H. Akkermans, Power load management as a computational market, Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS'96), Kyoto, IEEE Press, Piscataway, NJ, 1996.

[48] F. Zhong, S. Kimbrough, D.J. Wu, Cooperative agent systems: artificial agents play the ultimatum game, HICSS-35, 2002.

D.J. Wu is Associate Professor of MIS at LeBow College of Business, Drexel University, USA. Dr. Wu's research interest include: Models of Capacity Bidding, Contracting with Options and Applications in B2B Exchanges Design; Multi-Agent Gaming and Social Trust in E-Business; Performance Impact and Outsourcing of ERP Systems; Advertising in Electronic Markets; and Formal Models of Electronic Commerce. Dr. Wu's work has appeared in *Decision Support Systems*, *European Journal of Operational Research*, *International Journal of Electronic Commerce*, *International Journal of Electronic Markets*, among others. Dr. Wu served (or is serving) as a co-guest editor for *International Journal of Electronic Commerce* as well as (*INFORMS*) *Journal of Group Decision and Negotiation*. Dr. Wu is currently co-editing a book on *Formal Models of Electronic Commerce*, to be published by Springer-Verlag in its series of electronic commerce. Dr. Wu obtained his PhD (in Operations and Information Management) from the Wharton School, University of Pennsylvania; his BE in Computer Science and BE in Management Engineering from Tsinghua University in Beijing, China. Dr. Wu has been a three-time winner of the SAP University Alliance Award (1998, 1999 and 2001). Dr. Wu's most recent work can be found at http://www.lebow.drexel.edu/wu300/research.htm.

Yanjun Sun is a PhD student of Management Information Systems at the LeBow College of Business at Drexel University, Philadelphia, PA, USA. His research interests focus on multi-agent bidding, auction and contracting support systems, and models of IT investment decision making. Yanjun's work has appeared in *International Journal of Electronic Markets*, and *Proceedings of the Annual Hawaii International Conference on System Sciences* (*HICSS-34*, *HICSS-35*). Yanjun obtained his MM in Marketing, BE in Finance, AS in Mathematics, and AJ in Law from Tsinghua University in Beijing, China.