The Emergence of Trust in Multi-Agent Bidding: A Computational Approach

D.J. Wu, Yanjun Sun
LeBow College of Business, Drexel University, Philadelphia, PA 19104, USA
{Wudj, ys45}@drexel.edu

Abstract

We experiment conditions for the emergence of trust in a multi-agent bidding setting using a computational approach. This is a follow up of our study on multi-agent bidding and contracting for non-storable goods. In this framework, there is a long-term contract market as well as a back-stop spot market. Seller agents bid into an electronic bulletin board their contract offers in terms of price and capacity, while Buyer agents decide how much to contract with the Sellers and how much to shop from the spot market. The goal is the following: First, to investigate if agents using reinforcement learning (non-myopic bidding) can discover good and effective bidding, auction and contracting strategies when playing repeated non-linear games where there does not exist any equilibrium; Second, to explore the emergence of trust in the sense of what kind of mechanisms induce cooperation in the above repeated non-linear game.

1. Introduction

We study the emergence of social trust [3, 4, 7, 14, 15, 18, 19] in a setting that was motivated from real problems of Internet auction of non-storable goods such as electric power [25, 26, 27, 32, these references are grouped thereafter as WKZ for short]. WKZ papers set up the theoretical framework for the optimal bidding and contracting for non-storable goods. The framework models the interaction of long-term contracting and spot market transactions between Sellers and Buyers for non-storable 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. WKZ papers show the optimal contracting strategies for the Buyers and the optimal bidding and auction strategies for the Sellers. WKZ papers also characterize existence conditions and structure of market equilibrium for the associated competitive game between Sellers, under the assumption that Sellers know Buyers' demand functions. However, if these conditions are violated, there does not seem to exist any equilibrium. However, in practical Internet auctions, the conditions as required in the WKZ theoretical framework might frequently be violated, which gives us the motivation of this study. Here we are interested in using artificial agents based computational approach [10, 13, 16, 17, 23, 24, 27, 30, 31] to explore the dis-equilibrium behavior of such a dynamic bidding system. In doing so, we model each selling artificial agent as a learning automata [12, 17], where the learning mechanism is characterized in reinforcement learning [6, 21] and genetic algorithms [11]. Each buying artificial agent is embedded with an optimal contracting strategy as derived in WKZ papers, hence Buyer agents do not learn since they are already using the optimal strategies given available bids on the electronic market. We study cooperation based trust [8] in this multi-agent electronic marketplace, and experiment conditions for selling agents to learn to trust [5] via co-evolution of bidding strategies [9, 17]. We test reciprocity as a device for contract enforcement or trust building [7]. This paper is a follow up of our paper [28] on multi-agent bidding and contracting for non-storable goods, where we extend the myopic learning to non-myopic learning, in particular, with a focus on the emergence of trust. We note here in passing that although throughout the paper, agents auction non-storable goods, the approach studied here is quite general, it can be used to auction storable goods as well.

The goal of the study is the following: First, we investigate if artificial agents can discover good and effective bidding, auction and contracting strategies when playing repeated non-linear games where there does not exist any equilibrium; Second, we explore the emergence of trust in the sense of what kind of mechanisms induce cooperation in the above repeated non-linear game.

The following of the paper is organized as the following. Section 2 provides a brief literature review. Section 3 investigates non-myopic bidding strategies. Section 4 studies the emergence of trust using a computational approach. Section 5 summarizes our findings.

2. Literature Review

0-7695-0981-9/01 $10.00 (c) 2001 IEEE

The issue of social trust is not new, it has been the central debate of philosophers from the beginning of human civilization [20, 24]. This issue becomes extremely important and imperative in electronic communities since it is difficult to figure out to whom you can trust in electronic marketplace [18, 19]. Surprisingly, the notion of social trust has never been agreed upon among philosophers nor among economists [3, 4], some argue that trust might be the results of Bayesian learning [15]. This makes computational approach of trust [16] difficult if not possible. In this paper, we view trust as cooperation via reciprocity from the computational economics literature (see, e.g., [3, 7, 8, 9]), where most of the work are based on lab experiments of human beings playing various simple trust games [5, 6, 7, 8, 9, 22]. Miller [17] studies the co-evolution of artificial agents playing the prisoner's dilemma game [1, 2] where agents are automata equipped with genetic algorithms. Wu and his colleagues explore coordination and cooperation in multi-agent organizational systems such as multi-agent enterprise modeling and multi-agent supply chain management [13, 24, 29]. In this paper, we continue such exploration for a different problem: multi-agent bidding and trust building in an electronic marketplace.

In the next section, we describe our model in detail.

3. Non-Myopic Bidding

We use the same price-bidding model for the multi-Seller case as in [28]. There are N Sellers, which form a set $\Xi = \{1,\cdots,N\}$. Each Seller i maximizes its expected profit Ep_i by bidding a contract price x_i anticipating the Buyer's optimal contracting strategy Q_i. Each Seller has a capacity limit K_i and a minimum cost c_i for entering the forward contract market. In the following illustrative example, we assume as in standard economics literature, linear contract demand, however, our model is general in handling any demand functions, linear or non-linear. We assume the demand function as $D_i = (100 - x_i)^+$, where $y^+ = Max[y,0]$. In case there is a bid tie, following WKZ [27], we adopt the following bid-tie allocation mechanism: If there is a tie in bids among any subset of Sellers, then Buyers' total demand for that subset of Sellers is allocated to the Sellers in proportion to their respective bidding capacities.

The following model describes the Seller's problem. $\forall k \in \Xi$, we define the following sets:

$$M_k^1 = \{i \in \Xi | x_i < x_k\};$$
$$M_k^2 = \{i \in \Xi | x_i \le x_k\};$$
$$M_k^3 = \{i \in \Xi | x_i = x_k\};$$
$$M_k^4 = \{i \in \Xi | x_i > x_k\};$$

Seller's Model:

$$\underset{x_k}{Max}\ Ep_k$$
$$p_k = [x_k - c_k] \cdot Q_k$$
$$D_k = (100 - x_k)^+$$

where according to WKZ [27],

$$Q_k(x) = \begin{cases} K_k & D(x_k) > \sum_{i \in M_k^2} K_i \\ \dfrac{K_k}{\sum_{i \in M_k^3} K_i}\left(D(x_k) - \sum_{i \in M_k^1} K_i\right) & \sum_{i \in M_k^1} K_i < D(x_k) \le \sum_{i \in M_k^2} K_i \\ 0 & D(x_k) \le \sum_{i \in M_k^1} K_i \end{cases}$$

We use the same three-Seller repeated game with the same technology and capacity parameter settings (see Table 1) as in [28] to illustrate our approach. In the table, c_i is the technology index that indicates the minimum price requirement for Seller i to enter the forward contract market, and K_i is the Seller i's total available capacity.

Table 1: Technology and capacity parameters for the three-Seller contract market.

	i	1	2	3
Ex.1	c_i	10	10	18
	K_i	40	40	30

Since myopic pure strategy price bidding does not lead to cooperation [28], we study whether non-myopic pure strategy price bidding would do better. If so, we are interested to know the conditions for agents' cooperation. We begin with no agent learning (fixed strategy tournament), then move to one or more agents learning. It is natural for us to choose reinforcement learning [21] as the primer learning mechanisms for our artificial agents since we do not know what would be the optimal or good bidding strategies in this setting (as noted in the introduction section), rather, the artificial agents have to discover such good or reasonable bidding strategies, i.e., there does not exist a teacher who can provide the feedback to train the agents as what are the right directions.

No learning. First we consider agents with fixed non-myopic strategies (no learning over time). In this case, the bidding strategies of agents are pre-specified and fixed for each and every time period (*t*). We endow each Seller with strategies like Random, or Tit-for-Tat or other strategies (defined below), and conduct tournaments to find out the winning strategies, in the spirit of Axelrod's tournament of the Prisoner's Dilemma game [1, 2]. However, it should be clear that we are investigating a different class of problems that are of both theoretical and practical importance as discussed in the introduction as well as in WKZ papers.

Consider three Sellers (*i*, *j*, *k*), each Seller tries to maximize his own profit,

$$\underset{x_i(t)}{Maximize}\ Ep_i\ ,$$

while competing with the other two Sellers by bidding a price *x(t)* in an electric marketplace. The game is repeated for a fixed number of times (*t = 30*).

Each Seller is equipped with three possible strategies: Random, Nice, and Tit-for-Tat, which are defined as the following:

Seller 1:

Random: $x_i(t) \in [18{,}55]$;

Nice: $x_i(t) = 55$;

Tit-for-Tat: $x_i(1) = 55$;

If $x_j(t-1) = 55$ and $x_k(t-1) = 55$, then $x_i(t) = 55$;

If $x_j(t-1) \neq 55$ or $x_k(t-1) \neq 55$, then $x_i(t) = 18$;

Seller 2:

Random: $x_j(t) \in [18{,}55]$;

Nice: $x_j(t) = 55$;

Tit-for-Tat: $x_j(1) = 55$;

If $x_i(t-1) = 55$ and $x_k(t-1) = 55$, then $x_j(t) = 55$;

If $x_i(t-1) \neq 55$ or $x_k(t-1) \neq 55$, then $x_j(t) = 18$;

Seller 3:

Random: $x_k(t) \in [19{,}55]$;

Nice: $x_k(t) = 55$;

Tit-for-Tat: $x_k(1) = 55$;

If $x_i(t-1) = 55$ and $x_j(t-1) = 55$, then $x_k(t) = 55$;

If $x_i(t-1) \neq 55$ or $x_j(t-1) \neq 55$, then $x_k(t) = 19$.

The results are shown in Table 2.

Table 2: Computer tournament of non-myopic strategies. Each Seller has three possible strategies, N (Nice Strategy), T (Tit-for-Tat strategy) and R (Random strategy) over the 30 time periods. The tournament consists of 3*3*3 = 27 games. Winner strategies are Tit-for-Tat and Nice, however, Tit-for-Tat is Nash while Nice is not.

	Strategy	Profit
1	(R, R, R)	(17089, 14500, 5982)
26	(T, T, N)	(22091, 22091, 13623)
27	(T, T, T)	(22091, 22091, 13623)

Now we add a "Nasty" strategy to each Seller's strategy space and test the robustness of previous winner strategies. For every Seller, we define the Nasty strategy as the following:

Seller 1:

$x_i(1) = 55$;

If $x_j(t-1) = 55$ and $x_k(t-1) = 55$, then $x_i(t) = 54$;

If $x_j(t-1) = 55$ and $x_k(t-1) \neq 55$, then $x_i(t) = 40$;

If $x_j(t-1) \neq 55$ and $x_k(t-1) = 55$, then $x_i(t) = 35$;

If $x_j(t-1) \neq 55$ and $x_k(t-1) \neq 55$, then $x_i(t) = 18$;

Seller 2:

$x_j(1) = 55$;

If $x_i(t-1) = 55$ and $x_k(t-1) = 55$, then $x_j(t) = 54$;

If $x_i(t-1) = 55$ and $x_k(t-1) \neq 55$, then $x_j(t) = 40$;

If $x_i(t-1) \neq 55$ and $x_k(t-1) = 55$, then $x_j(t) = 35$;

If $x_i(t-1) \neq 55$ and $x_k(t-1) \neq 55$, then $x_j(t) = 18$;

Seller 3:

$x_k(1) = 55$;

If $x_i(t-1) = 55$ and $x_j(t-1) = 55$, then $x_k(t) = 54$;

If $x_i(t-1)=55$ and $x_j(t-1)\neq 55$, then $x_k(t)=39$;

If $x_i(t-1)\neq 55$ and $x_j(t-1)=55$, then $x_k(t)=39$;

If $x_i(t-1)\neq 55$ and $x_j(t-1)\neq 55$, then $x_k(t)=19$.

Our experiment shows that Tit-for-Tat is fairly robust, it remains to be the winner strategy, as shown in Table 3.

Table 3: Computer tournament of non-myopic strategies. Each Seller has four possible strategies, A (Nasty strategy), N (Nice strategy), T (Tit-for-Tat strategy) and R (Random strategy) over the 30 time periods. The tournament consists of 4*4*4 = 64 games. When introducing Nasty strategies, Tit-for-Tat remains to be the winner strategy.

	Strategy	Profit
1	(R, R, R)	(14718, 13427,7705)
42	(T, T, N)	(22091, 22091, 13623)
43	(T, T, T)	(22091, 22091, 13623)
64	(A, A, A)	(10656, 10656, 934)

One agent learning. Now we introduce learning agents into the game. First we let Seller 1 to learn, and to test the performance of this artificial intelligent agent in the above environment. The goal of this agent is:

$$\underset{x_i(t)}{Maximize}\ Ep_i$$

$$Ep_i=\sum_{t=1}^{30} p_{i,t}(x_i(t)\,|\,x_{-i}(t)),$$

where, $x_i(t)$ is Seller 1's bid at Time t, and $x_{-i}(t)$ is the others' bids at the same time period.

As shown in Table 4, the learning artificial agent discovers a strategy that outperforms all other strategies, and it has some characteristics of a Tit-for-Tat strategy. An interesting finding of this experiment is that Tit-for-Tat remains to be the best response for the other players when playing against an "intelligent" player.

Table 4: Computer tournament of non-myopic bidding strategies. Seller 2 and 3 each has four possible strategies, A (Nasty strategy), N (Nice strategy), T (Tit-for-Tat strategy) and R (Random strategy) over the 30 time periods, whereas Seller 1 is an intelligent agent (using a dynamic G strategy). The tournament consists of 5*4*4 = 80 games. The intelligent agent discovers a winner strategy that outperforms all other strategies, and it has some characteristics of a Tit-for-Tat strategy.

	Strategy	Profit
65	(G, R, R)	(26356, 9842, 4143)
75	(G, T, T)	(23115, 21483, 13248)

All agents learning. The above experiment suggests that strategies like "Tit-for-Tat" may lead to cooperation. Now we want to study what if all agents are "intelligent" in the sense that they can learn and adapt. We design a non-myopic bidding system, as depicted in Figure 1, to investigate.

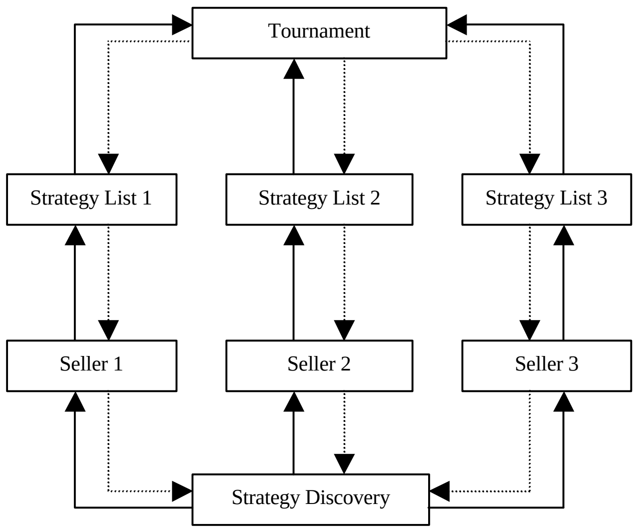

Figure 1: Non-myopic bidding system

We now formally define this game. Assume each agent has H rules in one generation, *i*, *j*, *k* indexes one of rules of Agent 1, 2, 3, respectively. Each rule specifies the agent's bid in each and every time period (*t* = *1*, …, *T*). Denote *t* as the t^{th} time period and *g* as the g^{th} generation for each Seller to discover new rules. The objective of each agent is:

Seller 1: $\underset{x_i(g,t)}{Maximize}\ Ep_i$

$$Ep_i=\frac{\sum_{j=1}^{H}\sum_{k=1}^{H}\sum_{t=1}^{T}p_{i,g,t}(x_i(g,t),x_j(g,t),x_k(g,t))}{H\times H};$$

Seller 2: $\underset{x_j(g,t)}{Maximize}\ Ep_j$

$$Ep_j=\frac{\sum_{i=1}^{H}\sum_{k=1}^{H}\sum_{t=1}^{T}p_{j,g,t}(x_i(g,t),x_j(g,t),x_k(g,t))}{H\times H};$$

Seller 3: $\underset{x_k(g,t)}{Maximize}\ Ep_k$

$$Ep_k=\frac{\sum_{i=1}^{H}\sum_{j=1}^{H}\sum_{t=1}^{T}p_{k,g,t}(x_i(g,t),x_j(g,t),x_k(g,t))}{H\times H}.$$

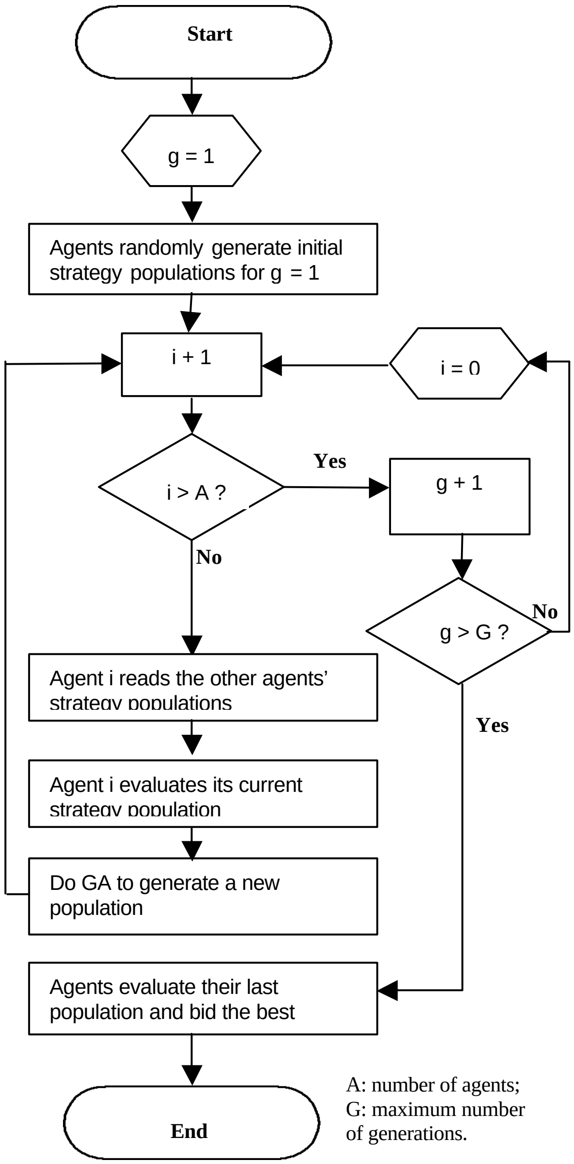

Figure 2: Flowchart of non-myopic bidding system.

Figure 2 describes the flowchart of our non-myopic price bidding system. Table 5 shows the result (*300* generations) for $T = 10$ weeks' bid. Once again, we observe that strategies similar to Tit-for-Tat lead to cooperation.

Table 5: Non-myopic bidding experiment: All agents learning.

	1	2	3	4	5	6	7	8	9	10
1	26	23	27	28	25	26	27	28	28	29
2	24	27	26	27	26	25	27	28	28	29
3	25	28	27	28	27	25	27	28	27	29

4. The Emergence of Trust

To further explore conditions that induce cooperation, we study the impact of the "climate" on agents' behavior. Here we define climate as percentage of nasty strategies used by the other players. For example, a nice climate would be that no other players (excluding the agent itself) use nasty strategies, a mild climate would that 50% of the other players are nice and the rest 50% are nasty, a bad climate would be that 100% of other players are nasty.

No learning. Similar to the previous section, we study the performance of different fixed strategies under various climates. In particular, we focus on four different strategies: Random, Nice, Tit-for-Tat, and Nasty. These strategies are defined in the previous section. The results are shown in Figures 3 and 4. An interesting finding is that the Nasty strategy seems to be the winner in various climates studied, followed by a random strategy. The reason why Tit-for-Tat strategy does not do well this time deserves further investigation, since it may do well in extremely bad climate. In any case, Niceness does not pay in the above experiment. The above findings suggest that the winning strategy depends on the agent's negotiation power as well. In this case, Seller 1 and Sell 2 enjoy a technology advantage over Sell 3, therefore, they can afford to behave nastily to extract as much profit margin as possible from the market. Seller 3 has a technology disadvantage, so he has no choice but to be nasty given the other players will surely take nasty positions. This can be further illustrated under the very bad climate where 100% of his opponent players are using nasty strategies. In this case, whatever strategy Seller 3 takes does not make much difference to his fortune.

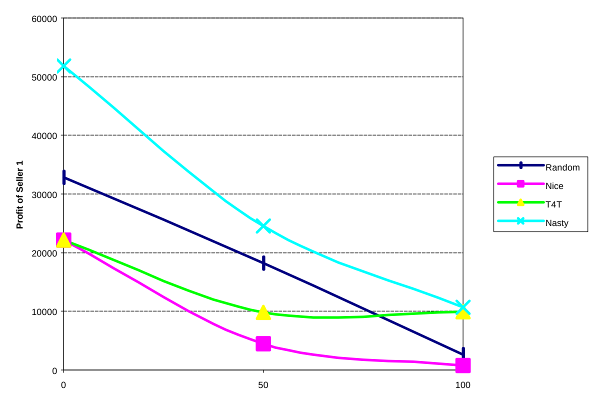

Figure 3: Seller 1's behavior in various climates.

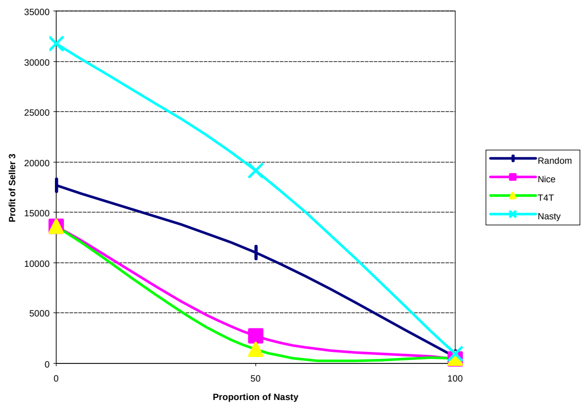

Figure 4: Seller 3's behavior in various climates.

One agent learning. Now we let Seller 1 to be an intelligent agent, and rerun the experiment. It discovers a better strategy that beats the "Nasty" strategy. It performs well under extremely violent circumstances. These results are depicted in Figure 5.

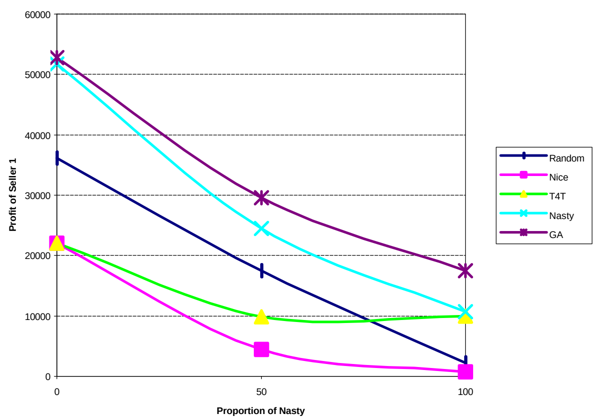

Figure 5: Intelligent agent Seller 1's performance in various climates.

We further test the behavior of this intelligent agent by allying Agent 2 and 3 in the sense that they will have identical strategy population with the same percentage of nasty strategies, ranging from 0%, 25%, 50%, 75%, to 100%. Figure 6 shows our intelligent Seller 1's profits over dynamic environments, and Table 6 shows his corresponding bids over the first 10 weeks.

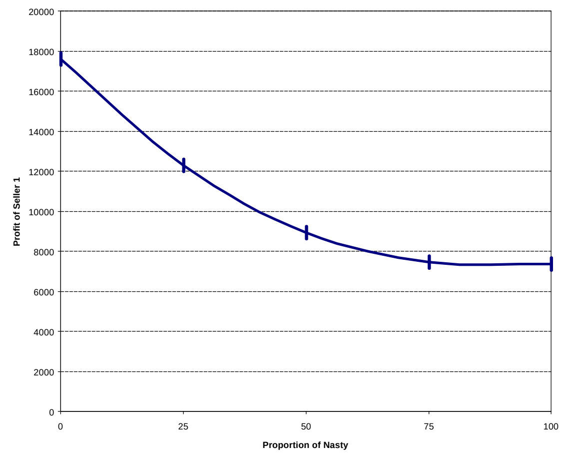

Figure 6: Seller 1's profit in dynamic climates.

Table 6: GA-agent behavior in various dynamic environments. (Row stands for 10 consecutive periods' bids in the same environment; Column stands for one week's bid in various environments.)

	1	2	3	4	5	6	7	8	9	10
0%	54	54	54	54	54	54	54	54	54	54
25%	54	54	54	54	54	54	54	54	54	54
50%	54	34	48	55	53	55	53	34	55	53
75%	54	34	55	39	34	55	39	34	55	39
100%	54	34	55	39	55	39	55	39	55	39

We observe that, from Figure 6 and Table 6, the intelligent agent learns immediately how to take advantage of a nice climate by being nasty (Row 1 and 2 in Table 6) and profit as much as it can. Even more interesting, the intelligent agent can learn a smart "trick" by "pretending" to be nice (by bidding 55 in Week 4) to induce the opponents' use of a greedy nasty strategy (bidding 54) in the next time period, Week 5. However, in Week 5, our intelligent agent, Seller 1, bids 53 to take full advantage of his opponent's greediness.

5. Summary

We now briefly summarize findings of various experiments conducted here. First, and most importantly, we find that artificial agents are viable in automated marketplace: they can find better strategies in a complex dynamic environment where such an equilibrium does not exist. Second, we find some preliminary conditions for the emergence of trust: non-myopic bidding can lead to cooperation. Strategies like Tit-for-Tat can induce Sellers to cooperate. Third, we find that the climate has an impact on agents' behavior. Perhaps one of the most interesting findings here is that "friendly" climate does not ensure cooperation. In fact, some agents tend to be "nasty" when the majority of the group is "nice". This in turn, disrupts the cooperation behavior of the whole community. Goodwill does not necessarily lead to cooperation or trust. We believe these results open the door of a novel approach on defining computational principles of trust, and in the long term, we hope a general theory on dis-equilibrium and trust would emerge.

Acknowledgement

Thanks to Tung Bui, Steven O. Kimbrough, Paul R. Kleindorfer and Yao-Hua Tan for many stimulating conversations, and to two anonymous referees of HICSS-34 and the participants of the FMEC 2000 workshop for their comments. File: CLNSS04.doc. 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. Corresponding author is D.J. Wu, his current address is: 101 North 33rd Street, Academic Building, Philadelphia, PA 19104. Email: wudj@drexel.edu.

References

1. Axelrod, R., *The Evolution of Cooperation*, Basic Books, New York, N.Y., 1984.
2. Axelrod, R., and Hamilton, W., "The evolution of cooperation," *Science*, 211, 27, 1390-1396, March 1981.
3. Berg, J., Dickhaut, J., and McCabe, K., "Trust, Reciprocity, and Social History", *Games and Economic Behavior*, 10, 122-142, 1994.
4. Dasgupta, P., "Trust as a Commodity", in: D. Gambetta (Ed.), *Trust: Making and Breaking Cooperative Relations*, 49-72, Blackwell, Oxford and New York, 1988.
5. Engle-Warnick, J., "Learning to Trust", Working Paper, Department of Economics, University of Pittsburgh, 2000.
6. Erev, E., and Roth, A., "Predicting How People Play Game: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria", *The American Economic Review*, 88, 848-881, 1998.
7. Fehr, E., Gachter, S., and Kirchsteiger, G., "Reciprocity as a Contract Enforcement Device: Experimental Evidence", *Econometrica*, 65, 833-860, 1997.
8. GÜth, W., Kliemt, H., and Peleg, B., "Co-Evolution of Preferences and Information in Simple Games of Trust", Working Paper, Department of Economics, Humboldt University of Berlin, 2000.
9. GÜth, W., Ockenfels, P., and Wendel, M., "Cooperation Based on Trust: An Experimental Investigation", *Journal of Economic Psychology*, 18, 15-43, 1997.
10. Holland, J., "Artificial Adaptive Agents in Economic Theory", *The American Economic Review*, 81, 365-370, 1975.
11. Holland, J., *Adaptation in Natural and Artificial Systems*, MIT Press, 1992.
12. Hopcroft, J., and Ullman, J., *Introduction to Automata Theory, Languages and Computation*, Addison-Wesley, Reading, MA, 1979.
13. Kimbrough, S., Wu, D.J., and Zhong, F., "Beer Game Computers Play: Can Artificial Agents Manage the Supply Chain?", *HICSS-34*, 2001.
14. Lahno, B., "Trust, Reputation, and Exit in Exchange Relationships", *Journal of Conflict Resolution*, 39 (3), 495-510, 1995.
15. Lahno, B., "Is Trust the Result of Bayesian Learning", Working Paper, University of Duisburg, 2000.
16. Marsh, S., "Formalizing Trust as a Computational Concept", Ph.D. Thesis, Department of Computing Science and Mathematics, University of Stirling, 1994.
17. Miller, J., "The Coevolution of Automata in the Repeated Prisoner's Dilemma", *Journal of Economic Behavior and Organization*, 29, 87-112, 1996.
18. Rea, T., and Skevington, P., "Engendering Trust in Electronic Commerce," *British Telecommunications Engineering*, 17 (3): 150-157, 1998.
19. Schillo, M., and Funk, P., "Who Can You Trust: Dealing with Deception", *Proceedings of the Workshop Deception, Fraud and Trust in Agent Societies at the Autonomous Agents Conference*, 95-106, 1999.
20. Shapiro, S., "The Social Control of Impersonal Trust", The American Journal of Sociology, 93 (3), 623-658, 1987.
21. Sutton, R., and Barto, A., *Reinforcement Learning: An Introduction*, MIT Press, 1999.
22. Tadelis, S., "What is a Name? Reputation as a Tradeable Asset", Working Paper, Department of Economics, Stanford University, 1997.
23. Weiss, G., (Ed.), *Multi-agent Systems: A Modern Approach to Distributed Artificial Intelligence*, Cambridge, MA: MIT Press, 1999.
24. Wu, D.J., "Software Agents for Knowledge Management: Coordination in Multi-Agent Supply Chains and Auctions", Forthcoming in, *Expert Systems with Applications*, 20 (1), January 2001.
25. Wu, D.J., Kleindorfer, P., and Zhang, J. E., "Optimal Bidding and Contracting Strategies for Deregulated Electric Power Market: Part I", *HICSS-33*, 2000a.
26. Wu, D.J., Kleindorfer, P., and Zhang, J. E., "Optimal Bidding and Contracting Strategies for Capital-Intensive Goods", Under Revision for, *European Journal of Operational Research*, 2000b.
27. Wu, D.J., Kleindorfer, P., and Zhang, J. E., "Optimal Bidding and Contracting Strategies for Deregulated Electric Power Market: Part II", *HICSS-34*, 2001.
28. Wu, D.J., and Sun, Y., "Multi-Agent Bidding and Contracting for Non-Storable Goods", *HICSS-34*, 2001.

29. Wu, D.J., Sun, Y., and Zhong, F., "Organizational Agent Systems for Intelligent Enterprise Modeling", Forthcoming in, *International Journal Electronic Markets*, 10 (4), November 2000.
30. Yan, Y., Yen, J., and Bui, T., "A multi-agent Based Negotiation Support System for Distributed Transmission Cost Allocation", *HICSS-33*, 2000.
31. Zacharia, G., Moukas, A., Boufounos, P., and Maes, P., "Dynamic Pricing in a Reputation Brokered Agent Mediated Knowledge Marketplace", *HICSS-33*, 2000.
32. Zhang, J. E., Kleindorfer, P., and Wu, D.J., "The Price of a Real Option on Non-storable Commodities with Heterogeneous Supply Costs", Working Paper, Department of Operations and Information Management, The Wharton School, University of Pennsylvania, 2000.