Automated Bargaining Agents
(Preliminary Results)

Ying Sun & Daniel S. Weld*
Department of Computer Science and Engineering, FR-35
University of Washington
Seattle, WA 98195
`ysun, weld@cs.washington.edu`

Keywords: multi-agent systems, negotiation

January 19, 1995

Abstract

Rapid commercialization of the Internet and development of the National Information Infrastructure are likely to change the nature of retailing and commerce in profound ways. These changes pose both challenges and opportunites for AI, such as how to design intelligent assistants that help humans cope with the increasing complexities of electronic commerce. In this paper, we present the first steps towards creating *automated bargaining agents* — intelligent assistants that can reason about the relative supply and demand for goods and services and negotiate to reach a good deal. We first review the economic and game theoretic background for such an endeavor. We then present a simple economic model of electronic commerce and describe several bargaining strategies. Preliminary experimental results suggest that sophisticated, adaptive strategies perform better than simple discounting approaches over a range of economic conditions.

*Many thanks to Jim Dana and Kathy Spier for discussions about game theory and economics. We would also like to thank Adam Carlson, Denise Draper, Oren Etzioni, Terrance Goan, Keith Golden, Nicholas Kushmerick, Neal Lesh, W. Derrick Weathersby, Mike Williamson, and Erh-Chun Yeh for suggesting bargaining strategies and providing helpful comments. This research was funded in part by Office of Naval Research Grant N00014-94-1-0060 and by National Science Foundation Grant IRI-9303461. This paper has not already been accepted by and is not currently under review for a journal or another conference. Nor will it be submitted to one during IJCAI's review period.

1 Motivation

Rapid growth of the Internet, development of the National Information Infrastructure (NII), and maturation of CommerceNet are likely to change the nature of retailing and commerce in profound ways. These changes pose both challenges and opportunities for AI. Electronic commerce is an important opportunity for AI, because without intelligent agents to assist them, humans will be swamped by the number of possible vendors and confused by the complexity of the contracts proposed.

For example, although AT&T, MCI, and Sprint have 87% of the $61 billion long distance market, nearly 400 companies provided long distance service in the United States in 1994 [13]. Furthermore, each of these companies offers multiple contracts. For example, *AT&T True USA Savings* provides a 10% discount for bills between $10–24.99 per month, 20% discount for bills between $25–74.99, and a 30% discount for bills over $75; there is no monthly fee. On the other hand, *Friends & Family from MCI* provides a 20% discount on all calls to specific numbers regardless of the total bill; if those numbers are MCI customers the discount is 40%, but there is a $3 per month fee. Other plans provide different discounts based on the time of call. With close to a thousand possible plans, it is virtually impossible for a human to choose the optimal contract.

Thus, the challenge for AI is to design intelligent assistants that help humans cope with the increasing complexities of electronic commerce. Even simple agents could have a huge payback. For example, it would be straightforward to design a computer agent that monitored one's telephone calling pattern for a month and then selected the most economical long distance carrier. In fact, such an agent could even route different calls to different carriers depending on the time of day and number being called. Of course, these opportunities are not limited to telephone services. An intelligent agent could poll the Apollo or SABRE airline reservation system to monitor the fare and seating availability of different flights.

Furthermore, the opportunities for intelligent assistance at electronic commerce are not limited to agents that search databases for the lowest price. Great potential lies in *automated bargaining agents* — intelligent assistants that can reason about the relative supply and demand of a good or service and negotiate to reach a good price. Today, many commodities, whether pork bellies, stocks, or airline seats, are traded in active markets; electronic trading programs manipulate these markets. As electronic commerce becomes widespread, it is likely that the markets for retail products will become equally dynamic, and automated bargaining tools may become widespread. Just as car dealers and electronics stores use sticker prices and sales to increase profits from naive buyers, electronic vendors will attempt to get the highest price for their product, lowering the price only when forced to do so. To get the best buy, a human will either need to be a skilled and knowledgable negotiator — or have an intelligent bargaining assistant. In this paper we present the first steps towards creating such an agent. Section 2 describes the economic and game theoretic background for the endeavor. Then in Sections 3 and 4 we define a simple model of electronic commerce and describe several bargaining strategies. Section 5 presents our preliminary experimental results. We conclude with a discussion of related and future work.

2 Bargaining as a Game

At an intuitive level, bargaining is quite simple. Take the simplest case in which there is one buyer (B), one seller (S), and one product (a "widget"). If the widget is worth more to the buyer than to the seller (a condition we will write as $\mathsf{value}(B) > \mathsf{value}(S)$), then both agents can benefit by a transaction that exchanges the widget for a dollar amount between the two values. The question is "What is the sale price?" One might argue that a "fair" price would be halfway between the two valuations, but it is impossible for an agent to determine this price if it doesn't know the value ascribed by the other (possibly lying) agent. If there is disparity in the agents' information (*e.g.*, S knows $\mathsf{value}(B)$ but B knows little about $\mathsf{value}(S)$), then S has the advantage in any negotiation. As a result, a common tactic is for buyer B to make an offer which is substantially less than $\mathsf{value}(B)$ in an attempt to conceal her true valuation.

One can analyze bargaining in terms of game theory [10]. *Players* are individuals, such as B and S above, who are capable of making decisions. An *action* by a player is a choice she can make; examples include making an offer or accepting the offer of another player. A player's *information set* is the set of values of different variables that the player thinks are possible but cannot distinguish; for example, B knows her value with certainty, but might only know that $\mathsf{value}(S)$ falls in the interval $[x, y]$. *Nature* is a non-player who makes random actions at specified points in the game. Different players might or might not know Nature's probability distribution, and they might or might not observe each other's (or Nature's) moves. A player's *strategy* specifies which action to take at each point in the game given the player's information set. A *mixed* strategy maps information sets to a probability distribution over actions; for example, a buyer might secretly toss a coin, raising her offer only if it came up heads. A *continuous strategy* selects from a continuum of actions; if players can choose any real number as an offer price, then their bargaining strategy is continuous. A player's *payoff* is the utility she receives when the game has been played out. In a multi-period game, payoffs are usually *discounted* in the sense that they are valued less if made at later rounds; discounting is calculated with the *discount rate*, a number between zero and one.

Games are typically classified on a variety of dimensions. If players can make *binding* commitments to each other during the game, the game is *cooperative*. A game has *complete information* if Nature does not move first, or if its initial move is observed by every player. A game is *certain* if Nature does not move after any player moves. A game has *symmetric information* if no player has information different from other players when she moves or at the end of the game. A game is *perfect* if every possible information set is a singleton (*i.e.*, no player ever has any uncertainty); this is a stronger condition than symmetry or completeness since it bans simultaneous moves.

To understand the dynamics of a game, it is essential to consider the possible strategy combinations, since these determine the expected payoffs. For some games there exists a combination which can be said to represent the "best" strategy for each player; such a combination is called an *equilibrium*. By changing the precise definition of "best," economists have devised several variants: dominant strategy equilibrium, Nash equilibrium, subgame perfect equilibrium, *etc.*; space precludes their precise definition in this paper. Note that many games don't have an equilibrium and some games have multiple equilibria.

One of the simplest models of bargaining is a two-player game in which A and B seek to divide a pie. The players alternate moves. First, A proposes an offer $0 \leq \theta_1 \leq 1$ corresponding to her desired share. Next, B accepts or makes a counter-offer. If A's offer is accepted in round m and the discount rate is r, then A's payoff is $\theta_m(\frac{1}{1+r})^m$ and B's is $(1-\theta_m)(\frac{1}{1+r})^m$. If no offer is ever accepted, both payoffs are zero. This game is noncooperative, with perfect, symmetric, complete, and certain information. Rubinstein [14] showed that this discounted, infinite game has a unique, perfect equilibrium outcome in which A receives $\frac{r^2+r}{r^2+2r}$. Thus if $r = .1$ then A could demand 52% of the pie and a rational B could gain nothing by arguing and so would accept.

In the more realistic case of incomplete information, equilibrium strategies for bargaining can lead to a sequence of actions lasting more than one time period. Consider the two-player game in which the seller knows only that the buyer has either valuation b_{low} or b_{high}. Similarly, the buyer knows the seller's valuation is either s_{low} or s_{high}. (The prior probabilities are common knowledge.) Assuming $s_{\text{low}} < b_{\text{low}} < s_{\text{high}} < b_{\text{high}}$, it is possible the widget is worth less to the buyer than to the seller and no sale is mutually desirable. On the other hand, if either the seller is *soft* (*i.e.*, $\mathsf{value}(S) = s_{\rm low}$) or the buyer is *soft* (*i.e.*, $\mathsf{value}(B) = b_{\rm high}$), a sale may occur. Players alternate offers, and the payoff from a sale (if one occurs) is discounted as before. The game is noncooperative, imperfect, asymmetric, incomplete, and certain. Intuitively, a viable strategy is for a soft buyer to randomly choose between masquerading as a *hard* buyer (*i.e.*, offering an amount $\leq b_{\rm low}$) and admitting it is soft. Every time the seller receives an apparently hard offer, it uses Bayes rule to revise the probability that the buyer is soft. If either the buyer or seller is soft, it will eventually reveal itself and the game will turn into a complete information one and terminate. Chatterjee and Samuelson [2] show that this strategy forms an equilibrium for the game, but they are forced to conclude that a multitude of other equilibria exist.

In the presence of multiple equilibria, it is rather difficult to analytically prefer one strategy to another. But analytic methods are simply one technique for designing bargaining strategies. In the next section we explain how empirical evaluation can compensate for intractable analytics. We note that while computer simulation has been used to evaluate agent strategies in simple games like the prisoner's dilemma [6] and double auctions [7], it has not been applied to the general case of unrestricted, multi-period bargaining.

3 Economic Model of Electronic Commerce

We elaborate the economic model of Chatterjee Samuelson [2], to account for multiple buyers and multiple sellers. There is a single commodity ("widgets"), but each agent's valuation for a widget is an arbitrary real number (*i.e.*, no limitation of two values as in the Chatterjee and Samuelson model). Buyers and sellers use a message-based protocol in an attempt to make one or more profitable deals.

3.1 Protocol

We assume a simple, discrete model of time. In each round (or *cycle*), an agent receives all messages sent to it in the previous cycle, performs arbitrary computation, and can send one message to every other agent if it desires. Our simple protocol provides four basic types of messages:

```
(request-offer <neg-id> <buy-or-sell?>)
(offer <neg-id> <buy-or-sell?> <price>)
(accept <neg-id>)
(abort <neg-id>)
```

When an agent sends `request-offer` to another agent, the receiver must respond with an `offer` message or `abort`. When an agent receives an offer from another agent, it can *accept* the offer, counter with a different offer, or abort the negotiation. Abort and accept messages require no response. All offers are *binding* for precisely one cycle, then they expire. We chose the protocol because it models the type of interactions that occur in markets and seems a fair abstraction of future electronic commerce applications; evaluation of alternative protocols is a topic for future research.

3.2 Agent Utility

Each agent has type *buyer* or *seller*. These classifications do not limit the type of messages that an agent can send, but they do determine what increases the agent's utility (or *profit*). We assume that each buyer, B_i, desires a certain number of widgets (called $\mathsf{demand}(B_i)$) and gains utility when it accepts an offer to sell at cost less than $\mathsf{value}(B_i)$. Each seller, S_j, can produce an arbitrary number of widgets and gains utility whenever a sale exceeds $\mathsf{value}(S_j)$. Note that this implies that the market is demand driven (rather than supply driven); sellers must compete amongst themselves for a share of limited demand.

Each agent knows which agents are in the game, but no agent knows the valuation, demand, type, or strategy of another agent. Furthermore, agents only get to see the content of messages addressed to them. The act of sending a message (regardless of type) incurs cost, c. The value of c and the discount rate, $0 < r < 1$, are common knowledge to all agents.

We define the payoff of a selling agent S_x from a sequence of deals as:[1]

$$\begin{aligned} \mathsf{profit}(S_x) &= \sum_i (\mathsf{price}(D_{x,i}) - \mathsf{value}(S_x))(\frac{1}{1+r})^{T(D_{x,i})} - \\ & c\sum_j (\frac{1}{1+r})^{T(M_{x,j})} \end{aligned} \quad (1)$$

where $D_{x,i}$ is the ith converged deal for agent S_x, and $T(D_{x,i})$ is the time when $D_{x,i}$ converged; $M_{x,j}$ is the jth outgoing message from agent S_x, and $T(M_{x,j})$ is the time when message $M_{x,j}$ was sent.

This economic model leads to a game that is noncooperative, with certain, yet imperfect, asymmetric, and incomplete information.

[1] The payoff of a buying agent is defined in a symmetric manner.

4 Bargaining Strategies

Since our model of electronic commerce is significantly more complex than previous models of bargaining, we do not attempt to predict equilibrium behavior analytically. Instead, we simulate the performance of a set of strategies empirically. We selected the strategies to be tested by soliciting suggestions from a group of Computer Science graduate students and faculty at the University of Washington. The result was a set of four strategies that can be used by buyers as well as sellers, and an additional strategy that applies only to sellers.

The first two strategies were motivated by the idea of discount chains. Instead of fancy bargaining, they stick to a low price.

- FIXED PRICE: When requested, a selling agent S_i always offers $(1+\epsilon)\mathsf{value}(S_i)$, and accepts any offer above that price; a buying agent B_j always offers $(1-\epsilon)\mathsf{value}(B_j)$, and accepts the lowest offer below that price. Here ϵ denotes a fixed profit margin (*e.g.*, 5%) that the agent commits to.

- FIXED OFFER: An agent makes the same offer as in the FIXED PRICE strategy, but the acceptance condition is relaxed. A selling agent S_i accepts any offer above $\mathsf{value}(S_i)$, and a buying agent B_j accepts the lowest offer below $\mathsf{value}(B_j)$.

The intuition behind the next two strategies suggests that an agent will gain utility by trying to bias the sales price in their favor. A seller starts by suggesting a high price and monotonically lowers it; buyers do the opposite. Specifically, sellers suggest initial offers based on the negotiation history with the opponent, while buyers use previous deal prices as an upper bound for the current negotiation. For brevity, we only detail the selling strategy below:

- MONOTONIC CONCESSION: When requested for an initial offer, a selling agent S_i offers $(1+\delta)\mathsf{value}(S_i)$ if it has never dealt with the opponent (typically $\delta = 50\%$). It offers the price of the most recent deal, if it recently converged with the opponent. Otherwise, it must have failed to reach agreement with the buyer. In this case, S_i sets ϵ to a number in $[.01, .05]$ and offers $\mathsf{max}((1-\epsilon)\mathsf{offer}(S_i, last), \mathsf{value}(S_i))$. In each subsequent round, S_i accepts any offer above $\mathsf{value}(S_i)$, or makes a counter offer by reducing the last offer price by a fixed percent (*e.g.*, 2%) (but not below $\mathsf{value}(S_i)$).

- Tit for Tat: This is an adaptive version of the Monotonic Concession strategy. In round n an agent makes a concession of its price by as much as its opponent made in round $(n-1)$. *I.e.*, for a selling agent S_i $\mathsf{offer}(S_i, n) = (1 - \mathsf{concession}(B_j, n-1))\mathsf{offer}(S_i, n-2)$, where B_j denotes S_i's opponent, and $\mathsf{concession}(B_j, n-1) = \frac{\mathsf{offer}(B_j, n-1) - \mathsf{offer}(B_j, n-3)}{\mathsf{offer}(B_j, n-3)}$. The action of a buying agent is symmetric.

The final strategy is based on the idea that a seller might gain market share if it knew what the other sellers were charging. The utility of this information might be so great that a seller should invest (*i.e.*, by sending messages) in order to learn it. The basic idea is to start out with a profitable price. Meanwhile, try to buy from all other sellers, trick them into revealing their best selling price, and adjust the price to be *just lower* than the best price one's competitor is offering. This strategy only applies to sellers because the economic model is demand limited (since sellers have infinite supply, buying agents don't care what other buyers are willing to pay).

- Undercut Competitor: When requested from a buyer, a selling agent S_i initially offers $\mathsf{price}(S_i) = (1 + \delta)\mathsf{value}(S_i)$ where $\delta = 50\%$; this price is updated as S_i finds out more about the price its competitors are offering. S_i accepts any buyer's offer above $\mathsf{value}(S_i)$.

  Meanwhile, S_i tries to buy from all other sellers, bargaining hard to find out how low the competitors will go. S_i always offers $\mathsf{value}(S_i)$. Whenever a competitive seller S_j offers $\mathsf{price}(S_j) < \mathsf{price}(S_i)$, S_i adjusts $\mathsf{price}(S_i) = \max((1 - \epsilon)\mathsf{price}(S_j), \mathsf{value}(S_i))$. If a competitive seller S_k does not lower its price for n (*e.g.*, 4) rounds, S_i aborts the negotiation with S_k. S_i aborts negotiations with all competing sellers after a preset time limit. [2]

5 Experimental Results

Each *experiment* consists of a series of *tests*, while each *test* corresponds to a specific parameter setting (*e.g.*, value for discount rate and cost per message). Since mixed strategies (*i.e.*, those involving randomness) are common, and the agents' valuations and demand are also randomly generated, the profit measured in a single run is often noisy. To ensure statistically significant

[2]These safeguards were designed to ensure an agent would not't spend too much when "spying" from its competitors.

results, we run 200 *trials* per test. For each test, we report the average value obtained in the trials as well as 90% confidence intervals.

At the beginning of each trial we create an economy of $k \times n$ buyers and $k \times m$ sellers, where n and m are the number of buying and selling strategies, respectively, and k is a constant — we create k buyers or sellers for each strategy. In our experiment, $n = 4$, $m = 5$ (the UNDERCUT COMPETITOR strategy is for sellers only), and $k = 2$. Valuations for buying agents are randomly generated from a Gaussian distribution, but all sellers are assigned the same valuation in a trial to facilitate comparison. The demand of the buying agents is also randomly generated (from the normal distribution $N(100, \delta)$). Since each transaction is restricted to one widget, buyers must go through a number of negotiations to complete their transactions. Each trial is simulated for a preset number of cycles (*e.g.*, 1000). To compare different strategies independent of the specific demand and valuation, we normalize the profit received by a *selling* agent, S_x, in a trial:

$$\mathsf{profit\text{-}percent}(S_x) = \frac{\mathsf{profit}(S_x)}{\sum_{i=1}^{km} \mathsf{profit}(S_i)} \tag{2}$$

over all $k \times m$ selling agents in the economy.

5.1 Experiment 1: Strategy Performance as a Function of Discount Rate

In our first experiment, we fix the cost per message, c, and see how the strategies perform as the discount rate, r, varies. Figure 1 shows the mean performance result (with 90% confidence intervals), when $c = 1.0$. [3]

Figure 1 yields two important lessons. First, the UNDERCUT COMPETITOR strategy performs significantly better than its competitors. Second, the advantage accrued by UNDERCUT COMPETITOR diminishes as the discount rate increases. We elaborate on these points below.

The UNDERCUT COMPETITOR strategy shows a significant lead among the five strategies. By using the information it "spies" from its competitors, the agent is able to offer a slightly lower price and thus attract more than 57% of the sales in every test. The FIXED PRICE and FIXED OFFER strategies attract around 12% of the sales each. By committing to a fixed profit margin, their performance is relatively stable compared with the other strategies. In

[3] The results in this section all refer to the selling strategies. The curves did not change significantly for other values of c.

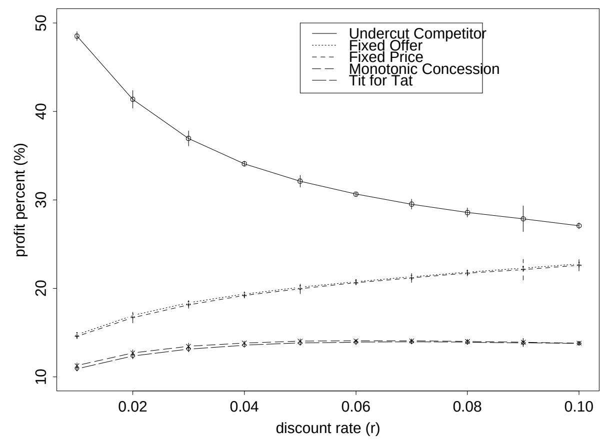

Figure 1: Strategy performance as a function of discount rate

fact, their performance almost matches the UNDERCUT COMPETITOR strategy when the discount rate is significantly large. The MONOTONIC CONCESSION strategy and TIT FOR TAT strategy yield the lowest market share (9%) and the lowest performance. We suspect that these two strategies do not interact well with all buying strategies.

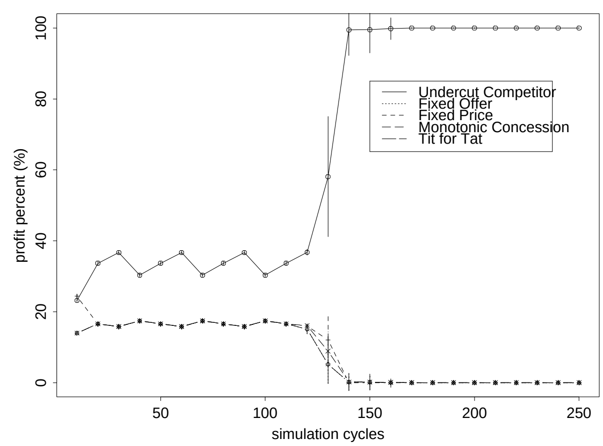

Figure 2: Performance of the UNDERCUT COMPETITOR strategy improved as it gathered information on it's competitors

While UNDERCUT COMPETITOR is dominant given a small discount rate, it loses its edge as r grows. We suspect that this is due to the fact that UNDERCUT COMPETITOR only beats the other strategies once it has gained information about their offering prices. This means that a disproportionate share of UNDERCUT COMPETITOR's profit is gained late in the simulation

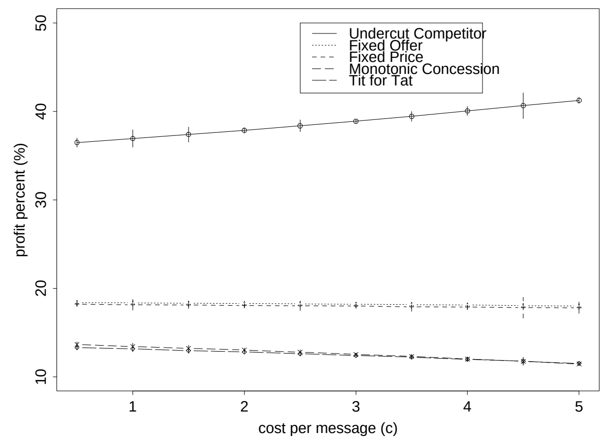

Figure 3: Strategy performance as a function of cost per message

(Figure 2). This makes sense in classical economic terms — UNDERCUT COMPETITOR is making an investment (by gathering information) which pays off after a time delay. A high discount rate always biases against delayed returns.

5.2 Experiment 2: Strategy Performance as a Function of Cost Per Message

In our second experiment, we fix the discount rate, r, in order to see how the strategies perform as the cost per message, c, varies. Figure 3 shows the mean performance result (with 90% confidence intervals), when $r = 0.03$.

As one can see, the cost per message, c, is not as a crucial factor as the discount rate, r. This matches expectations since c is only a linear factor whereas r is an exponential factor (Equation 1). The surprising fact is that UNDERCUT COMPETITOR performs *better* as message cost increases. Since UNDERCUT COMPETITOR agents send more messages (in their attempt to learn the asking price of their competitors), one might expect them to degrade as c rises.

Actually, the explanation is quite simple. As Figure 4 shows, an increase in the cost per message deteriorates the *absolute* performance of all strategies. However, in the case of UNDERCUT COMPETITOR, the overhead due to message cost is a small percentage of a relatively large profit. In contrast, the overhead hurts the other strategies more seriously since it represents a larger proportion of their (smaller) profit. As a result, the gap between the UNDERCUT COMPETITOR strategy and the other strategies enlarges as the message cost increases.

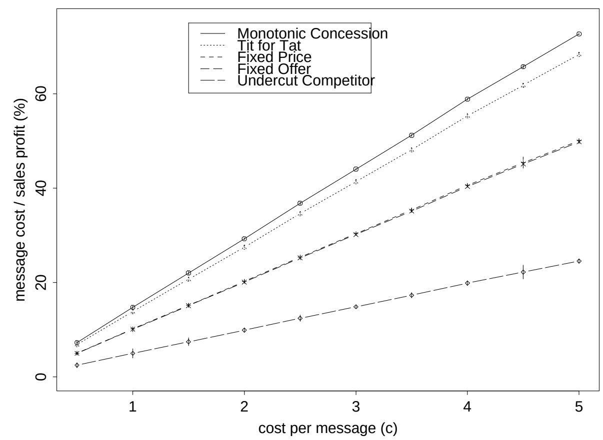

Figure 4: Ratio of message cost & sale profit as cost per message varies

6 Related Work

In Section 2 we describe relevant work from the field of economics, but there has also been considerable work on negotiations in the distributed AI community [1] as well. Davis and Smith's contract net protocol [4] provides an auction mechanism for agent negotiation using coordinated dialogue interactions comprising announcement, bidding, and award of contracts. Sandholm refined the bidding and awarding process using marginal cost calculations [15].

Shoham and Tennenholz claim that a society of automated agents require social laws to constrain their behavior; they present a general model of social law in a computational system and investigate its properties [17]. Shoham [16] advocates a computational framework called agent-oriented programming (AOP), where agents are defined in terms of beliefs, capabilities, decisions, and obligations. Agents communicate with each other via message passing, where message types (request, inform, promise, *etc.*) are drawn from the speech act theory [3]. Sidner extended AOP and proposed an artificial negotiation language that models human negotiations of collaboration [18]. Although we implemented our bargaining strategies in plain Common Lisp, we could have easily expressed them in the AOP framework instead.

Rosenschein and Zlotkin have applied economic and game theory to formal analysis of protocols for automated negotiation among agents [11]. They proposed evaluation schemes for protocols in terms of efficiency, stability, and simplicity. For example, the Vickrey auction mechanism (best bid wins, gets second price) is attractive since it provides no incentive for an agent to underbid or overbid [12]. Rosenschein and Zlotkin also identified the connection between protocols and domains, and categorized domains into a hierarchy of task-oriented domains, state-oriented domains, and worth-oriented domains. In addition, they address the case of incomplete information and consider the design of deception-free protocols. Our model of electronic commerce can be seen as a simple form of worth-oriented domain. However, unlike Rosenschein and Zlotkin, we allow explicit utility transfer. Since buyer and seller objectives are inherently adversarial, we conjecture that no stable, distributed protocol exists for the domain.

There are many existing approaches that we might adopt to make our bargaining agents more intelligent and sophisticated. For example, Kraus *et al.* built an agent that played the game of Diplomacy [8], which involved negotiation of sophisticated military agreements. It may be possible to adapt her structure and some of her negotiation heuristics into an automated bargaining agent. We might also adopt Gmytrasiewicz *et al.*'s recursive modeling method so our agent could use its predictions of other agents' responses when planning its actions [5,5]. In addition, Kraus *et al.* presented a strategic model of negotiation that takes time into account during the negotiation process [9]. Such mechanism could be applicable to our experiments as they get more dynamic and time-sensitive.

7 Conclusion

In this paper, we advocated a new challenging domain for AI applications, and made the first steps towards building an intelligent bargaining assistant for electronic commerce. We reviewed the literature in economics and game theory, presented a simple model of electronic commerce and a negotiation protocol, and developed several bargaining strategies under this model. To evaluate the bargaining strategies, we developed a general experimental framework and evaluation scheme. Our preliminary results show that complex, adaptive strategies (such as Undercut Competitor) can perform significantly better than simple discounting approaches. While an increased discount rate diminishes some of the advantage accrued by Undercut Competitor, it still dominates the other strategies tested. Perhaps surprisingly, the Undercut Competitor strategy's relative performance *increased* with higher message cost — despite the fact that it sent more messages than other strategies.

Since our investigation has just begun, we make no claims about the optimality of Undercut Competitor or other strategies. While complex adaptivity appears useful, there may exists simple, parasitic strategies which perform still better. We do believe, however, that our work provides a useful framework for future investigation. We are devising additional strategies and will evaluate their performance. We are also planning several new experiments. For example, we will measure how the performance of strategies varies as the economy becomes supply limited rather than demand driven; the optimal behavior is very different in these two cases so an adaptive strategy could be quite interesting. We will also consider the case in which product valuations vary dynamically and new agents are created as simulation progresses.

We also wish to consider more elaborate protocols and analyze their stability and convergence properties. At present, one agent can only communicate an offer or acceptance to another. It would be more interesting if buyer B_1 could tell seller S_2 that "Seller S_3 is only charging \$$N$ per widget." In fact, a seller might be willing to pay for this type of information (since similar information was useful in the UNDERCUT COMPETITOR strategy). Furthermore, an independent agent might be willing to pay for this information and sell it to others.

This leads to the idea of building a consumer reports agent. We hope to develop pricing formulae for determining the utility of accurate information about the valuation of adversarial agents. Generalizing the protocol to allow pricing of information in addition to widgets raises questions about the form of contracts agents can make with each other, and many knowledge representation issues are involved.

References

[1] A. Bond and L. Gasser. *Readings in Distributed Artificial Intelligence.* Morgan Kaufman, San Mateo, CA, 1988.

[2] K. Chatterjee and L. Samuelson. Bargaining Under Two-Sided Incomplete Information: The Unrestricted Offers Case. *Operations Research*, 36(4):605–618, 1988.

[3] P. R. Cohen and C. R. Perrault. Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science*, 3(3):177–212, 1979.

[4] R. Davis and R. G. Smith. Negotiation as a metaphore for distributed problem solving. *Aritificial Intelligence*, 20(1):63–109, 1983.

[5] P. J. Gmytrasiewicz, E. H. Durfee, and D. K. Wehe. A decision-theoretic approach to coordinating multiagent interactions. In *Proc. 12th Int. Joint Conf. on A.I.*, pages 62–68, August 1991.

[6] D. R. Hofstadter. The prisoner's dilemma computer tournaments and the evolution of cooperation. In *Metamagical Themas: Questing for the Essence of Mind and Pattern*, chapter 29, pages 715–734. Basic Books, New York, 1985.

[7] Rust J., J. Miller, and R. Palmer. Characterizing effective trading strategies: Insights from a computerized double auction tournament. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.

[8] S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *To appear in Computational Intelligence*, 1994.

[9] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *To appear in Artificial Intelligence*, 1994.

[10] E.. Rasmusen. *Games and Information.* Blackwell, Oxford, UK, 1989.

[11] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* The MIT Press, Cambridge, MA, 1994.

[12] J. S. Rosenschein and G. Zlotkin. Designing conventions for automated negotiation. *AI Magzine*, 15(3):29–46, Fall 1994.

[13] A. Rosenthal. Telephone wars: Who offers the best long-distance rates? *Consumers Digest*, pages 30–33, December 1994.

[14] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

[15] T. Sandholm. An implementation of the contract net protocol based on marginal cost. In *Proc. 11th Nat. Conf. on A.I.*, pages 256–262, 1993.

[16] Y. Shoham. Agent-oriented programming. *Artificial Intelligence*, 60(1):51–92, March 1993.

[17] Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies (preliminary report). In *Proc. 10th Nat. Conf. on A.I.*, 1992.

[18] C. Sidner. An artificial discourse language for collaborative negotiation. In *Proc. 12th Nat. Conf. on A.I.*, pages 814–819, 1994.