



Emergent Properties of a Market-based Digital Library with Strategic Agents¹

SUNJU PARK

spark@business.rutgers.edu

Management Science and Information Systems Department, Faculty of Management, Rutgers University, Newark, NJ 07102

EDMUND H. DURFEE

durfee@umich.edu

Artificial Intelligence Laboratory, University of Michigan, Ann Arbor, MI 48109-2110

WILLIAM P. BIRMINGHAM

wpb@umich.edu

Artificial Intelligence Laboratory, University of Michigan, Ann Arbor, MI 48109-2110

Abstract. The University of Michigan Digital Library (UMDL) is designed as an open system that allows third parties to build and integrate their own profit-seeking agents into the marketplace of information goods and services. The profit-seeking behavior of agents, however, risks inefficient allocation of goods and services, as agents take strategic stances that might backfire. While it would be good if we could impose mechanisms to remove incentives for strategic reasoning, this is not possible in the UMDL. Therefore, our approach has instead been to study whether encouraging the other extreme—making strategic reasoning ubiquitous—provides an answer.

Toward this end, we have designed a strategy (called the p-strategy) that uses a stochastic model of the market to find the best offer price. We have then examined the collective behavior of p-strategy agents in the UMDL auction. Our experiments show that strategic thinking is not always beneficial and that the advantage of being strategic decreases with the arrival of equally strategic agents. Furthermore, a simpler strategy can be as effective when enough other agents use the p-strategy. Consequently, we expect the UMDL is likely to evolve to a point where some agents use simpler strategies and some use the p-strategy.

Keywords: multi-agent systems, digital libraries, strategic reasoning, emergent behavior

1. Introduction

Mechanism design and individual-agent design are two primary design issues when building an open multi-agent system consisting of self-interested agents. System designers devise a mechanism that defines who can communicate with whom and what message types are allowed. Then, agent designers (who do not necessarily share common goals with system designers) develop self-interested agents to achieve their individual goals under the given mechanism.

Of course, these two design issues are interdependent; a well-designed mechanism can simplify the design of individual agents (and vice versa). Over a quarter of a century ago, for example, Vickrey devised an auction mechanism (now sometimes

called the Vickrey auction) where the best bid would win the auction and the payment to or by its maker would be the amount of the best losing bid. As he analyzed [24], the Vickrey auction mechanism makes rational agents bid their true reservation prices, such that even self-interested agents, if they are rational, will behave honestly and not try in vain to outsmart other agents. Therefore, the Vickrey auction mechanism trivializes the design of each individual agent's decision making.²

However, designing a good mechanism that exhibits certain properties—which is called incentive engineering—is difficult, especially for dynamic systems where the participants and their interactions evolve over time. The example of such dynamic systems that we use throughout this paper is the University of Michigan Digital Library (UMDL). In UMDL, we aim to provide an infrastructure for rendering library services in a networked information environment [4]. The UMDL is designed as a multi-agent system, where agents (representing users, collections, and services of the digital library) sell and buy information goods and services through auctions. While supporting flexibility and scalability, the open multi-agent architecture and market infrastructure create dynamics (agents participating in an auction change, matches between buyers and sellers vary, and auctions themselves evolve), which adds additional complexity to mechanism design.

As system architects, we strive for an efficient system. Although the UMDL market allows the self-interested agents to seek profits, we do not want strategic agents to undermine the overall system performance (efficiency in market), nor such agents to reap excessive profits from other agents (efficiency in allocation). That is, we want an incentive-compatible mechanism that makes strategic reasoning unnecessary [27]. Unfortunately, we do not have such a mechanism yet for the UMDL system; we fully expect strategic agents who try to take advantage of other agents to enter the system. Note that the potential of having strategic agents is not limited to the UMDL auction. Most open multi-agent systems expect strategic behavior by the agents, who try to maximize their utilities given the current mechanism.

So, does this mean the UMDL will become inefficient? Will (and should) a UMDL agent spend much of its computational power trying to outsmart other agents? What happens if all the agents behave strategically? In this paper, we answer these questions, by studying the properties of the UMDL with strategic agents. Developing a mechanism that prevents strategic thinking is a hard problem, so instead we use a bottom-up approach. We design a strategy that the UMDL agents may use, and experiment with such strategic agents to learn about the system properties. In particular, we are interested in knowing whether making strategic reasoning ubiquitous (instead of preventing it) reduces its negative effects.

In this paper, we briefly examine the target system, the UMDL service market society. We describe an agent bidding strategy called the *p-strategy* and demonstrate its advantages over other simpler strategies in the context of the UMDL auction. Then, by experimenting with multiple *p-strategy* agents, we investigate some emergent properties of the UMDL auction. Finally, we review some of the previous work on multi-agent system design issues and conclude our paper.

2. The UMDL service market society

The UMDL provides digital library services in a distributed information environment. Given the magnitude of the number of contents and services available, the rate of change in what is available, the size of a user population, and the evolving nature of that population, a digital library faces a great challenge in its administration efforts. One possible approach is to move as much of the administration into the infrastructure as possible. The infrastructure should encourage the following things [5].

- **Flexibility:** It should be able to embody a wide variety of policies to realize different types of libraries (e.g., public, corporate, university, personal, etc.).
- **Extensibility:** Providers and consumers of information goods and services should have incentives to join the library and should be capable of finding their counterparts.
- **Scalability:** As the number of users, goods, and services grows, the underlying, computerized administration of the library should not bog down.

Toward this end, the UMDL is structured as a collection of agents that buy and sell services from each other using an auction infrastructure. We refer to this infrastructure as the Service Market Society (SMS). Instead of relying solely on internally-designed agents, the UMDL can attract outside agents to provide new services, which are motivated by the profit they might accrue by participating in the system. Since the UMDL is open, all agents are treated as self interested.

We distinguish three broad classes of agents populating the UMDL: User Interface Agents (UIAs), Collection Interface Agents (CIAs), and Mediator Agents. Users, or library patrons, have individual UIAs, which interact with the UMDL on their behalf to acquire library services. CIAs, representing the publishers and other owners of these collections, provide access and search services for the library collections. Various middlemen, or Mediator Agents, perform a variety of services that end-users or other agents desire. An example is Query Planning Agents (QPAs), which act as middlemen, accepting queries from UIAs and returning the list of CIAs related to the queries. Another type of mediator agent is the Service Classifier Agent (SCA), which provides ontology services.

Figure 1 shows a simple agent-interaction scenario in the UMDL Service Market Society (SMS). In this scenario, users want to find sources of information for various topics (e.g., history, science, or mathematics) and various audience levels (e.g., middle school, high school, or professional). Initially, the user sends a query to the UIA. The UIA must then find an agent that can service this query. The SCA recommends the most appropriate services out of those that are currently available using its base ontology (Step 1). Once the UIA has acquired the appropriate service-label from the SCA for the service it is looking for, it needs to locate the sellers of that service (QPAs in this example). The UIA sends the desired service-label to the Auction Manager Agent (AMA), which returns a list of auctions that sell the appropriate query planning services (Step 2). Then, the UIA sends buy offers to a particular Auction Agent (Step 3). The Auction Agent collects




Figure 1. Agent interaction in the UMDL Service Market Society. Steps 1 and 2 are the finding phase, where agents find the auction to participate in. Step 3 is the negotiation phase, where agents settle on the price of the query service they are buying or selling. Steps 4 and 5 are the query phase, where the actual query is forwarded and the documents are returned. This paper focuses on the negotiation phase of Step 3.

offers from the participating agents and determines matches consistent with those offers. Once a match is found, the UIA will send its query to the matching QPA (Step 4). The QPA returns the list of appropriate CIAs, and the UIA then forwards its query to them (Step 5). The CIAs can translate UMDL queries to a variety of protocols (e.g., http, Z39.50, FTP, etc.) and return the appropriate documents.

Each activity in Figure 1 addresses different research issues, such as how to describe what agents buy or sell, when and how to create an auction, how to locate the right auction to participate in, and so on. Since the focus of this paper is on multi-agent auctions, we assume that agents have found the right auction and are ready to buy and sell in an auction, while ignoring other issues.

Selling and buying of services are done through auctions, operated by Auction Agents. Figure 2 shows an example of a UMDL auction, where User Interface Agents (UIAs) want to find sources of information for some topic (say, science) on behalf of certain kinds of users (say, high school), and some Query Planning

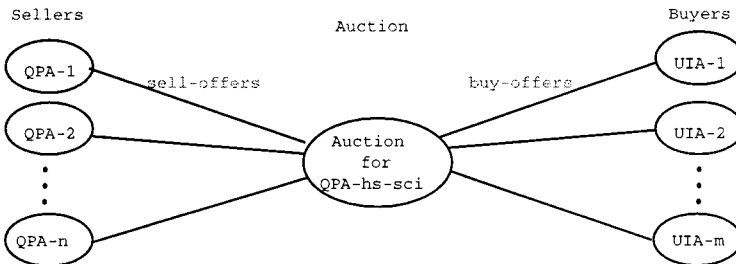


Figure 2. The UMDL double auction for “QPA-service_high-school_science.” In this auction, there are n sellers (QPAs) and m buyers (UIAs). Additional buyers and sellers may join, and the existing buyers and sellers may leave the auction at any time (i.e., entry-and-exit is allowed).

Agents (QPAs) sell the services for finding such collections. In this auction, there are n sellers (QPAs) and m buyers (UIAs). Additional buyers and sellers may join, and the existing buyers and sellers may leave the auction at any time.

The UMDL auction is a *continuously-clearing double auction* (CDA). Compared to more familiar one-sided auctions, the UMDL is a double auction where both buyers and sellers post their buy prices and sell prices, respectively. It is a continuously-clearing auction, as every bid could trigger a potential clear, depending on whether it matches an existing bid. In the UMDL, the Auction Agent continuously matches the highest buyer to the lowest seller, given that the buy price is greater than the sell price. The clearing price is based on the seller's offer price (i.e., buyers receive the entire surplus).³ The CDA supports dynamic pricing, and is well suited for frequent, timely transactions needed in information economies, such as document-delivery services provided by the UMDL.

An incoming offer that cannot be matched with the existing offers becomes a new standing offer. Since buyers (sellers) with bid prices higher (lower) than any standing sell (buy) offer get matched, the buyers' standing offers in the auction (if any) always have lower offer prices than the sellers' standing offers (if any). That is, standing offers ordered by lowest to highest bid are always in a (bbb ... bsss ... s) sequence, as shown in Figure 3.

3. P-strategy: An agent bidding strategy based on stochastic modeling

In the previous section, we have examined how the UMDL auction works. In this section, we describe an agent bidding strategy based on stochastic modeling (called the p-strategy).

The four-step p-strategy is as follows. First, a p-strategy agent models the auction process using a Markov chain (MC) with two absorbing states (success and failure). Second, it computes the transition probabilities between the MC states. Third, it computes the probabilities and the payoffs of success and failure. Finally, it finds the best offer price to maximize its expected utility.

In this section, we describe how to build the MC model for the UMDL auction, while skipping the details of how to define the exact transition probabilities between the MC states (step 2), how to compute the probabilities and the payoffs of success and failure (step 3), and how to find the best payment (step 4). More details can be found in one of our papers [15].




Figure 3. Standing offers in the auction. At first, one buy offer (b1) and one sell offer (s1) are standing at the auction. When a sell offer (s2) with a higher price than the current highest standing buy offer arrives, it becomes a new standing offer. When a buy offer (b2) that is higher than the current lowest standing sell offer (s1) comes, the auction agent matches those two offers.

The main idea behind the p-strategy is to capture the factors that influence the expected utility in the MC model of the auction process. For instance, a seller is likely to raise its offer price when there are many buyers, or when it expects more buyers to come. The MC model takes those factors into account in the MC states and the transition probabilities. The number of buyers and sellers standing at the auction, the arrival rates of future buyers and sellers, and the distribution of buy and sell prices are among the factors modeled in the MC model.

Note that, in the UMDL auction where the clearing price is determined at the seller's offer price, sellers have somewhat stronger incentives to bid above cost, as they always set the market price and this affects their tradeoff between the probability of trading and the profit earned [11, p. 288]. Therefore, we use the p-strategy *seller* when explaining the MC model of the auction process. Of course, buyers have incentives, although somewhat weaker, to behave strategically as well (i.e., to bid below valuation), and the p-strategy buyer's MC model can be constructed in a similar fashion.

Each state in the MC model represents the status of the auction. The (bbs^{*}) state, for example, represents the case where there are two standing buy offers and two standing sell offers and the sell offer of the p-strategy seller doing the reasoning (represented as s^{*}) is higher than the other seller's offer. If we assume that offers arrive at most one at a time, the auction can go to any of the following states from the (bbs^{*}) state (see Figure 4):

- (bbs^{*}): No offer arrives during the clearing interval.
- (bbs^{*}): A buy offer arrives, and it is matched with the lowest seller.
- (bbbs^{*}): Because of no match, a new buy offer becomes a standing offer.
- (bss^{*}): A sell offer arrives, and it is matched with the highest buyer.
- (bbs^{*}s): Because of no match, a new sell offer becomes the highest standing offer.
- (bbbs^{*}s): A new sell offer becomes a standing offer, but the p-strategy seller's offer is still the highest.

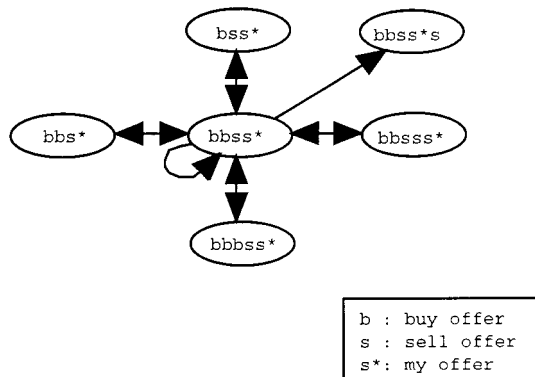


Figure 4. Transitions from the (bbs^{*}) state. Each transition represents the clearing of the auction and the arrival of a new buy or sell offer, if any.

Figure 5 depicts the MC model for the UMDL auction with a maximum number of standing buy and sell offers of five each. That is, to manage the size of the auction, we limit the number of buy and sell offers standing in the auction not to exceed five each. When an additional buyer (seller) arrives, the buyer (seller) with the lowest (highest) offer will be kicked out of the auction first.

The number of MC states increases with the size of the auction. When the maximum number of standing offers is limited to m buyers and n sellers, the number of MC states is $(m + 1) \times ((n + 1)n)/2 + 2$. Of course, one may shrink the size of the MC model, while sacrificing the accuracy of the model. For example, the p-strategy seller may consider the states with more than a certain number of sellers with the offer prices lower than its offer price as a failure. To have a chance of getting matched, the p-strategy seller should wait until all the sellers with lower offer prices are matched (which may take a long time), so considering those states as a failure is a good approximation in practice.

Using the MC model and its transition probabilities, the p-strategy agent is able to capture various factors that influence the utility value and tradeoffs associated with those factors. Figure 6 shows an example of tradeoffs between the number of standing buyers and sellers when deciding on the best offer price. In general, the seller raises its best offer price when there are more standing buyers (to increase the profit of a possible match). When the number of standing sellers is five (at the right end of the graph), however, the p-strategy seller bids a lower price when there is one buyer than when there is no buyer. That is, the p-strategy seller lowers its offer to increase the probability of a match (instead of increasing the profit of a match). Offering a higher price in this case would have served to price it out of the auction when it might otherwise have been able to trade profitably.

Intuitively, agents with complete models of other agents will always do better in the auction, but without repeated encounters complete models are unattainable. In the UMDL, an agent in its lifetime meets many different agents, and as a result its model of other agents is incomplete. Instead of modeling individual agents, therefore, the p-strategy uses a model of the auction process.

4. Advantages of the p-strategy

The p-strategy models the auction process stochastically and thus ignores the fact that the other agents in the auction may also behave strategically. This might sound like a bad engineering decision, especially for the UMDL auction where agents do try to outsmart other agents.

In this section, we demonstrate that the p-strategy works well in the dynamic UMDL auction, by comparing the profit of the p-strategy seller (p-seller) with three different types of sellers. They are:

- A seller who bids its cost plus some fixed markup (FM-seller),
- A seller who bids its cost plus some random markup (ZI-seller),⁴ and
- A seller who bids the clearing price of the next transaction (CP-seller).⁵

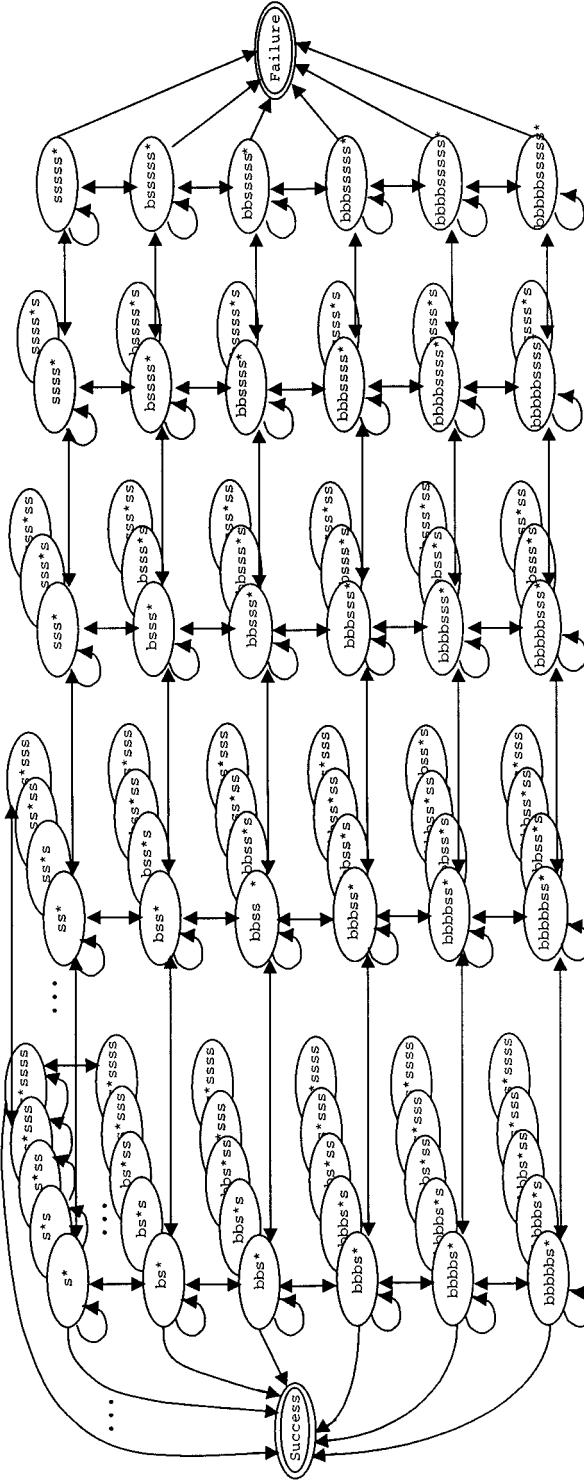


Figure 5. The MC model for the UMDL auction. A transition to the success state happens when the p-strategy seller's offer (s^*) gets matched with an incoming buy offer. A transition to the failure state happens when the p-strategy seller's offer (s^*) is kicked out of the auction.




Figure 6. Tradeoffs between the number of buyers and sellers standing at the auction. The best offer price is determined based on the tradeoffs between the probability of a match and the profit of a match.

4.1. Experimental setting

The experimental settings are as follows.

- Auction

The auction clears every three seconds. The value of the clearing interval does not affect the results from our experiments, provided that on average at most one offer arrives at the auction each interval. When more than one offer arrives at an auction, they are saved in a queue and retrieved by the Auction Agent one by one.

- Buyers

A single agent simulates multiple buyers by submitting multiple bids. In our experiments, every six seconds the buyer submits its bid with a probability of 0.8. By adjusting the offer interval and the offer rate of the single buyer, we can change the arrival rate of buy offers to the auction.⁶ The buyer offers valuations that are drawn randomly from a uniform distribution between 10 and 30.

- Sellers

For each experiment, we compare the profits of two sellers: p-seller and the opponent (either FM-seller, ZI-seller, or CP-seller). Both sellers submit their bids every 24 seconds on average. In addition, similar to the buyer case, a single agent simulates all the other sellers at the auction. Its offer interval and offer rate are set at 12 and 0.8, respectively.

The costs of all the sellers are based on their loads, which are computed from the message traffic and the current workload. That is, $\text{cost} = \alpha \times (\text{number of messages per minute}) + \beta \times (\text{number of matches per minute})$. The cost function




Figure 7. Comparison of the profits of the sellers.

represents economies of scale. As we expect the workload from matches to be higher than that from communication in most cases, we set α to one, and β to five for our experiments.

4.2. Experimental results

In the first set of experiments, we have compared the profits of the p-seller and the FM-seller. When competing with FM-sellers with various markups, the p-seller always gets a higher profit. This is not surprising, since the p-seller is able to use extra information about the auction. Figure 7(a) shows the accumulated profits of the p-seller and the FM-seller who bids its cost plus seven as its offer (i.e., fixed-markup = 7). The best markup will vary depending on the auction situations, but seven is the best markup for this particular auction setting.

Figure 7(b) shows the profits of the p-seller and the ZI-seller. The ZI-seller works poorly against the p-seller, which indicates the randomization strategy does not work. The ZI-seller can be thought of as a naive strategy that fails to take advantage of potential profit margin [8].

In the final experiment, we have compared the p-seller with the CP-seller. The CP-seller receives a price quote from the auction—the clearing price were the auction to clear at the time of the quote—and submits it as its offer as long as it is higher than its cost.⁷ Note that the clearing price is hypothetical, since it changes when new offer(s) arrive during the time between the clearing-price quote and the CP-seller's offer.

As shown in Figure 7(c), the p-seller usually gets a higher profit when competing with the CP-seller. Since the CP-seller gets more matches than the p-seller (but its profit per match is smaller) on average, however, the CP-seller works better when getting more matches does not increase its cost much (e.g., when β is zero). Bidding the next clearing price seems like a good heuristic, especially when clearing prices do not change much (i.e., the hypothetical clearing-price quote is more likely to be the actual clearing price).

The p-strategy works well in the UMDL auction due to its dynamics. No agent can have a complete, deterministic view about the current and future status of the

auction, and naturally, an agent strategy should be able to take into account the dynamics and the resulting uncertainties. In our experiments, the p-strategy that models the auction process stochastically receives higher profit than the other strategies. In particular, the p-strategy seller works well when the auction is dynamic and the negotiation zone between the buyer and the seller is wide.

5. Collective behavior of p-strategy agents in the UMDL auction

Given that the p-strategy is effective in the UMDL auction (from the previous section), nothing prohibits any self-interested agent from adopting the p-strategy. We expect many p-strategy agents to coexist in the UMDL, and thus are interested in the collective behavior of such agents. In this section, therefore, we investigate (1) how the absolute and relative performance of a p-strategy agent changes against other p-strategy agents, and (2) how the UMDL is affected by multiple p-strategy agents.

5.1. Experimental setting

Figure 8 shows six experimental settings with seven buyers and seven sellers. Although we fix the number of buyers and sellers, by changing their offer rates and offer intervals, we can simulate a large number of agents and different levels of activities in the auction. By increasing the offer rates or decreasing the offer intervals, for instance, we can simulate a more dynamic auction. In this set of experiments, the portion of p-strategy agents in the total population is important. It is also the case that a large number of agents who bid infrequently have the same effect on the auction as a small number of agents who bid more frequently.

In this experiment, we deliberately set supply to be higher than demand to emphasize competition among sellers; each buyer submits its bid every 30 seconds with a probability of 0.5, while the offer interval and the offer rate of each seller are set to 30 seconds and a probability of 0.7, respectively. Each experiment ran for five hours.

Session Description		Bidding strategy						
		Seller1	Seller2	Seller3	Seller4	Seller5	Seller6	Seller7
(1)	All competitive sellers	C	C	C	C	C	C	C
(2)	1 p-strategy seller	C	C	C	C	C	C	P
(3)	2 p-strategy sellers	C	C	C	C	C	P	P
(4)	3 p-strategy sellers	C	C	C	C	P	P	P
(5)	5 p-strategy sellers	C	C	P	P	P	P	P
(6)	All p-strategy sellers	P	P	P	P	P	P	P

Figure 8. Experimental setting.

The buyers bid their true valuations, while the sellers bid their sell prices based on their strategies. In Session 1, all seven sellers bid their true costs. Since traders honestly report their reservation prices, Session 1 gets the most matches and serves as a benchmark for comparing market efficiency. From Session 2 through Session 6, we introduce more p-strategy agents into the auction.

In the experiments, we measure efficiency of the system in two ways. First, we measure the efficiency of allocation, by comparing the p-strategy seller’s absolute and relative profits. This measurement indicates how the p-strategy seller’s performance changes with an increasing number of other p-strategy sellers in the system. Second, we measure the efficiency of the market, using the number of matches made and the total profit generated. The market-efficiency result shows how much having multiple p-strategy sellers undermines the overall system performance.

5.2. *Experimental Results*

Although we have shown that the p-strategy agent has an upper hand over other-strategy agents, this advantage may not hold in the presence of other p-strategy agents. To test this, we compare the profits of Seller 7 (p-strategy agent) across Sessions 2 to 6. As shown in Figure 9, the profit of the p-strategy (smart) agent decreases as the number of p-strategy (equally smart) agents increases.

Now that we have established that the profit of the p-strategy agent decreases as more agents use p-strategy, another question arises. How will a simpler strategy agent perform in the presence of multiple p-strategy agents? In Figure 10, by replacing Seller 1 with the fixed-markup seller (with markup = 5), we measure the relative performance of the FM-seller and the p-seller. The FM-seller’s profit is generally less than that of the p-seller, but the difference decreases with an




Figure 9. Profits of p-seller.




Figure 10. Profits of FM-seller and p-seller.




Figure 11. Efficiency measured by the number of matches in the auction process.

increase in the number of p-strategy agents.⁸ That is, the disadvantage of being “less smart” decreases as the number of smart agents increases.

This result indicates that an agent may want to switch between using the p-strategy and using a simpler strategy depending on what the other agents are doing. By dynamically switching to a simpler strategy, an agent can achieve a similar profit to that of using the p-strategy, while exerting less effort (time and computation) on computing bids.

We also measure the number of total transactions made in the auction process, as shown in Figure 11. When the number of p-strategy agents increases, the number of matches decreases, since the p-strategy agents usually get fewer matches but more profit per match. The number of matches is not an exact indication of the market efficiency, as the average profit per match differs across the different settings. If the Auction Agent requires a service charge per transaction, however, it may serve as an indication of how much it can profit from the auction process.

In addition, we measure the market efficiency using the total profit accrued by buyers and sellers (see Figure 12). The buyer’s profit is its valuation minus the clearing price when the deal is made, and zero otherwise. The seller’s profit is the clearing price minus its cost when deal is made, and zero otherwise. The total profit eventually decreases with increasing numbers of p-strategy agents, as the market becomes inefficient due to strategic misrepresentation of p-strategy agents (and, therefore, missed opportunities of matches).

The total profit, however, does not decrease as sharply as we had expected. It turns out that the CDA, even without strategic agents, is not very efficient because it continuously clears rather batching offers and finding the most efficient matches possible. Hence, the introduction of more strategic agents does not make a significant difference. A periodic, clearing-house style double auction is unrealistic




Figure 12. Efficiency measured by the total profit.

for the UMDL system where the participants of the auction change and matches should be made quickly.

We conjecture that having multiple p-strategy sellers (compared to naïve-strategy sellers) poses interesting tradeoffs between strategic inefficiency and surplus extraction. By misrepresenting their true costs, the strategic sellers miss possible transactions (i.e., strategic inefficiency). By anticipating the future arrival of buyers, on the other hand, they are able to seize more surplus (i.e., surplus extraction).

5.3. *Lessons learned*

A conventional way of designing a system that exhibits certain properties is to engineer it. Incentive engineering, however, is unsuccessful in the UMDL system because of the system's complexity and dynamics. Instead, by making the p-strategy available to the agents, we have studied the effects of ubiquitous strategic agents in the UMDL system.

We summarize our observations as follows. First, although a self-interested agent in the UMDL system has the capability of complex strategic reasoning, our experiments show that such reasoning is not always beneficial. As shown in Figure 9, the advantage of being smart decreases with the arrival of equally smart agents.

Second, if all the other agents use the p-strategy, an agent with a simple strategy (e.g., fixed markup) can do just as well, while incurring less overhead to gather information and compute bids. This result indicates that an agent may want to switch between a complex strategy and a simple one depending on the behavior of other agents. As the overhead of complex reasoning becomes more costly, an adaptive strategy that dynamically decides on which strategy to use will be more desirable.

Third, we expect the UMDL is likely to evolve to a point where some agents use simpler strategies while some agents use more complex strategies that are more knowledge intensive (such as the p-strategy). It follows from the above observation that if enough other agents use complex reasoning, an agent can achieve additional profit even when it continues using a simple strategy.

Finally, the market efficiency of the UMDL (represented by the total profit) will not decrease sharply with increasing number of strategic agents. As shown in Figure 12, having multiple p-strategy agents increases the market efficiency slightly up to some point. Moreover, the profit-seeking behavior of self-interested agents will keep the UMDL agent population mixed with agents of various strategies. Even though the market efficiency eventually decreases with an increase in the number of p-strategy agents, because of its mixed agent population, the UMDL will not suffer market inefficiency of the worst case.⁹

6. Related work

Early distributed artificial intelligence (DAI) research described how to apply the power of multiple, distributed agents to a problem. The emphasis of this research

was getting multiple agents to work *together* to solve a problem in a coherent, robust, and efficient manner. Typically, individual agents were assumed to be cooperative (e.g., share tasks, communicate truthfully), and exchange of tasks among agents was based on suitability—adjacency, processing capability, current agent load, etc. This line of work is often called cooperative distributed problem solving (CDPS). In mid eighties, however, a group of researchers challenged CDPS' assumption of cooperative agents. That is, agents are basically self-interested. This branch of work is called MAS.¹⁰

In cooperative systems (like the early DAI systems), a designer (or a group of designers sharing common goals) set shared goals and control both the global mechanism design and the individual-agent design. For instance, agents and mechanism can be built in certain ways (e.g., share information, be honest, and so on) conducive to cooperation. In MAS, on the other hand, the system architects have no (or limited) control over the preferences and goals of individual agents. Note that the goal of CDPS (achieving the competence of a system) differs from that of MAS (studying the properties of a system). The design problem in MAS is different from CDPS, as a result of relaxing the assumption on cooperative agents. In MAS with self-interested agents, system architects need to design a proper rule to ensure some desired property (honesty of agents, for example) rather than assuming it.¹¹

It is no surprise then that many researchers turned to game theory to design and analyze strategic interaction in MAS. A multi-agent interaction is in essence a game played among a set of agents, and such settings have been studied in economics, especially in noncooperative game theory [14, 22]. Many examples of game-theoretic approaches can be found in MAS. Rosenschein and Zlotkin identified two building blocks of a multi-agent system, protocol and strategy,¹² and focused on designing a protocol that "...motivates agents towards telling the truth..." [17]. Game-theoretic approaches for global mechanism design for MAS can be found in various papers [1, 20, 21, 23].

While game theory provides a symmetric, stable solution (e.g., a set of equilibrium strategies where no agent wishes to change its part of the strategy), it has some weaknesses, such as the rationality assumption (what happens if agents are not rational?), multiple equilibria (how to choose among them?), and so on. Furthermore, game theory tends to be applied to highly abstract and simplified settings, and can be difficult to apply to more complex situations. We do not use game theory because the UMDL auction is too complex to model and because we cannot assume rationality in the game-theoretic sense.

Note that the information required in game theory and our stochastic-modeling method, respectively, is different. In the former, all the information about strategies and payoffs of all the agents involved should be known. In the latter, we require less information by ignoring strategic behavior of other agents. This kind of strategy seems well-suited for the dynamic systems where modeling the dynamics is more important than modeling the individual behavior.

The use of microeconomics theory is not limited to the game-theoretic analysis of the MAS with self-interested agents. The MAS community has borrowed many other concepts from economics, such as the concepts of utility functions, payments,

marginal-cost calculations, supply and demand, consumer and producer (buyer and seller), auctions, and so on. Market-based approaches have become increasingly attractive to MAS researchers because of the ready availability of underlying formal models from microeconomics. Microeconomics provides techniques and theories that can be used in designing and analyzing both systems for CDPS and MAS with self-interested agents. In addition, market-based approaches have conceptual advantages; the decentralized approach of microeconomics corresponds well to MAS, and it is a natural fit for e-commerce transactions, as in the UMDL system.

We are particularly interested in multi-agent auctions because of their potential role in e-commerce. In physical markets dominated by standardized industrial goods and posted prices, examples of elaborate negotiation are limited to high-value items (e.g., houses, automobiles, collectibles, etc). In the electronic commerce, on the other hand, several factors, e.g., the trend toward customization, proliferation of agents, increasing interaction between buyers and sellers, favor negotiated prices [3]. As a result, developing mechanisms for negotiation and designing individual agents for automated negotiation has become an important research area. In this context, auctions are an important research paradigm, as they have regular yet flexible structure, the interaction protocols are clearly defined, and many variations of protocols can be instantiated.

The agents in the auctions are assumed to be self-interested—agents will try to maximize their utilities in all possible ways given the current interaction protocols. There is a potential gain for an agent in taking other agents into account. Research questions then are how to design an efficient strategy for self-interested agents and what are the system-wide properties of markets populated with self-interested agents.

To address the question of what an agent should do in multi-agent interactions, some researchers have developed a new solution concept based on a recursive model of other agents (i.e., what they think what I think and so on) [7]. While game-theoretic agents are ultra-smart and super-rational, such that they can reason about this whole recursive hierarchy ad infinitum [2], the recursive modeling method (RMM) assumes that agents can only build a finite nesting of models (because of practical limitations on acquiring such knowledge).

Basically, RMM is a decision-theoretic approach to game theory, which is advocated by some as a practical solution concept for developing an agent strategy. Our p-strategy can be viewed as 1-level RMM, as it models the behavior of other agents (but not the other agents' "interior" reasoning processes).

Rust et al. [18] have carried out a double auction tournament to investigate appropriate agent strategies. Their results show that a very simple "waiting in the background" trading strategy emerges as the winner of the tournament. We have not used this winning strategy in our experiments, as the auctions are differently configured. In Rust et al., the timing of a bid is an important decision factor, while timing is not a factor in the UMDL auction setting.

Another approach to designing multi-agent systems with self-interested agents is to let manipulation from agents happen and live with it. Preventing strategic behavior of individual agents is unrealistic for many complex systems. Instead, by

studying how the individual, strategic agents impact the overall system behavior, we gain insights on properties of agent societies, such as characterizing the types of environments and agent populations that foster social and antisocial behavior [25, 26]. Our approach falls into this category, as one of our goals is to explain how strategic agents (in particular the p-strategy sellers) affect the UMDL in terms of market and allocation efficiency.

Studies on system-level analyses of the multiple agents, investigating the roles of system-wide properties are found in other papers [9, 10, 12, 13, 26]. Some results show that in large markets with significant uncertainty, taking ones' effect on prices into account does not pay off. This is, of course, a positive indication, since it implies that agents in large markets are not encouraged to speculate about their effect on prices. At the same time, we can expect many different, complex interactions among agents that are hard to predict, as shown in this paper.

7. Conclusion and future work

In this paper, we have used the p-strategy to examine the collective behavior of strategic agents in the UMDL system. We have first shown the performance of the p-strategy seller against other types of sellers. Having established that the p-strategy works better, we have then examined the market and allocation efficiency with varying numbers of p-strategy agents.

The findings are useful to both system designers and agent designers. It is reassuring from the system designers' viewpoints that the market efficiency of the UMDL does not decrease sharply and that the worst-case market inefficiency is not likely to be realized (since even though self-interested agents have the capability of complex strategic reasoning, not all of them will behave so). At present, we cannot determine the exact demographics of the agent population for the best market efficiency, but we are continuing experiments on many different types of agent populations to get a better understanding of the overall system behavior.

From the allocation-efficiency perspective, on the other hand, agent designers learn that using the p-strategy does not always pay off, and that a simple strategy is sometimes as effective. We are currently developing an adaptive p-strategy to dynamically determine when to use the p-strategy and when not to. An adaptive p-strategy will be beneficial not only to a self-interested agent but also to the overall system efficiency.

Acknowledgments

This research has been funded in part by the joint NSF/DARPA/NASA Digital Libraries Initiative under CERA IRI-9411287 and by NSF grant 9872057. The first author has been partially supported by the Horace H. Rackham Barbour scholarship, and the ITECC Center at Rutgers University.

Notes

1. This is an extended version of our ICMAS-98 paper [16].
2. Vickrey's auction mechanism may be inappropriate for certain settings. For limitations of Vickrey auctions, see the work by Sandholm [19].
3. The clearing price can be set in many different ways. For example, it can be set at the middle of the buyer's and the seller's offer prices. By setting the clearing price at the seller's offer price, we deliberately make the sellers to be more strategic. See Section 3 for more details.
4. ZI stands for Zero-Intelligence. The ZI-seller is a "budget-constrained zero-intelligence trader," who generates random bids subject to a no-loss constraint [8].
5. The clearing price is the price quoted by the auction agent. The auction agent hypothetically answers the question of "what price would a seller need to offer to get matched, given the state of the other agents' bids at present?"
6. $\text{Arrival-rate-of-buy-offers} \cong (\text{offer-rate}/\text{offer-interval}) \times \text{clearing-interval}$. The arrival rate varies, however, since the agent is allowed to submit a new offer immediately after a match, even when it has not reached the next offer interval.
7. When the cost is higher than the clearing-price quote, the CP-seller bids its cost plus some fixed markup. That is, it behaves like the FM-seller when its cost is higher than the clearing-price quote.
8. The FM-seller's profit exceeds that of the p-seller in Session 5 (in Figure 9), but at present we cannot conclude whether this is statistically significant. We are statistically analyzing the experimental results with many different auction settings (by varying the number of buyers and sellers, arrival rates, and negotiation zone).
9. The worst-case inefficiency happens when all the sellers use the p-strategy, as shown at Session 6 in Figure 12.
10. Lately the whole field of DAI has been called MAS, while CDPS is often viewed as a subfield. Our definition of MAS captures the self-interestedness of agents, in contrast to cooperative agents in CDPS. Readers may refer to the paper by Durfee and Rosenschein for more discussion about CDPS and MAS [6].
11. We note, however, the agents in cooperative setting may end up competing for resources as well. For instance, if agents do not know how their actions will affect the global state or have only partial (incorrect) information, they may end up working at cross purposes based on their incomplete view of the overall problem.
12. Protocol and strategy correspond to mechanism and individual agent strategy in our terminology, respectively.

References

1. R. I. Brafman and Tennenholtz, "On partially controlled multi-agent systems," *Journal of Artificial Intelligence Research*, vol. 4, pp. 477–507, 1996.
2. A. Brandenburger and E. Dekel, "Hierarchies of beliefs and common knowledge," *Journal of Economic Theory*, vol. 59, pp. 189–198, 1993.
3. CERC, *NSF Workshop: Research Priorities in Electronic Commerce*, Technical Report, Center for Research in Electronic Commerce (CREC), University of Texas at Austin, Austin, TX, 1998.
4. E. H. Durfee, D. L. Kiskis, and W. P. Birmingham, "The agent architecture of the University of Michigan Digital Library," *IEE British Computer Society Proceedings on Software Engineering (Special Issue on Intelligent Agents)*, 1997, vol. 144(1), pp. 61–71.
5. E. H. Durfee, et al., "The Dynamics of the UMDL Service Market Society," in *ICMAS-98 the Second International Workshop on Cooperative Information Agents (CIA-98)*, 1998.
6. E. H. Durfee, and J. S. Rosenschein, "Distributed problem solving and multi-agent systems: Comparisons and examples," in *DAI Workshop*, 1994.
7. P. J. Gmytrasiewicz, and E. H. Durfee, "A rigorous, operational formalization of recursive modeling, in *First International Conference on Multi-Agent Systems (ICMAS-95)*, 1995, pp. 125–132.

8. D. K. Gode and S. Sunder, "Lower bounds for efficiency of surplus extraction in double auctions," in *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust (Eds.), Addison-Wesley: Reading, MA, 1993, pp. 199–219.
9. J. E. Hanson, and J. O. Kephart, "Spontaneous specialization in a free-market economy of agents," in *Autonomous Agents '98 Workshop, Artificial Societies and Computational Markets*, Minneapolis/St. Paul, 1998, pp. 19–26.
10. J. Hu, and M. P. Wellman, "Online learning about other agents in a dynamic multiagent system, in *Second International Conference on Autonomous Agents*, 1998, 239–246.
11. J. H. Kagel and W. Vogt, "Buyer's bid double auctions: Preliminary experimental results," in *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust (Eds.), Addison-Wesley: Reading, MA, 1993, pp. 285–305.
12. J. O. Kephart, et al., "Dynamics of an information-filtering economy," in *Second International Workshop on Cooperative Information Agents*, Paris, 1998.
13. J. O. Kephart, J. E. Hanson, and J. Sairamesh, "Price-war dynamics in a free-market economy of software agents," in *ALIFE VI*, Los Angeles, 1998.
14. D. M. Kreps, *Game Theory and Economic Modelling*, Oxford University Press: Oxford, New York, 1990.
15. S. Park, E. H. Durfee, and W. P. Birmingham, "Advantages of strategic thinking in multiagent contracts (a mechanism and analysis)," in *International Conference on Multi-Agent Systems (ICMAS)*, Kyoto, Japan, 1996, pp. 259–266.
16. S. Park, E. H. Durfee, and W. P. Birmingham, "Emergent properties of a market-based digital library with strategic agents, in *International Conference on Multi-Agent Systems (ICMAS)*, Paris, France, 1998, pp. 230–237.
17. J. S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, The MIT Press, 1994.
18. J. Rust, J. Miller, and R. Palmer, Behavior of trading automata in a computerized double auction market, in *The Double Auction Market*, D. Friedman and J. Rust, (Eds.), 1993, pp. 155–198.
19. T. W. Sandholm, "Limitations of the Vickrey auction in computational multiagent systems," in *International Conference on Multi-Agent Systems*, 1996, pp. 299–306.
20. T. W. Sandholm and V. R. Lesser, "Advantages of a leveled commitment contracting protocol," in *AAAI*, Portland, OR, 1995, pp. 126–133.
21. Y. Shoham and M. Tennenholtz, "Social laws for artificial agent societies: Off-line design, *Artificial Intelligence*, vol. 73(1–2), pp. 231–252, 1995.
22. J. Tirole, *The Theory of Industrial Organization*, The MIT Press: Cambridge, MA, 1998.
23. H. R. Varian, "Economic mechanism design for computerized agents," in *USENIX Workshop on Electronic Commerce*, New York, NY, 1995, pp. 13–21.
24. W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, vol. 16, pp. 8–37, 1961.
25. J. M. Vidal and E. H. Durfee, "The impact of nested agent models in an information economy," in *Second International Conference on Multi-Agent Systems (ICMAS96)*, 1996, pp. 377–384.
26. M. P. Wellman and J. Hu, "Conjectural equilibrium in multiagent learning," *Machine Learning*, vol. 33, pp. 179–200, 1998.
27. G. Zlotkin and J. S. Rosenschein, "Mechanism design for automated negotiation and its application to task oriented domains," *Artificial Intelligence*, vol. 86(2), pp. 195–244, 1996.