An Evolutionary Framework for Determining Heterogeneous Strategies in Multi-Agent Marketplaces

Alexander Babanov

Dept of Computer Science and Engineering
University of Minnesota, USA

Wolfgang Ketter

Department of Decision and Information Sciences
RSM Erasmus University, NL

Maria Gini

Dept of Computer Science and Engineering
University of Minnesota, USA

Abstract

We propose an evolutionary approach for studying the dynamics of interaction of strategic agents that interact in a marketplace. The goal is to learn which agent strategies are most suited by observing the distribution of the agents that survive in the market over extended periods of time. 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 and co-exist in the same market.

Key words: Trading agents, multi-agent systems, genetic algorithms, simulation

1 Introduction

The approach to the study of agent decision making that we propose in this paper is based on a large-scale evolutionary simulation environment [2]. The

Email addresses: `babanov@cs.umn.edu` (Alexander Babanov), `wketter@rsm.nl` (Wolfgang Ketter), `gini@cs.umn.edu` (Maria Gini).

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 [17].

Evolutionary frameworks have been used extensively in Economics [12,6,16,19]. One reason is that economists have long recognized the usefulness of computational systems for rigorous studies through controlled experimentation (see, for example, [10]). 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.

A drawback of most evolutionary systems is that they need a homogeneous representation of the strategies used by the agents. The homogeneity requirement is dictated by the agent reproduction rules. We propose an extension that allows one to use strategies that are represented heterogeneously, and we illustrate our ideas with a case study of a society of service providers and customers.

Our proposed extension adds a layer of evolutionary learning atop disjointly evolving types of agents that use different strategies. Every type of agent maintains its 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; it could even be a neural network, which is continuously trained on the performance of its "children."

The purpose of the additional evolutionary layer is to learn the probabilities with which representatives of each type should enter the market, by observing the distribution of surviving agents by type. Each time a new agent is introduced to the market, the top layer decides on its type and uses the corresponding reproduction rule. Agents who fail to satisfy a predefined performance criterion are removed at regular time intervals and, eventually, replaced by more fit entities.

We introduce our case study, a test model called *Citysim*, in Section 3 and we present an analytical model. We illustrate the idea of the two-layered evolutionary framework in greater detail in Section 4. In Section 5 we present experimental results obtained using the evolutionary framework and examine the distribution of different strategies over a wide range of initial conditions. The results show that it is possible for heterogenous strategies to coexist in the market and to evolve to occupy specific niches. For reading convenience, we present a summary of our notation in the Appendix.

2 Related Work

Joshua Epstein and Rob Axtell have developed an evolutionary bottom-up agent-based computational framework, Sugarscape [6], for the study of human social phenomena such as trade, migration, group formation, combat, interaction with an environment, transmission of culture, propagation of disease, and population dynamics to grow an agent economy from scratch. Sugarscape showed successfully that an evolutionary framework without any traditional human made assumptions can grow and deliver meaningful results.

Understanding collective interactions among agents that dynamically price services or goods is discussed in [7], where several pricing strategies are compared. Because of the complexity of providing an analytical treatment, the study is limited to a small number of agents. 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.

A simulation based approach to study dynamic pricing strategies in finite time horizon markets is described in [5]. 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 like we do is that it simplifies experimenting with a variety of parameters.

An important result of these studies is that a Nash equilibrium is unlikely to exist in a market where customers look for the lowest price and suppliers adjust their prices dynamically. Another important result is that even simple strategies can be quite effective. For instance, Cliff's [3] 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 [4].

The use of evolutionary methods is proposed in [14], where the evolution of the agent population as they adapt their strategy by observing what happens in the environment is simulated. Cliff [**?**] 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 [15]. The major difference with the work presented here, is that we aim at providing a methodology for extending theoretical modeling to deal effectively with complex multi-agent systems with a large number of agents. In [**?**], an evolutionary system is used to model changes in the use of land, in particular to model deforestation in Yucatan.

Studying such systems analytically is often impossible. There have been a few attempts to model very large multi-agent systems at the macroscopic level. Shehory [18] 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 [8] is limited to systems that obey the Markov property, i.e. such that the agent's future state depends only on its present state.

Our model has some similarities to the model for electronic service markets presented in [9], in the sense that we also take into account customer waiting time and not only the price of the service. In addition to modeling the waiting times, our model includes a geographical component. Transportation costs affect customer choices and consequently affect the decision of where each supplier locates its business.

3 Concept Check: Citysim

In this Section we present a test model called *Citysim*. We demonstrate that in this model it is possible for heterogenous strategies to coexist in the market and to evolve to occupy specific niches. We examine the emergence of such a behavior over a wide range of initial conditions and support the experimental results with theoretic analysis of the model.

3.1 The Model

Citysim is a model of a city inhabited by two types of economic agents: providers of a particular service, for example, gas stations or hairdressers, and their customers. We assume that the properties of the customer side of the society do not change in the course of a simulation. At the same time the society of suppliers is expected to evolve to meet the demands of the customers.

The assumptions of our test model are similar to those from the urban economics model of a monocentric city [13], and from the model of market areas with free entry [11]. We intentionally built 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 essential, they are offered mainly to demonstrate that the results obtained by the means of the proposed approach are, in fact, meaningful.

One issue with using evolutionary models is that the number of control variables might easily become uncontrollable. Typically, a researcher can guide an evolutionary model by changing: (a) the parameters of the environment, (b) the rules of interaction between agents, (c) the reproduction rules, (d) the initial parameters and distributions, and/or (e) the internal settings of the agents and agents' strategies.

The number of parameters is high and it is a daunting task to study the space of parameters thoroughly (and convincingly). The fact that the evolution of a system might take different paths depending on random events, such as those used in the reproduction scheme or in pairing agents during market activities, adds to the complexity. Worse yet, it is not always clear what parameters control the major trends in the system and what merely cause small perturbations.

In this Section we tackle the problem of managing control variables by backing up the evolutionary system by the corresponding analytical model.

We proceed by building a model of a monocentric city that serves as a market for suppliers of some service. Suppliers are given free entry to the market, and are allowed to position themselves when they enter the market and set prices. Suppliers may enter the market with different *sizes*, meaning that they can serve a different number of customers simultaneously. They use scheduling to provide service to the customers at the earliest possible time. The net price to customers accounts for the supplier's price, transportation cost, and service delay.

Next we apply a market area analysis to find what comparative advantage each size supplier has, and under which circumstances.

3.2 General Terms

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

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

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:

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

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 looks for suppliers and chooses the one that provides the service at the minimum cost. We define the net 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:

$$c = p + d \times c^{\mathrm{mile}} + \Delta t \times c^{\mathrm{hour}} \tag{3}$$

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

Suppliers, in turn, enter the market at intervals governed by a random process similar to the one used for the customers. A supplier is characterized by its pricing strategy and the number of customers it can serve simultaneously, also referred as the *supplier's size*. 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:

$$c^{\mathrm{work}}(s) = c^{\mathrm{work}}(1) \times (1-g)^{s-1} \quad \text{for} \quad s \geq 2 \tag{4}$$

$$c^{\mathrm{idle}}(s) = c^{\mathrm{idle}}(1) \times (1-g)^{s-1} \quad \text{for} \quad s \geq 2 \tag{5}$$

where $c^{\mathrm{work}}(1)$ and $c^{\mathrm{idle}}(1)$ are constants and g determines a gain due to the supplier's size.

[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, [1,13]). We adopt a hyperbolic distance-density relation for convenience of the analysis.

A supplier maintains a schedule of all future deliveries and their prices; it is paid after delivering the service. Each supplier is audited at regular periods and removed from the market if its profit becomes negative.

3.3 Analytical Model

We impose the following set of assumptions on the equilibrium state of the market:

Assumption 1 *Suppliers operate at zero profit without idle periods, and do not discriminate between customers.* This assumption implies that a size s supplier sets its price at a constant level of the work cost $c^{\mathrm{work}}(s)$.

Assumption 2 *The market area of a single supplier is small relative to the size of the whole city,* hence the density of customer arrivals might be considered constant over each supplier's market area.

Assumption 3 *Market areas are circular and do not interfere.*

Assumption 4 *Inside each market area customers arrive at regular intervals.*

Under assumption 1 a supplier of size s should attract a minimum of s customers in its market area per hour. We refer to the smallest circular area that can support a particular supplier as its *support*. The relation between the radius of the support ρ and the frequency of customers' arrival per square mile per hour ν is expressed as

$$s = \pi \rho^2 \nu \tag{6}$$

We can use the equations (2) to find the frequency ν as a function of the distance from the center r. Consider a part of a city ring with middle radius r and width Δ. For small Δ-s a part of such ring with angular size Δ/r is a nearly square region that accounts for a fraction $\frac{\Delta}{R} \times \frac{\Delta}{2\pi r}$ of all customer arrivals. The frequency ν then can be expressed as the limiting ratio of the fraction of customers appearing in the region to the area of this region by a square:

$$\nu(r) = \lim_{\Delta \to 0} \frac{\frac{\Delta}{R} \times \frac{\Delta}{2\pi r}}{\Delta^2} \lambda^{\mathrm{c}} = \frac{\lambda^{\mathrm{c}}}{2\pi r R} \tag{7}$$

From equations 6 and 7 we derive the radius of the support as

$$\rho(s, r) = \sqrt{\frac{2rRs}{\lambda^{\mathrm{c}}}} \tag{8}$$

The next step is to note that under assumptions 1 and 4 in equilibrium the average delay that customers experience inside the support of size s supplier is equal to $1/2s$. Hence the net price that customers face on the boundary of the support is

$$p(s,r) = c^{\text{work}}(1-g)^{s-1} + c^{\text{mile}}\rho(s,r) + c^{\text{hour}}\frac{1}{2s} \tag{9}$$

Now we can equate prices at the boundaries of the supports of supplier sizes s_1 and s_2 to find the distance from the center r at which they don't have any advantage over each other. Assuming $s_1 > s_2$ and solving for r we get

$$\bar{r}(s_1,s_2) = \frac{\lambda^{\text{c}}}{2R}\left[\frac{c^{\text{work}}\left[(1-g)^{s_1-1} - (1-g)^{s_2-1}\right] + c^{\text{hour}}\left[\frac{1}{2s_1} - \frac{1}{2s_2}\right]}{c^{\text{mile}}\left[\sqrt{s_2} - \sqrt{s_1}\right]}\right]^2 \tag{10}$$

Although Eq. 10 allows us to find the boundaries between each pair of different supplier sizes, we also need a criterion that will show the comparative advantage of one strategy over another as a function of the distance from the center. To derive such criterion we equate prices at the boundaries with size s_2's radius "inflated" by a coefficient γ to show the extra (over strictly necessary) area where s_2 is preferred by customers over s_1. Solving for γ we get a coefficient that is equal to 1 at $\bar{r}(s_1,s_2)$, and is increasing with size s_2 gaining a comparative advantage over s_1:

$$\gamma(s_1,s_2,r) = \frac{c^{\text{work}}\left[(1-g)^{s_1-1} - (1-g)^{s_2-1}\right] + c^{\text{hour}}\left[\frac{1}{2s_1} - \frac{1}{2s_2}\right]}{c^{\text{mile}}\sqrt{\frac{2rRs_2}{\lambda^{\text{c}}}}} + \sqrt{\frac{s_1}{s_2}} \tag{11}$$

3.4 Analysis of the Analytical Model

Note that there are four degrees of freedom in both equations 10 and 11, namely: (λ^{c}/R), $(c^{\text{work}}/c^{\text{mile}})$, $(c^{\text{hour}}/c^{\text{mile}})$ and g. The first one determines the size of the simulated city and it is only of interest from the scalability point of view. The other three variables influence the shapes and relative positioning of the city zones where having a particular size of service is relatively advantageous.

Figure 1 illustrates how the size gain parameter g and the cost ratios influence the preferences of suppliers. The figure shows three graphs in which size 1 to 5 has the highest advantage over size 3 as a function of the distance from the center and some other parameters. In each graph we restrict (λ^{c}/R) and two other degrees of freedom such that the horizontal line drawn at the middle of each graph is the common intersection of all three.

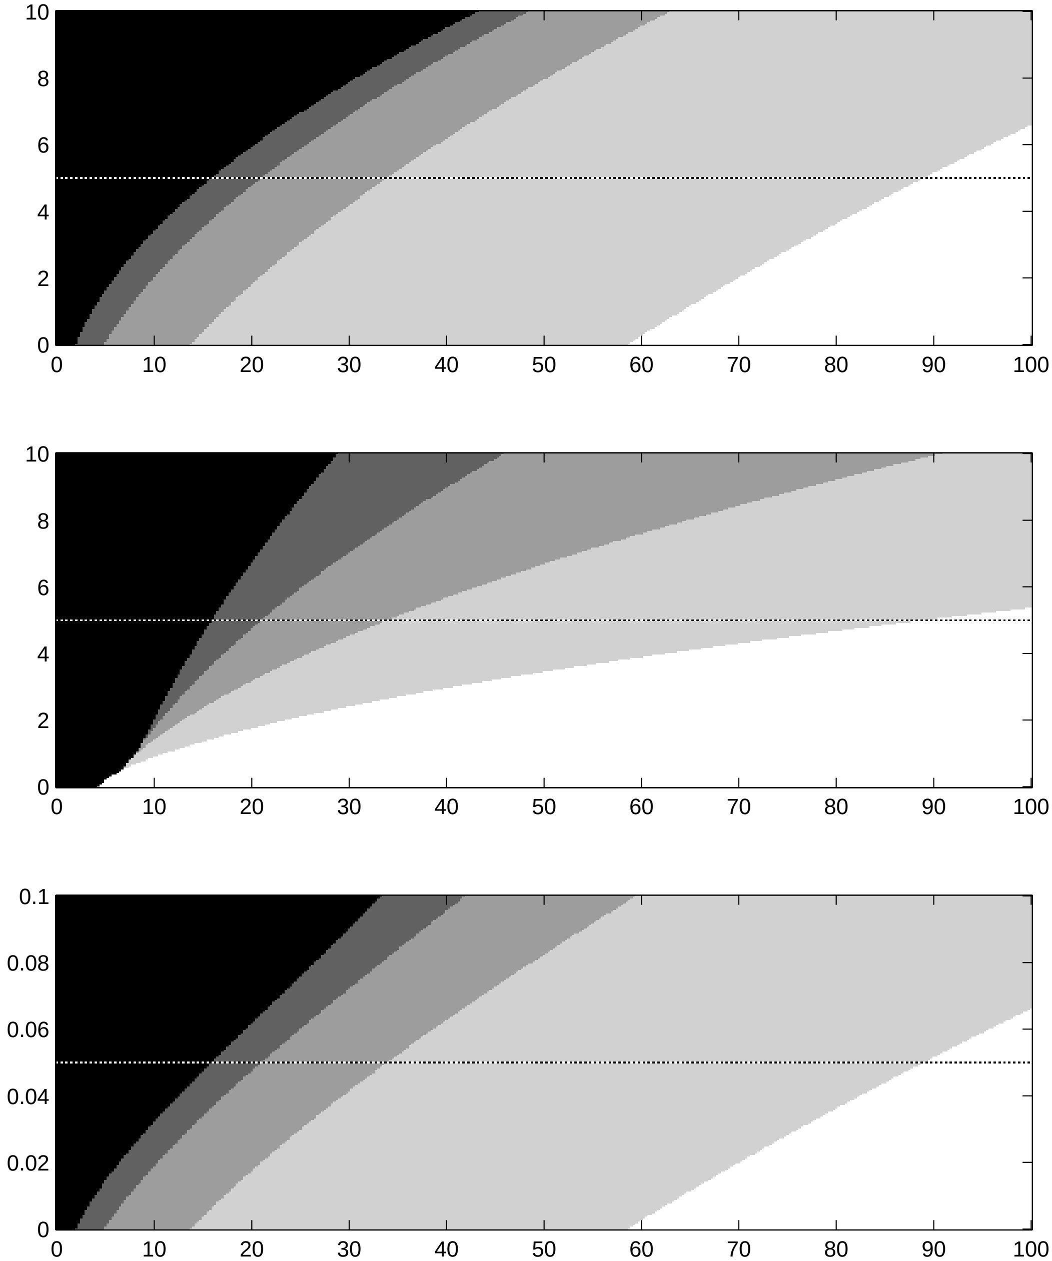

Fig. 1. Graphs displaying which of sizes from 1 (white) to 5 (black) has the highest comparative advantage over size 3 (medium gray). The x-axis is the distance from the center; the y-axis is $(c^{\text{work}}/c^{\text{mile}})$ for the top figure, $(c^{\text{hour}}/c^{\text{mile}})$ for the middle, and g for the bottom one. The horizontal line in the middle of each graph is the common intersection of all.

By studying the graphs in Figure 1 we may choose a set of parameters that is of particular interest for us. For example, the common intersection line seems to have all five sizes occupying noticeable niches in the market. Drawing such conclusion, however, requires that the theory and the experimental results have a reasonable degree of compliance. Therefore the next section is devoted to showing that this, in fact, is the case in the suggested model.

4 Evolutionary Environment

Our test model is a continuous time discrete-event simulation of a society of economic agents: suppliers of a service and their customers. The customer side of the society is assumed to be in equilibrium and does not change its properties in the course of a simulation. The society of suppliers is the one that evolves over time to 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 the 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.

4.1 Supplier Strategies and Generators

We define a pair of supplier's size and pricing strategy to be a *supplier type*. 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 an 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 again, perhaps, at a more favorable time. It also suggests a way for completely new types to enter the market, as we demonstrate later in Section 5.2.

In the experiments referred to in this chapter we have used two strategies that exhibit sufficiently different behaviors. In both strategies a 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 the potential revenue flow given the state of the market.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 on the distance from the center of the 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.

5 Experimental Results

5.1 *“Reality Check” Experiments*

Our first series of experiments with the Citysim model are aimed at verifying if the derivations from the analyitical model are close to the results from actual simulations.

The experiments were conducted for the following values of the previously defined parameters: $R = 100$ miles, $\lambda^{\mathrm{c}} = 1200$ or 1600 customers per hour, $c^{\mathrm{mile}} = 2$ \$/mile, $c^{\mathrm{hour}} = 10$ \$/hour. Costs c^{idle} and c^{work} for size 1 suppliers were set to 5 and 10 \$/hour respectively and reduced by $g = 5\%$ with each unit size gain.

In the following we will refer to four experiments: two with λ^{c} equal to 1200 and two with λ^{c} equal to 1600. For each frequency, one experiment involved price seekers of sizes 1 to 5, and the other market samplers of the same five sizes. The setup with customer frequency λ^{c} equal to 1600 corresponds to the common intersection line in Figure 1.

All experiments were sampled after 200 periods [2], which corresponds to ap-

[2] Technical note: we used Java 1.4 (Sun JDK) for the simulation software. One 200 “period” experiment takes between 3 and 4 days of Pentium III 650 MHz processor

proximately 2.4 billions of customer entries for $\lambda^{\rm c} = 1200$ and 3.2 billions for 1600.

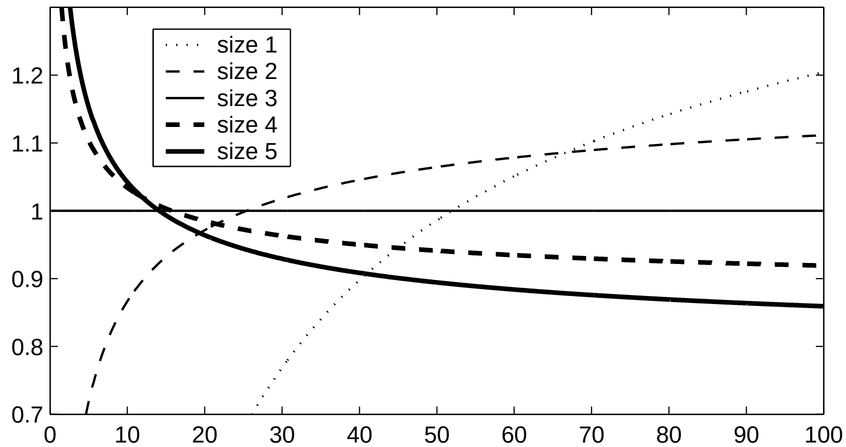

$\gamma(3, s, \cdot)$ for $s = 1, \ldots, 5$

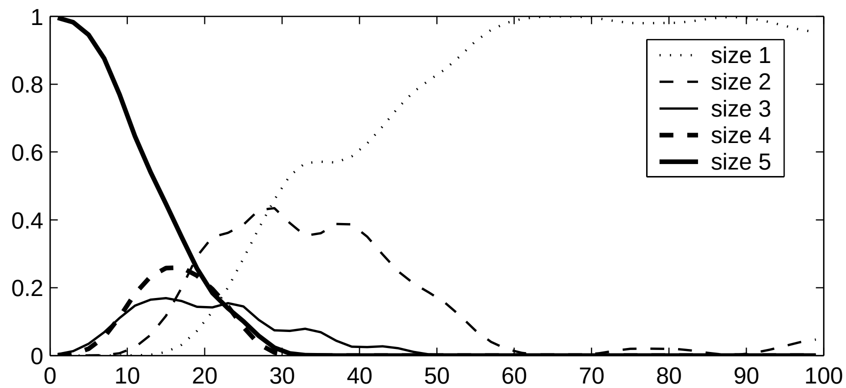

Percentage of capacity provided by suppliers of each size for price seekers

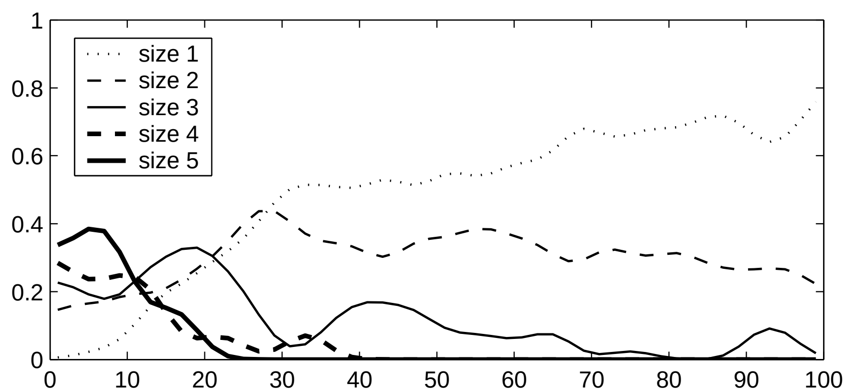

Percentage of capacity provided by suppliers of each size for market samplers

Fig. 2. Results for customer frequency $\lambda^{\rm c} = 1200$ and simulation time of 200 periods. Graphs of $\gamma(3, s, \cdot)$ for $s = 1, \ldots, 5$ (top graph), and of approximate percentages of capacity provided by suppliers of each size as a function of the distance from the center. The middle graph is for price seekers, the bottom is for market samplers.

computational time.

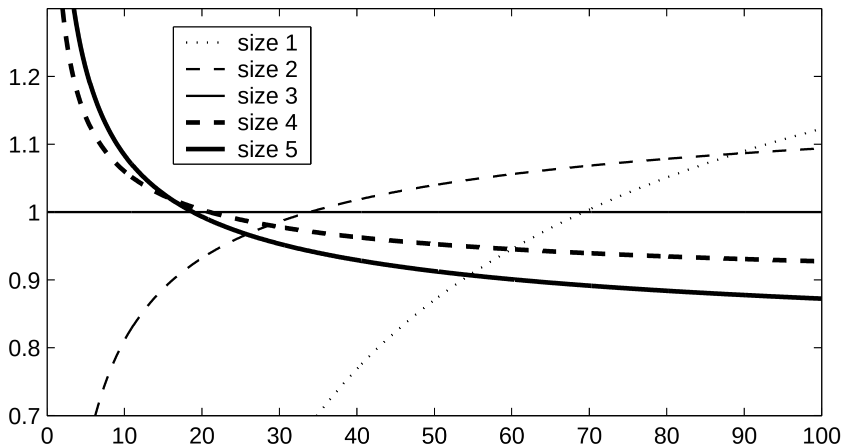

$\gamma(3, s, \cdot)$ for $s = 1, \ldots, 5$

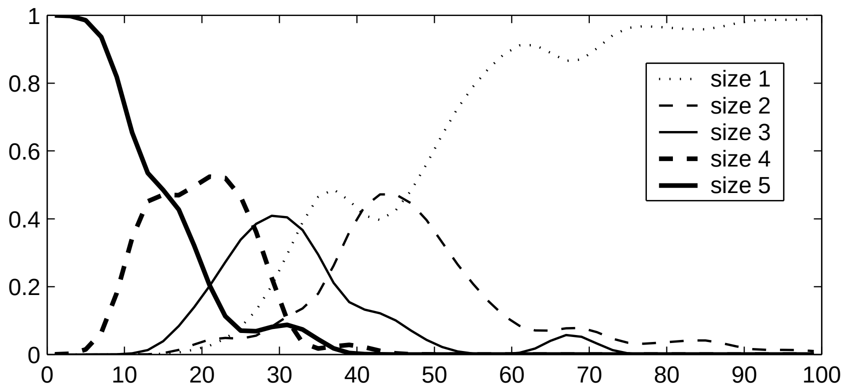

Percentage of capacity provided by suppliers of each size for price seekers

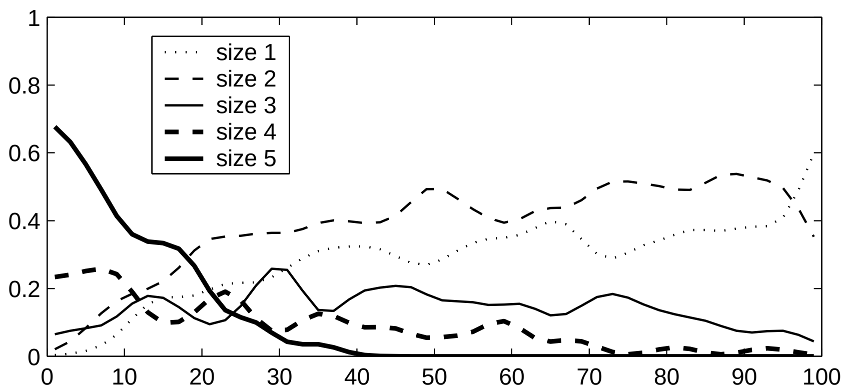

Percentage of capacity provided by suppliers of each size for market samplers

Fig. 3. Results for customer frequency $\lambda^{\mathrm{c}} = 1600$ and simulation time of 200 periods. Graphs of $\gamma(3, s, \cdot)$ for $s = 1, \ldots, 5$ (top graph), and of approximate percentages of capacity provided by suppliers of each size as a function of the distance from the center. The middle graph is for price seekers, the bottom is for market samplers.

Figures 2 and 3 show the results of the four reference experiments together with theoretic predictions for $\gamma(3, s, \cdot)$ for $s = 1, \ldots, 5$. We use the percentage of supplied capacity as a measure of supplier size's dominance at a particular

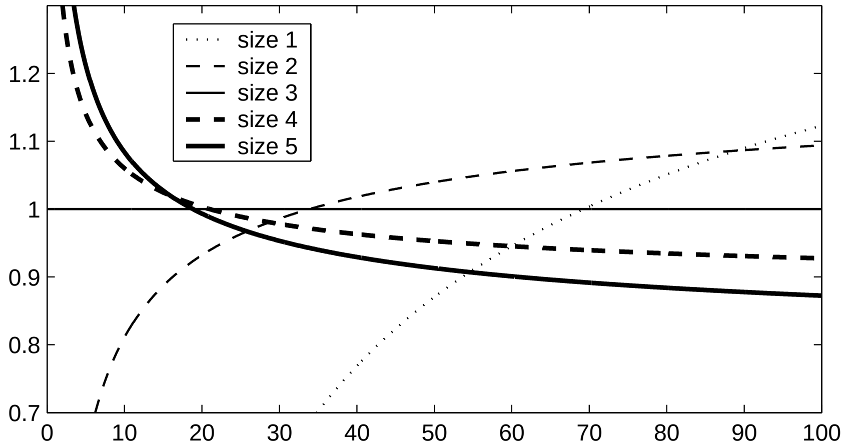

$\gamma(3, s, \cdot)$ for $s = 1, \ldots, 5$

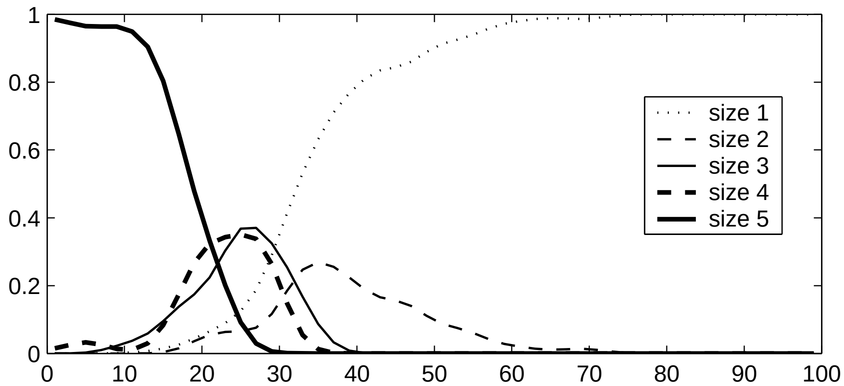

Percentage of capacity provided by suppliers of each size for price seekers

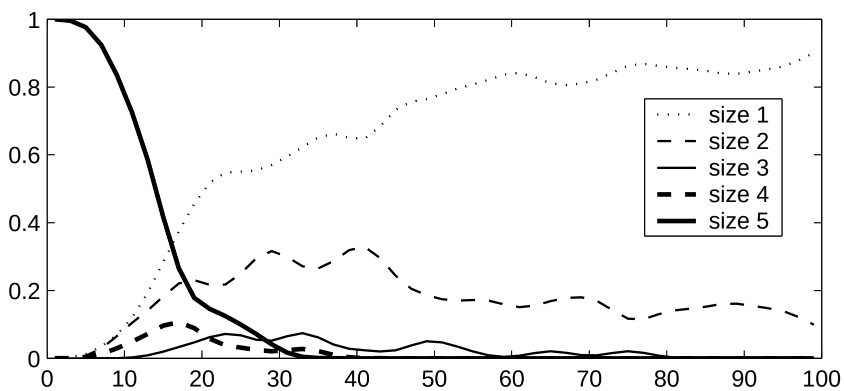

Percentage of capacity provided by suppliers of each size for market samplers

Fig. 4. Results for customer frequency $\lambda^{\mathrm{c}} = 1600$ and simulation time of 600 periods. Graphs of $\gamma(3, s, \cdot)$ for $s = 1, \ldots, 5$ (top graph), and of approximate percentages of capacity provided by suppliers of each size as a function of the distance from the center. The middle graph is for price seekers, the bottom is for market samplers.

distance from the center.

Observe that the trends in the experimental results follow the theory rather closely. First, all five sizes acquire considerable shares of the market, with larger suppliers tending to stay in the center of the city. Second, the interval around which size 3 has the highest representation in the market in the experimental results is not much different from the region where $\gamma(3, s, \cdot) \leq 1$ for all $s \neq 3$.

It is important to note that although the theory predicts the general state of the market quite satisfactorily, there exist noticeable discrepancies between the results obtained using different supplier strategies. Most spectacular, the separation of zones where one size is preferred over others is much more distinct in the price seeker experiments than in the market sampler ones. This, however, is a sign that the evolutionary experimenting provides a richer framework than the original theoretic model did.

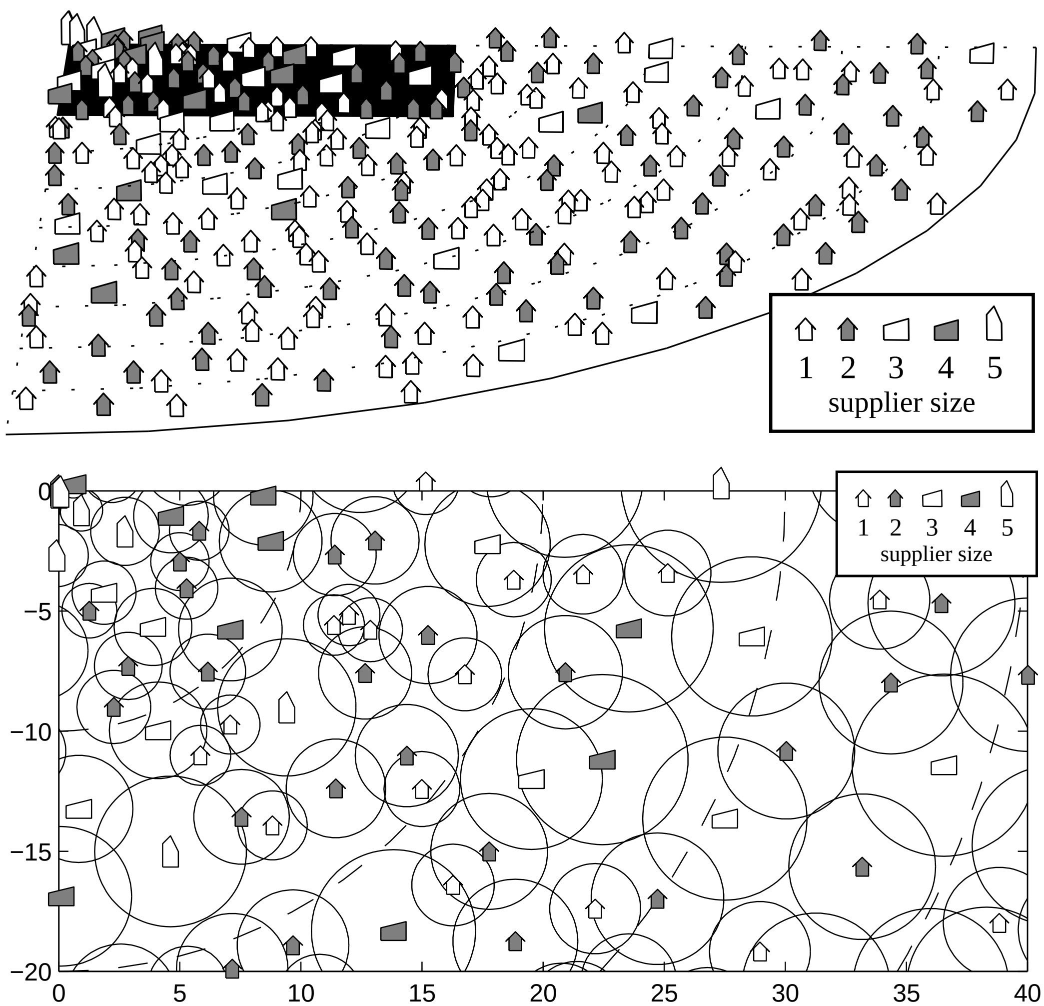

Fig. 5. One quarter of the city populated by 5 sizes of market samplers after 200 period s (top graph). The bottom graph shows a black area of the top graph with theoretic prediction of supplier supports (refer to Eq. (10)). Customer frequency λ^{c} is 1600.

The key to the differences in the results of simulating our two test strategies is in their deployment and pricing mechanisms. In particular, market samplers are designed such that upon deployment they explore the market in many spots and choose the one that promises them a higher revenue flow. This leads to a market structure where smaller suppliers can survive in between market areas of the bigger rivals, much alike a basket of ping-pong balls mixed with a basket of tennis balls fits in less than a two basket volume (see Figure 5).

Price seekers, in turn, lack the ability to pick the exact location of their deployment, hence their survival depends greatly on the ability to secure the supporting area by offering the right price. This explains why they follow the presented theory closer — the theory suggests at what distance from the center one should deploy to be able to sell at the minimum reasonable price. Indeed, observing the example of the price seekers' city in Figure 6 we may note that it is more common for suppliers to deploy near one another, thus stressing the ability to choose the right price.

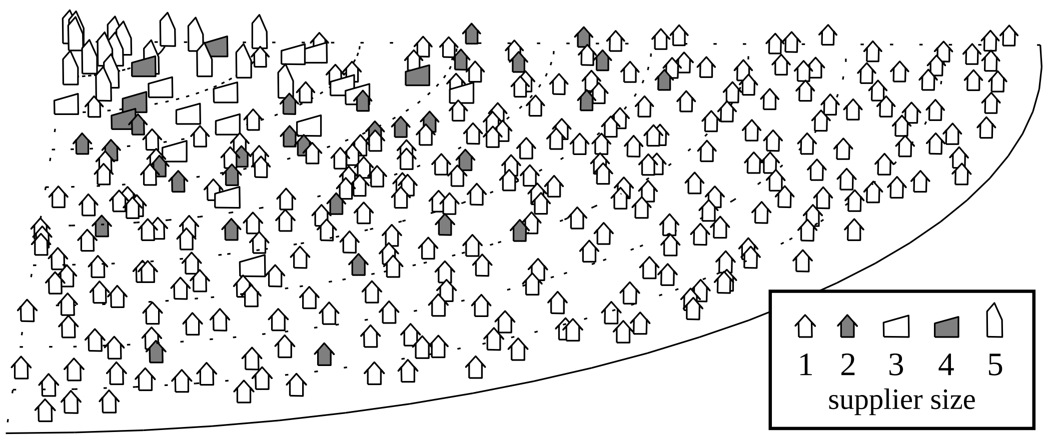

Fig. 6. One quarter of the city populated by 5 sizes of price seekers after 200 period s. Customer frequency λ^c is 1600.

A more interesting trend is that from the way the supports, depicted in the bottom part of Figure 5, overlap, it is clear that in the simulation the suppliers are able of surviving in higher density then explained by the theory. One reason for this is that in the simulation suppliers are not restricted by the assumption 1 of the theory, hence they charge higher prices and derive non-zero profits. Such behavior is necessary for surviving in the evolving market, since the arrival of customers is random, and the positions and prices of rivals are not guaranteed.

Our final observation is the distribution of average service delays faced by the customers upon entry. Figures 7 and 8 show the distributions of delays for customer entry frequency 1600, two supplier types, and two time points: 200 and 600 periods. We can derive several conclusions from the data:

(1) average delays are longer near city boundaries than in the center;
(2) in some areas delays are close to zero, confirming our guess that suppliers

are packed closer than predicted, since at least some of them have idle periods;

(3) in general, average delays are not much different from what is assumed in theoretical derivations, i.e. $1/2s$ for a supplier of size s;

(4) in market seekers' case, time delays become considerably more evenly distributed, as a result of more even distribution of suppliers.

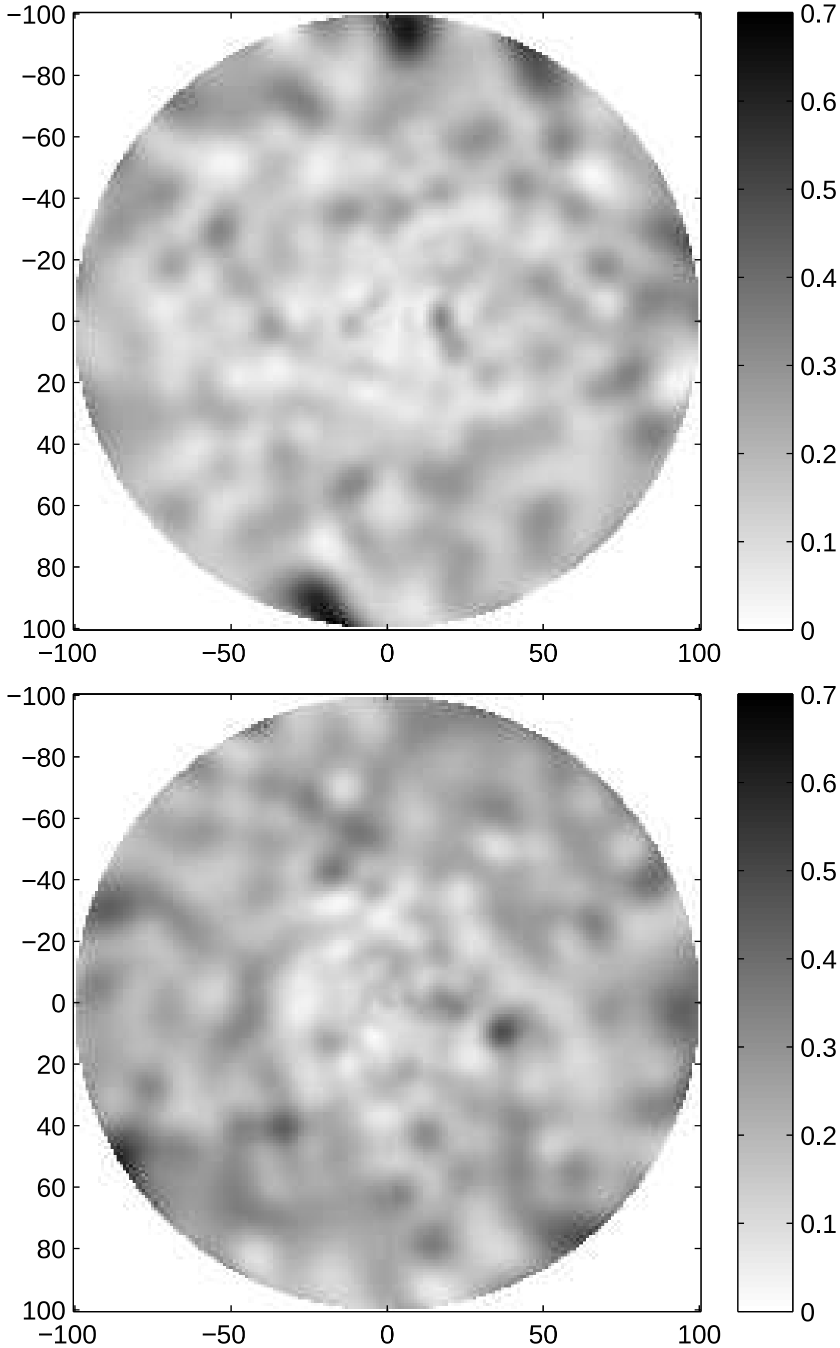

Fig. 7. Distributions of delays for 5 price seeker strategies, for 200 periods (top) and 600 periods (top) ($\lambda^{c} = 1600$).

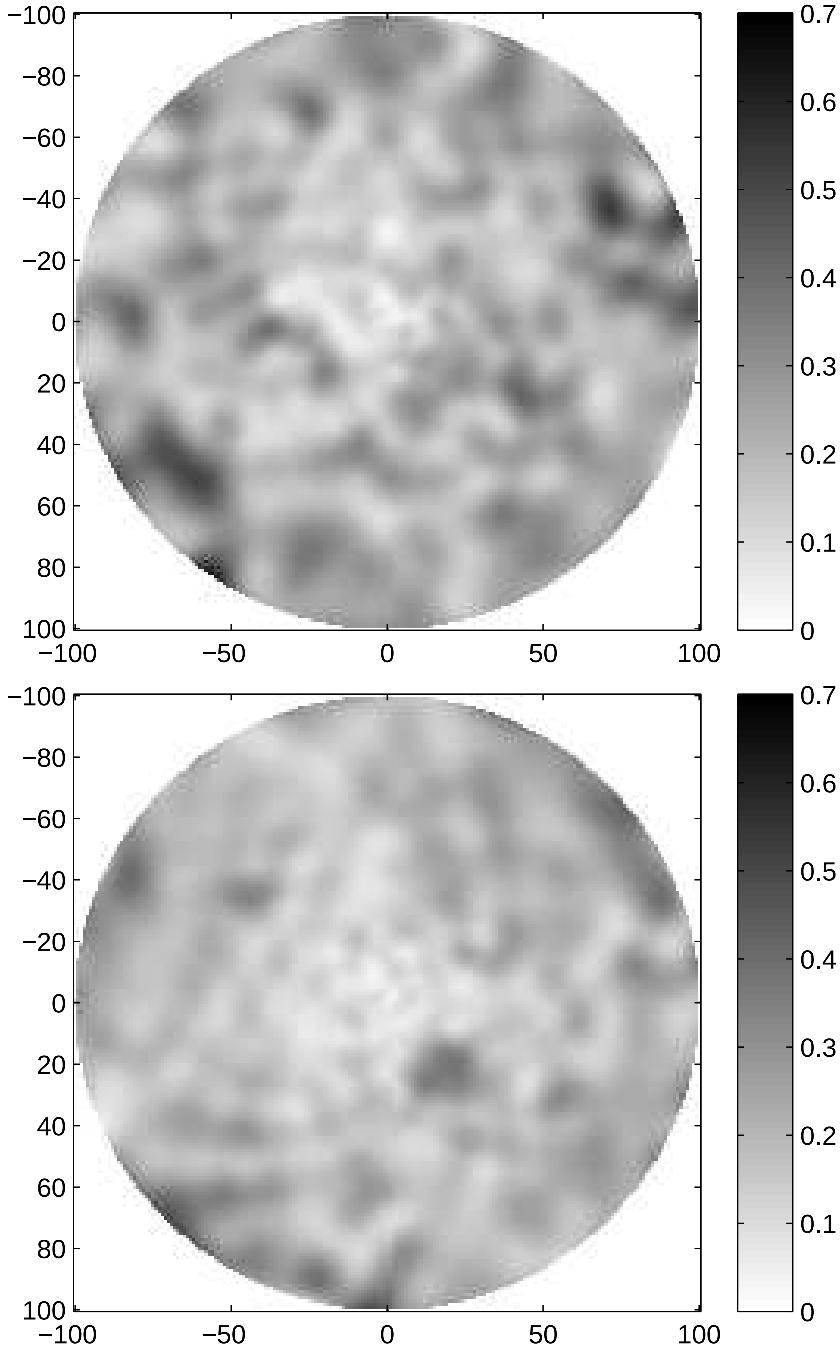

Fig. 8. Distributions of delays for 5 market sampler strategies, for 200 periods (top) and 600 periods (top) ($\lambda^{\mathrm{c}} = 1600$).

5.2 Supplier Strategy Entry Experiments

We devote our next set of experiments to verifying the expectation that new and retired strategies can enter the market by the means of randomly selecting one of them with a small probability, or noise (refer to Section 4). In a properly functioning model it should be possible for new types of suppliers 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.

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 an 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 9. Each periodin the figure corresponds to 10,000 simulated hours and roughly 2.5 million of transactions [4]. The market samplers enter the market at period 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 periods: 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 periods are depicted in Figure 9 by vertical lines.

Figure 10 suggests a schematic view of the city at period 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 10. Firstly, price seekers dominate the market and, indeed, their size 1 suppliers tend to 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 11 presents a view of the city at period 290 when the market is domi-

---

[3] For the previously introduced parameters we used the following values: $R = 50$ miles, $\lambda^{\mathrm{c}} = 250$ customers per hour, $c^{\mathrm{mile}} = 0.5$ \$/mile and $c^{\mathrm{hour}} = 1$ \$/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 periods to let the generators adjust their (or rather our) pretty volatile initial guesses about the market situation.

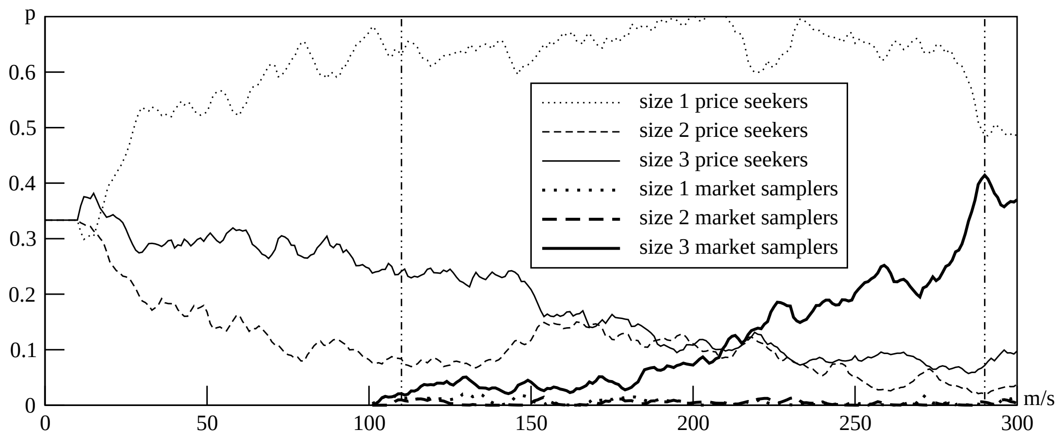

Fig. 9. Probabilities of a new supplier entry for different supplier types as a function of period numbers. Market sampler suppliers are introduced at period 100. Vertical lines denote periods 110 and 290.

nated by size 3 market samplers and size 1 price seekers. The structure of the market is evidently different from the one we witnessed 10 periods 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 12 with information on the prices and the costs that customers face at the two considered periods. The left graph shows average prices and standard deviations for the ten concentric city zones. The right graph shows the *decaying averages* [5] 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.

[5] 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.

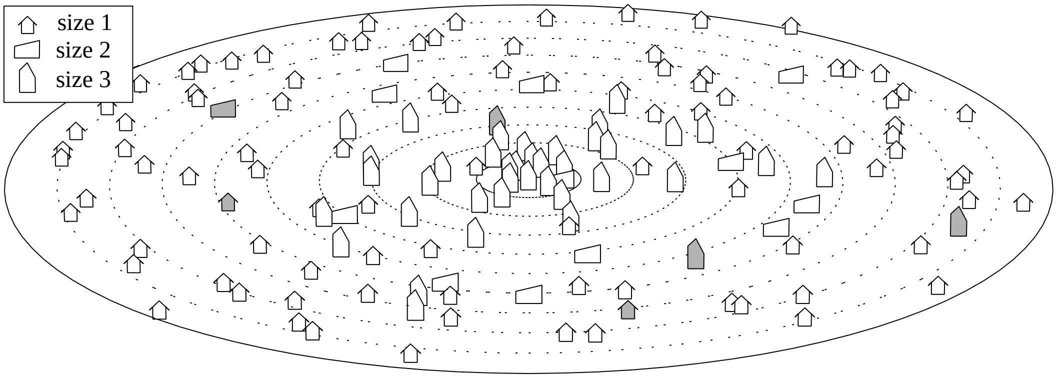

Fig. 10. City snapshot at period 110. Price seeker suppliers are white, market samplers are gray.

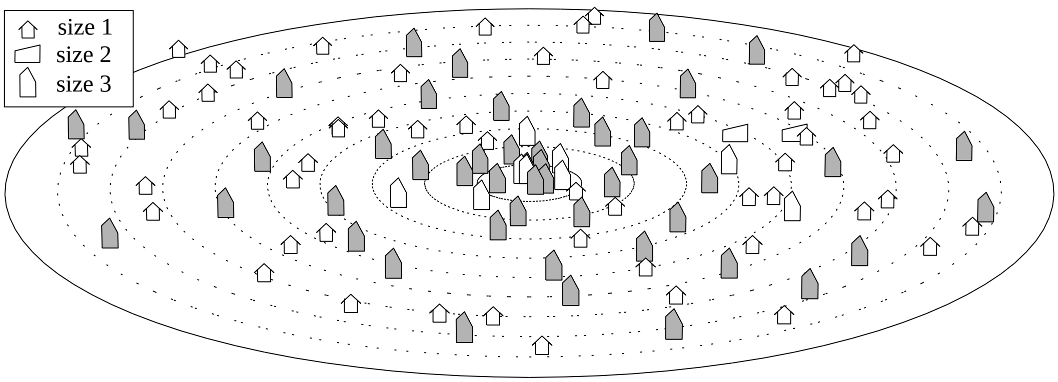

Fig. 11. City snapshot at period 290. Price seeker suppliers are white, market samplers are gray.

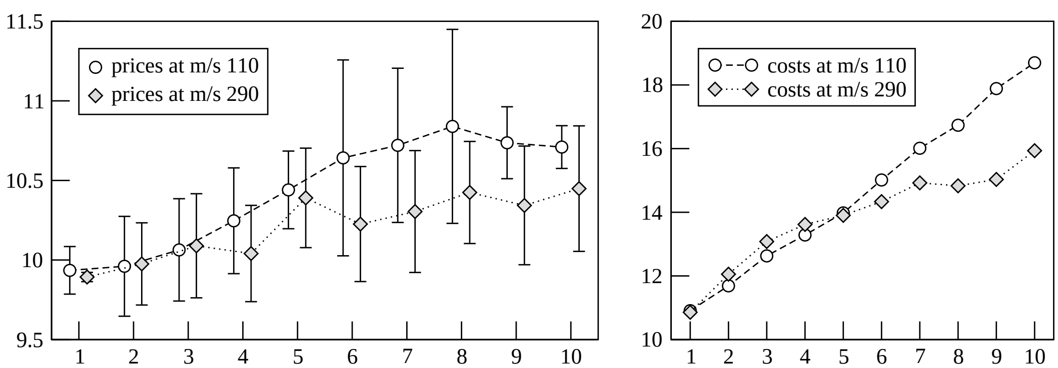

Fig. 12. Average supplier prices with standard deviations (left) and 25 hour half-life decaying averages of customer costs (right) for 10 concentric city zones at period 110 and 290.

6 Conclusion and Future Work

We introduced the idea of building a large-scale evolutionary framework for experimentation with heterogenous strategies on top of existing multi-agent systems. We have proposed an evolutionary framework where agents with different pricing and location selection strategies compete in a market. New agents are introduced 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.

7 Appendix: Summary of notation

Symbol	Definition
λ^{c}	Fixed frequency of the Poisson process for customers arrival
R	Radius of circular city
c^{mile}	Cost per mile of travel
c^{hour}	Cost per hour of delay
d	Distance to customer
Δt	Delay due to servicing previously scheduled customers
c	Net cost of a service
s	Size of service provider
$c^{\mathrm{work}}(s)$	Cost of serving a customer for a size s service provider
$c^{\mathrm{idle}}(s)$	Idle cost for a size s service provider
g	Gain due to supplier size
da	Decaying average of customer costs

References

[1] A. Anas, R. Arnott, K. A. Small, Urban spatial structure, Journal of Economic Literature 36 (3) (1998) 1426–1464.

[2] A. Babanov, W. Ketter, M. Gini, An evolutionary approach for studying

heterogenous strategies in electronic markets, in: Engineering Self-Organising Applications (ESOA 2003), Kluwer, 2003.

[3] D. Cliff, J. Bruten, Minimalintelligence agents for bargaining behaviors in marketbased environments, Tech. Rep. HPL-97-91, Hewlett Packard Labs (1997).

[4] R. Das, J. E. Hanson, J. O. Kephart, G. Tesauro, Agent-human interactions in the continuous double auction, in: Proc. of the 17th Joint Conf. on Artificial Intelligence, Seattle, WA, USA, 2001.

[5] J. M. DiMicco, A. Greenwald, P. Maes, Dynamic pricing strategies under a finite time horizon, in: Proc. of ACM Conf on Electronic Commerce (EC'01), 2001.

[6] J. M. Epstein, R. Axtell, Growing Artificial Societies: Social Science from the Bottom Up, The MIT Press, Cambridge, MA, 1996.

[7] J. O. Kephart, J. E. Hanson, A. R. Greenwald, Dynamic pricing by software agents, Computer Networks 32 (6) (2000) 731–752.

[8] K. Lerman, Design and mathematical analysis of agent-based systems, in: Lecture Notes in Artificial Intelligence (LNAI 1871), Springer Verlag, 2001, pp. 222–233.

[9] P. M. Markopoulos, L. H. Ungar, Shopbots and pricebots in electronic service markets, in: Game theory and decision theory in agent-based systems, Kluwer Academic Publishers, 2001.

[10] D. McFadzean, D. Stewart, L. Tesfatsion, A computational laboratory for evolutionary trade networks, IEEE Transactions on Evolutionary Computation 5 (5) (2001) 546–560.

[11] E. S. Mills, M. R. Lav, A model of market areas with free entry, Journal of Political Economy 72 (3) (1964) 278–288.

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

[13] A. O'Sullivan, Urban Economics, McGraw-Hill Higher Education, 1999.

[14] S. Park, E. H. Durfee, W. P. Birmingham, An adaptive agent bidding strategy based on stochastic modeling, in: Proc. of the Third Int'l Conf. on Autonomous Agents, 1999.

[15] S. Phelps, P. McBurney, S. Parsons, E. Sklar, Co-evolutionary mechanism design: a preliminary report, in: J. Padget, O. Shehory, D. Parkes, N. Sadeh, W. Walsh (eds.), Agent Mediated Electronic Commerce IV, vol. LNAI2531, Springer-Verlag, 2002, pp. 123–142.

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

[17] T. C. Schelling, Micromotives and Macrobehavior, W. W. Norton & Company, Inc., 1978.

[18] O. Shehory, S. Kraus, O. Yadgar, Emergent cooperative goal-satisfaction in large scale automated-agent systems, Artificial Intelligence 110 (1).

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