An Evolutionary Approach for Studying Heterogeneous Strategies in Electronic Markets

Alexander Babanov
Dept of Computer Science and Engineering
Dept of Economics
University of Minnesota
babanov@cs.umn.edu

Wolfgang Ketter
Dept of Computer Science and Engineering
University of Minnesota
ketter@cs.umn.edu

Maria Gini
Dept of Computer Science and Engineering
University of Minnesota
gini@cs.umn.edu

ABSTRACT

We propose an evolutionary approach for studying strategic agents that interact in electronic marketplaces. We describe how this approach can be used when agents' strategies are based on different methodologies employing incompatible rules for collecting information and for reproduction. We present experimental results from a simulated market, where multiple service providers compete for customers using different deployment and pricing schemes. The results show that heterogeneous strategies evolve in the same market and provide useful research data.

Keywords

multi-agent systems, genetic algorithms, simulation.

1. INTRODUCTION

Online marketplaces are gaining popularity among producers seeking to streamline their supply chains [8], and among consumers looking for good opportunities. Intelligent software agents can significantly facilitate human decision processes either by helping to select strategies to increase profit or by making autonomous choices. In either case, agents need methods for making decisions under uncertain and changing market conditions (see [14] for an analysis of pricing strategies).

The approach to the study of agent decision making that we propose is based on a large-scale evolutionary simulation environment [4]. The goal of the simulation is to reveal a set of parameters and to find a corresponding niche in the market where a strategy has a competitive advantage over other strategies.

The rationale behind our choice of an evolutionary framework is that it provides results without requiring an overly complex theory of agent motivation, optimization criteria, or strategic interaction. The framework is determined by the motives of individual agents, the rules of agents interactions, and the laws governing survival and creation of new agents. Given that, the evolution of the system provides dynamic information on the macroscopic behaviors of the agents society. For an explanation of the relation between micromotives of agents and macrobehavior see [25].

Evolutionary frameworks have been used extensively in Economics [20, 11, 24, 27]. One reason is that economists have long recognized the usefulness of computational systems for rigorous studies through controlled experimentation (see, for example, [17]). Evolutionary systems are relatively straightforward to construct, and, at the same time, they provide a viable tool for experimentation. Using an evolutionary approach allows one to analyze how different strategies change as the result of interactions among many agents over long periods of time.

We start by considering in Section 2 the issues and standard methods for studying the dynamics of interaction and strategies in electronic marketplaces, and we outline our proposed methodology. In Section 3 we present a case study of a simulated market, where multiple service providers compete for customers, and where profitability is the criterion used to stay in business. The experimental results we show conform to expectations. Service providers with different strategies and capacities end up occupying different niches in the market. Finally, in Section 4 we compare our approach with other related methods.

2. AN EVOLUTIONARY FRAMEWORK

The goal of our research is to usefully employ intelligent agents to support human decision making in real markets. A theoretical analysis of sample strategies used by agents could help us understand what benefits could be derived from their use and what strategies would be most appropriate.

To enable a theoretical analysis, an assumption commonly made in multi-agent systems is that an individual agent does not have a large impact on the market. Under such assumption, the analysis can be conducted as if no other agent changes its strategy. Another common approach is to assume that everyone adopts the same strategy, and use the representative agent approach to solve for an equilibrium.

These approaches become problematic when we want to analyze a market in which many different strategies are present, and none of them clearly dominates the others. In this case, the relative strength of each strategy depends greatly on the other strategies that are present in the market as well as on the general market situation. This suggests the need for studying agent strategies in a dynamic market environment.

There is an interplay between theory and computational experiments. A good example to illustrate the issue is the case reported in [24] about his study of the labor market clearinghouse for American physicians. As reported in the study, none of the theorems from the theoretical model could be applied directly to the real market data, since the models where much simpler than reality. However, the theory proved useful to design more complex experiments and to validate the experimental results. The computational methods allowed to compare the experimental results with the theory. Showing that departures from the theory were small was an important result of the entire study.

We propose to use an evolutionary approach as a computational framework for assessing multi-agent systems. The major advantage of an evolutionary framework is that it is a natural choice for studying complex society of entities, where the structure of the society changes during the simulation, and the entities change as well in an effort to adapt to the changing environment. Evolutionary game theory [29] provides tools for analysis in dynamic environments. Examples studied by other researchers range from the emergence of cooperation in an otherwise selfish society [2, 3] with possible formation of spatial patterns of strategic interaction [16], to ostracism and neighborhood effects [13], and design of auction mechanisms [23].

A drawback of most evolutionary systems is that they require a homogeneous representation of the strategies used by the agents, because the major mechanism to maintain diversity is crossbreading [18].

The homogeneity requirement is dictated by the agent reproduction rules, and causes most evolutionary environments to be in-house projects that do not benefit from the cooperation of several research teams. In our approach, we extend the way agents are reproduced in a way that allows the use of strategies that are represented heterogeneously. We illustrate our approach with a case study of a society of service providers and customers.

Our method adds a layer of evolutionary learning atop disjointly evolving types of agents that use different strategies. Every type of agent maintains its own separate source of "genetic" information. This information could be a gene pool, if the type is based on the genetic algorithms paradigm; it could be some statistical data, as it is the case in the test model we introduce later in this paper; or it could be a neural network, which is continuously trained on the performance of its "children."

Agents who fail to satisfy a predefined performance criterion are removed at regular time intervals and, eventually, replaced by more fit entities. The purpose of the additional evolutionary layer is to learn the probabilities with which new agents of each type should should enter the market. By observing the distribution of the surviving agents by type, the system can assess which types of agents are more successful and give them additional market space by creating new agents of their same type. Once the type of each new agent has been decided, the corresponding reproduction rule is applied.

We illustrate this two-layered evolutionary framework in greater detail in the next Section where we introduce our case study.

3. TEST MODEL

In this Section we present a case study of how the suggested approach is used to assess two different pricing and deployment strategies in a market with multiple suppliers and their customers.

Our test model is a continuous time discrete-event simulation of a society of economic agents: suppliers of a service and their customers. The assumptions of the model are similar to those from the urban economics model of a monocentric city [21], and from the model of market areas with free entry [19]. We intentionally build our model upon classical urban economics models to ensure its practical meaning. It is important to note though that the observations of the model's behavior are not paramount, they are offered mainly to demonstrate that the results obtained by the proposed approach are, in fact, meaningful.

3.1 General Terms

The agents live and interact in a circular city of radius R. Customers come to the market for a single transaction at random intervals governed by a stationary Poisson process with a fixed frequency λ^{c}:

$$t_{i+1}^{\mathrm{c}} = t_i^{\mathrm{c}} - \frac{1}{\lambda^{\mathrm{c}}} \log U[0,1]$$

where $U[x,y]$ is a random variable distributed uniformly on the interval $[x,y]$. The location of a new customer in polar coordinates is determined by the following rules:

$$r \sim U[0,R] \qquad \text{and} \qquad \alpha \sim U[0,2\pi)$$

The rules imply that the density of customers is inversely proportional to the distance from the center of the city[1].

Upon entry, a customer observes different suppliers and chooses the one that provides the service at minimum cost. We define the cost of service c as a function of the supplier's price p, distance to the customer d, and delay due to servicing previously scheduled customers Δt:

$$c = p + d \times c^{\mathrm{mile}} + \Delta t \times c^{\mathrm{hour}}$$

where c^{mile} and c^{hour} are cost per mile of travel and cost per hour delay respectively.

The customer side of the society is assumed to be in equilibrium and does not change its properties in the course of a simulation. At the same time the society of suppliers is expected to evolve and meet the demands of the customers. The restriction on the customer side is imposed to fix the scale of the simulation as well as to avoid imposing extra assumptions on the behavior of customers. In fact, it is customary for urban economics models to assume a fixed distribution of the population density and a particular form of the customers' utility function. In our model though we have the freedom of changing the parameters of the customer side during the simulation.

Suppliers, in turn, enter the market with a frequency that is positively correlated to the average profit derived in the market. A supplier is characterized by its pricing strategy and the number of customers it can serve simultaneously, also called the supplier's *size* for brevity. Serving one customer takes size s supplier one continuous hour and costs $c^{\mathrm{work}}(s)$, staying idle for an interval of time of any length costs $c^{\mathrm{idle}}(s)$ per hour. Both costs decrease with size to simulate economies of scale.

Each supplier is audited at regular periods and removed from the market if its profit becomes negative.

[1] Studies in urban economics suggest and support by empirical evidence the use of a reversed exponential relation between population density and distance from the city center (see, for example, [1, 21]). We adopt a hyperbolic distance-density relation for convenience of the analysis and subsequent evolutionary experiments.

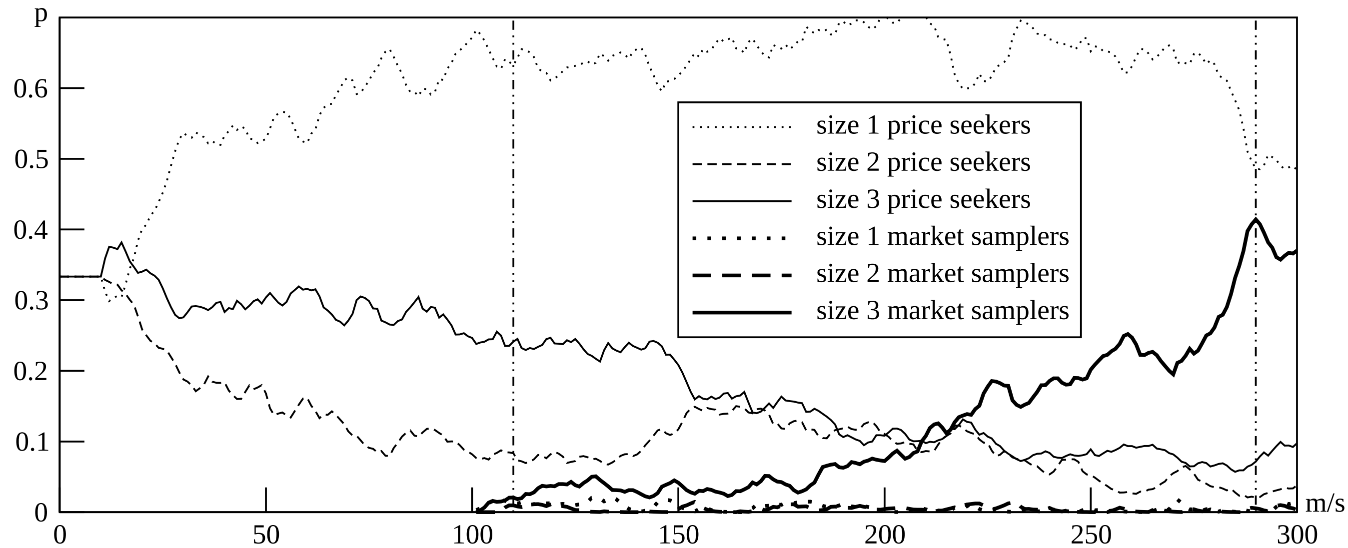

Figure 1: Probabilities of a new supplier entry for different supplier types as a function of milestone numbers. Market sampler suppliers are introduced at milestone 100. Vertical lines denote milestones 110 and 290.

3.2 Supplier Strategies and Generators

We define a *supplier type* to be a pair of supplier's size and pricing strategy. A type is represented in the market by the corresponding *supplier generator*. Each generator maintains a pool of information concerning the history and the current state of its type suppliers. It uses the collected information to create new suppliers of its type and provide them with initial location, price and, perhaps, other parameters specific to the type.

The probability that a supplier of a particular type will enter the market next is proportional to the number of suppliers of its type that are surviving in the market. There is also a small probability, referred to as *noise*, that a new supplier is assigned a type at random. The noise allows types that were once retired from the market completely to enter it at a more favorable time. It also suggests a way for new types to enter the market, as it will be shown later in the experimental results.

In the experiments we describe later, we have used two strategies that exhibit sufficiently different behavior. We designed the strategies so that neither strategy has a strict advantage over the other. Because of that, the strategies can coexist and evolve in the market at the same time. In both strategies, the supplier accepts whatever location and price was suggested by its generator and never alters them. Such restriction simplifies the analysis of the results, since only generators are capable of learning and adapting to the market situation.

The first strategy, code named *market sampler*, involves sampling the city in several locations to maximize a potential revenue flow given the state of the market[2]. The price and the number of samples it takes to place a supplier are assumed to be distributed normally. The corresponding generator periodically collects information on the price and the number of samples from the current suppliers that passed at least one audit. The information is used to estimate the normal distributions mentioned above and merge them with the historical ones.

The second strategy, *price seeker*, assumes that the "right" price and density of the suppliers depend solely on the distance from the center of a city. The price seeker generator attempts to estimate the most beneficial distributions of price and density by observing the number and prices of the surviving suppliers. Consequently, the location of a new supplier is chosen to reflect the beliefs of its generator about the density, and the price is selected depending on the distance from the center.

To summarize the important properties of the strategies, the market sampler strategy is capable of pinpointing the best location to deploy a new supplier, yet it assumes the same price distribution for every location. The price seeker strategy can avoid unattractive locations and choose better prices on a global scale, while lacking the precision in positioning its suppliers relative to their rivals.

3.3 Expected Properties of the Model

There are a few properties one might expect to observe in a reasonably behaving simulation. First, large size suppliers shall eventually concentrate in the center of the city where their production cost advantage will attract customers and where the customer demand is sufficiently high to give enough work to larger suppliers. In the outer parts of the city, the scarcity of customers shall result in a low viable density of suppliers and, hence, in a relatively high influence of transportation costs on the customer preferences. That should make small suppliers prefer the rim, where they could stay closer to the customers than their larger counterparts.

The second expectation is that in a market dominated by market samplers the distribution of prices as a function of distance from the center should be flatter than in a market dominated by price seekers. The reason is that market samplers assume the same distribution of prices suits all city locations, while price seekers can crowd in zones in which their size and price selections are advantageous. The same

[2]Under a set of assumptions it amounts to finding a maximum of $(p_c + d_c \times c^{\text{mile}})^3$, where p_c and d_c are the price and the distance to the cheapest supplier at the location.

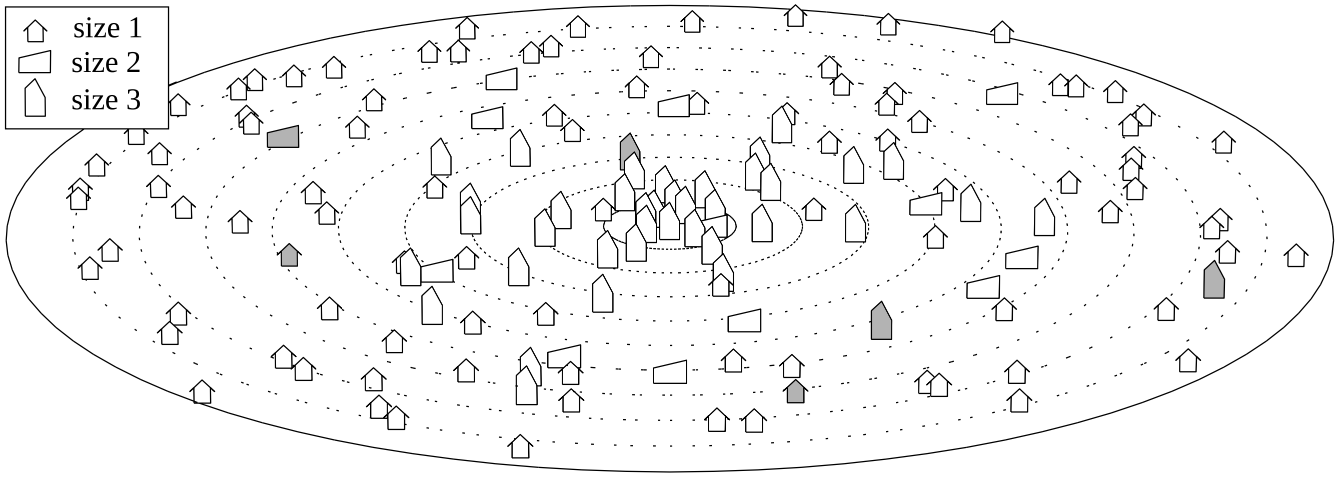

Figure 2: City snapshot at milestone 110. Price seeker suppliers are white, market samplers are gray.

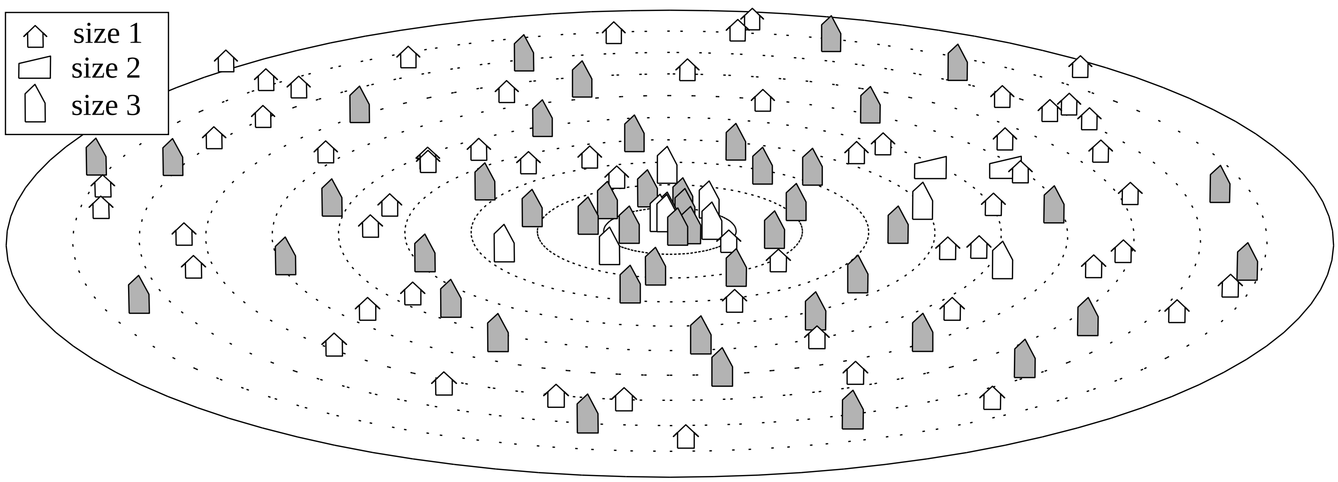

Figure 3: City snapshot at milestone 290. Price seeker suppliers are white, market samplers are gray.

argument implies that the concentration of small suppliers in the outer zones and large suppliers in the central ones shall be expected to be more distinct in the society of price seekers.

The final expectation is related to the introduction in the market of new and retired types. In a properly functioning model it should be possible for new types to acquire a niche in an existing market. This does not imply that every strategy must win a noticeable representation in the market, only that a reasonably competitive strategy should eventually get its share when the market state is favorable.

3.4 Experimental Results

In order to verify the viability of our ideas we have conducted a large number of experiments using different market parameters, such as frequencies and costs defined previously. The following analysis applies to a typical experiment[3], which starts with a city populated by price seeker suppliers of sizes 1, 2 and 3. Later, market sampler suppliers of the same three sizes are introduced to the market via a 10% noise factor (i.e. initially each of the new types gets less than 2% chance every time a supplier is created).

The resulting distribution of entry probabilities for each of the six types is shown in Figure 1. Each *milestone* (m/s) in the figure corresponds to 10,000 simulated hours and roughly 2.5 million of transactions[4]. The market samplers enter the market at milestone 100 when the situation is relatively stable and try to find their niche. Eventually, the size 3 market sampler type proves itself to be competitive and captures a sizable share. In the following we consider two simulation milestones: one at 110, soon after the market samplers' entry, and the other at 290, at the point of the major success of size 3 market sampler types. These two milestones are depicted in Figure 1 by vertical lines.

Figure 2 suggests a schematic view of the city at milestone

[3]For the previously introduced parameters we used the following values: $R = 50$ miles, $\lambda^c = 250$ customers per hour, $c^{\text{mile}} = 0.5$ \$/mile and $c^{\text{hour}} = 1$ \$/hour. Costs c^{idle} and c^{work} for size 1 suppliers were set to 5 and 10 \$/hour respectively and reduced by 3% with each size unit gain.

[4]The probabilities are kept constant for the first 10 milestones to let the generators adjust their (or rather our) pretty volatile initial guesses about the market situation.

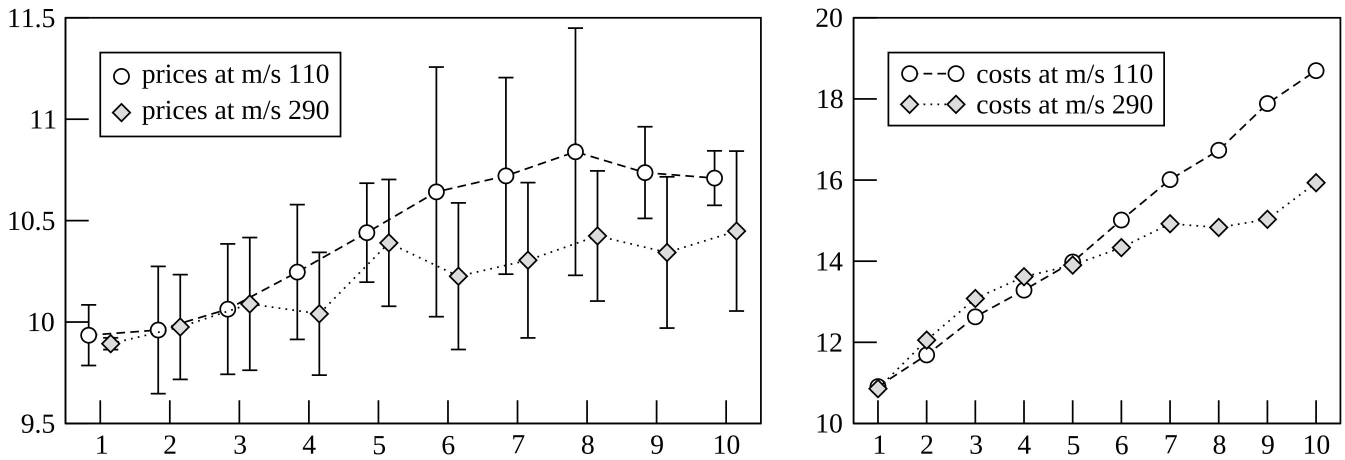

Figure 4: Average supplier prices with standard deviations (left) and 25 hour half-life decaying averages of customer costs (right) for 10 concentric city zones at milestones 110 and 290.

110. Each supplier is depicted by a house of a particular shape, depending on the supplier's type; price seeker types are colored in white and market samplers are gray. The concentric circles divide the city in ten equally wide zones with each zone getting roughly the same number of customer entries per unit of time.

There are two important observations to be made from Figure 2. Firstly, price seekers dominate the market and, indeed, their size 1 suppliers tend to congregate in the rim, while size 3 operate mostly in the middle of the city and size 2 form a ring in between the other two types. Secondly, the distribution of suppliers is quite uneven with dense clusters and wide open areas situated at the same distance from the center. The summary of the observations confirms that the price seeker generators have little regard to the exact placement of suppliers, while converging to the right distributions as a whole.

Figure 3 presents a view of the city at milestone 290 when the market is dominated by size 3 market samplers and size 1 price seekers. The structure of the market is evidently different from the one we witnessed 10 milestones after the market samplers were introduced — the size 3 market samplers are distributed regularly across the city with the size 1 price seekers surviving in between and in several remaining clusters.

To complete the comparison of the market situations we introduce Figure 4 with information on the prices and the costs that customers face at the two considered milestones. The left graph shows average prices and standard deviations for the ten concentric city zones. The right graph shows the *decaying averages*. The concept of the decaying average is introduced to avoid storage requirements of calculating moving averages in a large scale discrete event simulation. We define the decaying average da_i with a half-life T for the time interval $[t_i, t_{i+1})$ iteratively as

$$w_i = 1 + w_{i-1} 2^{-\frac{t_i - t_{i-1}}{T}} \quad \text{and} \quad da_i = \frac{v_i + (w_i - 1) da_{i-1}}{w_i}$$

where t_i and v_i are respectively the value and the time of the event i, and w_i is the weight of all events from 1 to i at time t_i. for the customer costs in the same zones. Both graphs suggest that the market populated with market sampler types exhibits more even price and, subsequently, customer cost distributions.

It should be emphasized here that the most important conclusion of our experiments is that a society of agents that use strategies based on different approaches, information pools, and reproduction methods does evolve as one society and produces meaningful results.

4. RELATED WORK

One of the core concepts in the evolutionary theory is the concept of an Evolutionary Stable Strategy [28]. A strategy like this has two main characteristics: (i) it can survive against itself, and (ii) no mutation adopted by an arbitrarily small fraction of individuals can invade by getting at least a comparable payoff. In our system the first requirement is satisfied, the price seeker strategy survived against itself. The second characteristic is not true in our case, the market sampler strategy invaded and was able to obtain a large presence in the market, yet it was unable to take over the market completely.

In the real world it is common for heterogeneous strategies to coexist. On the contrary, in a simulation environment, it is difficult to come up with a single encoding for strategies, even for well studied problems [12]. Our approach allows for the interaction and evolution of heterogeneous strategies.

Much research has been done in the last few years in designing pricing strategies for agents, assessing their performance, and seeing how well they adapt to changing environmental situations [5].

Understanding collective interactions among agents that dynamically price services or goods is discussed in [14], where several pricing strategies are compared. The major difference between their pricing strategies and the strategies our agents use is that transportation costs and geographical locations are important for us, but are not relevant when buying and selling information. Examples of price-wars caused by agents that dynamically change their posted price for information bundles are described. Because of the complexity of the problem, experiments are limited to a small number of agents.

A simulation based approach to study dynamic pricing strategies in finite time horizon markets is described in [10].

The study uses a market simulator and simple strategies. The results are evaluated in terms of overall profit, but there are so many variables in the simulation that it is hard to assess the generality of the results obtained. One of the advantages of using an evolutionary framework is that it simplifies experimenting with a variety of parameters.

An important conclusion of many studies is that even simple strategies can be very effective. For instance, Cliff's [6] Zero-Intelligence Plus trader agents have minimal intelligence, yet they have been successfully used in continuous double auctions, where they performed very well even when compared to human traders [9].

The use of evolutionary methods is proposed in [22], who simulates the evolution of the agent population as they adapt their strategy for continuous double auctions by observing what happens in the environment. Cliff [7] uses genetic algorithms to learn the parameters that control how his trader agents evolve their pricing strategies. Along similar lines, an evolutionary system based on Genetic Programming is presented in [23]. In this system, agents evolve auction strategies to bid in the electricity market.

The major difference between these and the work presented here, is that we are interested in understanding how strategies of individual agents interact in the market, as opposed to study specific types of auctions to learn auction rules. We are also interested in providing a methodology for studying effectively multi-agent systems with a large number of agents.

Studying such systems analytically is often impossible. There are a few attempts to model very large multi-agent systems at the macroscopic level. Shehory [26] models them using ideas from classical mechanics. Goal satisfaction is modeled by collisions between dynamic particles, the agents, and static particles, the goals. The method requires a measure of distance to the goal, which is hard to do except in cases where agents operate in a cartesian environment. The methodology presented in [15] is limited to systems that obey the Markov property, i.e. such that the agent's future state depends only on its present state.

5. CONCLUSIONS

We have proposed an evolutionary framework where agents with different pricing and location selection strategies compete in a market. New agents are introduced to the market with a probability proportional to the number of agents with the same strategy already in the market. In this way, successful agents have a greater probability that new agents with their same strategy will be introduced in the market, and so have a better chance of creating a market niche.

The evolutionary system produces several coexisting *optimal* strategies, not only a single *optimum* one. In our test case, optimal strategies vary according to the geographical location and size of suppliers, and depend on the state of neighboring suppliers, and the general market situation.

In the experiments, we initially allowed suppliers of different sizes and with different strategies to select their own location. During the simulation we observed how different types of suppliers occupied different market niches and how new behavior patterns emerged over time.

Current work is aimed at developing a more detailed theoretical model of the agents, and analyzing how close the experimental results obtained with the evolutionary framework are to the results predicted by the model.

6. ACKNOWLEDGMENTS

Partial support for this research is gratefully acknowledged from the National Science Foundation under awards NSF/IIS-0084202 and NSF/EIA-9986042.

7. REFERENCES

[1] Alex Anas, Richard Arnott, and Kenneth A. Small. Urban spatial structure. *Journal of Economic Literature*, 36(3):1426–1464, September 1998.

[2] R. M. Axelrod. *The evolution of cooperation.* Basic Books, 1984.

[3] Robert Axelrod. *The complexity of cooperation.* Princeton University Press, 1997.

[4] A. Babanov, W. Ketter, and M. Gini. An evolutionary framework for large-scale experimentation in multi-agent systems. In T. Wagner, editor, *MAS Problem Spaces and Their Implications to Achieving Globally Coherent Behavior.* Kluwer, 2003.

[5] Christopher H. Brooks, Robert Gazzale, Rajarshi Das, Jeffrey O. Kephart, Jeffrey K. MacKie-Mason, and Edmund H. Durfee. Model selection in an information economy: Choosing what to learn. *Computational Intelligence*, April 2002.

[6] D. Cliff and J. Bruten. Minimalintelligence agents for bargaining behaviors in marketbased environments. Technical Report HPL-97-91, Hewlett Packard Labs, 1997.

[7] Dave Cliff. Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets. Technical Report HPL-2001-99, Hewlett Packard Labs, 2001.

[8] John Collins, Wolfgang Ketter, and Maria Gini. A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. *Int'l Journal of Electronic Commerce*, 7(1):35–57, 2002.

[9] Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauro. Agent-human interactions in the continuous double auction. In *Proc. of the 17th Joint Conf. on Artificial Intelligence*, Seattle, WA, USA, August 2001.

[10] Joan Morris DiMicco, Amy Greenwald, and Pattie Maes. Dynamic pricing strategies under a finite time horizon. In *Proc. of ACM Conf on Electronic Commerce (EC'01)*, October 2001.

[11] Joshua M. Epstein and Robert Axtell. *Growing Artificial Societies: Social Science from the Bottom Up.* The MIT Press, Cambridge, MA, 1996.

[12] Stephanie Forrest. Genetic algorithms: Principles of natural selection applied to computation. *Science*, 261:872–878, 1993.

[13] Richard J. Gaylord and Louis J. D'Andrea. *Simulating Society.* Springler-Verlag, 1998.

[14] Jeffrey O. Kephart, James E. Hanson, and Amy R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, 2000.

[15] Kristina Lerman. Design and mathematical analysis of agent-based systems. In *Lecture Notes in Artificial Intelligence (LNAI 1871)*, pages 222–233. Springer Verlag, 2001.

[16] Kristian Lindgren. Evolutionary dynamics in game-theoretic models. In *The Economy as an*

Evolving Complex System II, pages 337–367, 1997.

[17] D. McFadzean, D. Stewart, and Leigh Tesfatsion. A computational laboratory for evolutionary trade networks. *IEEE Transactions on Evolutionary Computation*, 5(5):546–560, October 2001.

[18] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs.* Springer, 1996.

[19] Edwin S. Mills and Michael R. Lav. A model of market areas with free entry. *Journal of Political Economy*, 72(3):278–288, 1964.

[20] Richard R. Nelson. Recent evolutionary theorizing about economic change. *Journal of Economic Literature*, 33(1):48–90, March 1995.

[21] Arthur O'Sullivan. *Urban Economics.* McGraw-Hill Higher Education, 1999.

[22] Sunju Park, Edmund H. Durfee, and William P. Birmingham. An adaptive agent bidding strategy based on stochastic modeling. In *Proc. of the Third Int'l Conf. on Autonomous Agents*, 1999.

[23] Steve Phelps, Peter McBurney, Simon Parsons, and Elizabeth Sklar. Co-evolutionary mechanism design: a preliminary report. In J. Padget, Onn Shehory, David Parkes, Norman Sadeh, and William Walsh, editors, *Agent Mediated Electronic Commerce IV*, volume LNAI2531, pages 123–142. Springer-Verlag, 2002.

[24] Alvin E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, July 2002.

[25] Thomas C. Schelling. *Micromotives and Macrobehavior.* W. W. Norton & Company, Inc., 1978.

[26] Onn Shehory, Sarit Kraus, and O. Yadgar. Emergent cooperative goal-satisfaction in large scale automated-agent systems. *Artificial Intelligence*, 110(1), May 1999.

[27] Leigh Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.

[28] Fernando Vega-Redondo. *Evolution, Games and Economic Behavior.* Oxford University Press, New York, 1996.

[29] Jörgen W. Weibull. *Evolutionary Game Theory.* The MIT Press, 1995.