Learning Curve: Analysis of an Agent Pricing Strategy Under Varying Conditions

Joan Morris, Pattie Maes

MIT Media Laboratory
20 Ames Street, E15-305
Cambridge, MA 02139 USA
{joanie, pattie}@media.mit.edu

Amy Greenwald

Computer Science Department
Brown University, Box 1910
Providence, RI 02912 USA
amygreen@cs.brown.edu

Abstract - *By employing dynamic pricing, the act of changing prices over time within a marketplace, sellers have the potential to increase their revenue by selling goods to buyers "at the right time, at the right price." As dynamic pricing systems become necessary as a competitive maneuver and as market mechanisms become more complex, there is a greater need for pricing agents to be used, and also a greater challenge for sellers to understand what is the best agent pricing strategy for their marketplaces. This paper addresses these issues by presenting a market simulator designed for analyzing agent pricing strategies for a market in which a seller has a finite amount of time to sell a finite number of goods. Through an in-depth analysis of our 'Goal-Directed' pricing strategy, we demonstrate the use of our market simulator as a means for understanding the relevant factors in determining an effective dynamic pricing strategy.*

Keywords : Dynamic pricing, agent simulation, electronic marketplaces.

1. Introduction

With the increasing sophistication of market analysis and pricing tools available to sellers, dynamic pricing is becoming more common. While models and practices exist today for setting optimal prices, such as in the airline industry [1], there is a limit to the potential of dynamic pricing if human sellers have to make individual pricing decisions for each transaction. For a seller to compete in a rapidly changing, increasingly competitive marketplace, we suggest that sellers use software agents representing their interests to deploy intelligent dynamic pricing strategies.

If a seller's agent is aware of constraints on time and inventory and can observe changes in buyers' demand levels and purchasing behaviors, then it can attempt to sell all of a seller's goods at the highest possible prices over time. One of the difficulties in employing these real-time agents is understanding the costs and benefits to different agent pricing strategies. Our approach to studying dynamic pricing and the application of agents in a marketplace is to use a market simulator for testing and comparing different agent strategies.

Our market simulator, called the *Learning Curve Simulator*, is designed for testing agent pricing strategies under varied market conditions and realistic buyer behaviors. While a theoretical approach to agent pricing strategies could be taken, we believe that a theory-based solution is often difficult to apply to a real-world marketplace because of the overly simplifying assumptions that typically need to be made in developing a theoretical model. Simulated marketplaces are able to model more diverse and complex scenarios, rather than the general case. By producing tangible, numerical results, our Learning Curve Simulator can be used as a tool for understanding real-world scenarios.

This paper demonstrates how the Learning Curve Simulator can be used to analyze dynamic pricing strategies. Our Goal-Directed strategy, which we present and analyze here in detail, is an example of an adaptive pricing strategy that could be applied to a finite market – a market in which there is a finite amount of inventory and buyers under a finite time horizon. Through our analysis, we will demonstrate the strength a simulator provides in producing tangible guidelines for dynamic pricing strategies in finite markets.

2. The Pricing Strategy

In our previous study of agent pricing strategies [2], we analyzed the effectiveness of two different pricing strategies within a specific market scenario of an airline auctioning airline tickets. While one of the strategies, the Goal-Directed strategy[1], was extremely successful in that airline scenario, it was not tested under enough conditions to extend it to a general pricing strategy conclusion. In this paper, we present an analysis of the Goal-Directed strategy using the Learning Curve Simulator.

We call this strategy "goal-directed" because by adjusting price, the pricing agent attempts to reach a *goal* by the end of the market. In this case, its goal is to sell its entire inventory by the last day of the market, and not before. The agent accomplishes this by observing its success in selling goods each day and responding with incremental changes in price. For example, in the airline scenario we presented in [2], an airline has thirty days to sell 100 airline seats. On each day of the market, the airline receives bids from buyers and accepts the highest bids above its reserve price, and at the end of the day, the airline's pricing agent calculates its reserve price for the next day using the Goal-Directed strategy. To do this, the agent compares the number of seats it has sold to the amount of seats it expected to sell. If too few seats have been sold, the agent responds by lowing its original offer price by the percentage it is off-target. If too many seats have been sold, then the agent compensates by raising the price.

$$price_{i+1} = price_0 + price_0 * á * \frac{\sum_{n=1}^{i} goodsSold_n - \exp GoodsSold_i}{\left(\exp GoodsSold_i\right)}$$

$$\exp GoodsSold_i = i * \left(initialInventory \big/ daysInMarket \right)$$

Figure 1: Goal-Directed strategy calculation

In this manner, the agent fine-tunes the price of the good to the level of demand that enables the seller to sell all the goods by the last day of the market, and not before. This Goal-Directed strategy calculation is presented in Figure 1.

3. The Market Simulator

The Learning Curve Simulator is designed to model real-world markets and buyer behaviors, for the purpose of testing dynamic pricing strategies. To do this, the simulator requires three categories of inputs: the Market Scenario, the Buyer Behavior, and the Seller Strategies, enumerated in Table 1. Using these inputs, the simulator constructs and runs a simulated market in which buyer and sellers match on price, perform transactions, and change their behavior each day based on different market conditions. At the end of the simulation, the success of a pricing strategy is determined by the total revenue earned and the amount of inventory sold by each seller. The following sections explain the simulator input categories in more detail.

3.1 Market Scenario

Our research focuses on *finite markets*. In these markets, a seller has a certain amount of inventory it must sell by a certain date. There are many examples of this type of market in today's economy; a few examples are airline tickets, rental cars, theatre tickets, perishable items, and seasonal retail goods. The finite elements of the market are defined by the Market Scenario inputs (see Table 1). Unlike our previous investigation, we are analyzing a posted-price market mechanism, not an auction mechanism, in this version of the simulator.

3.2 Buyer Demand over Time

An integral part of a finite market is that the value of the good changes over time, whether by a change in buyer perception or the good's publicly known value. Thus, the Learning Curve Simulator models this change in buyers' perception of price, otherwise known as valuation, through a series of valuation/time curves. To test the robustness of a pricing strategy, we test it under five different buyer valuation/time curves: flat, increasing, decreasing, mid-peak, and mid-dip.

3.3 Variation Among Buyers Each Day

Another important aspect of the buyers' behavior is how the individual buyers differ from each other on a single day. We have modeled this in several different ways.

---

[1] In "*Sardine: Dynamic Seller Strategies in an Auction Marketplace*," this strategy was referred to as the Reserve Pricing Strategy.

First, there is a variance among buyers in their willingness to pay for a good, so the simulator calculates a distribution of buyer prices each day based on the input values for variance and private vs. public valuation (see Table 1). A private value good is one in which the buyer's willingness to pay is derived from his/her personal utility of the good. A public good is one in which the buyer's willingness to pay is based on the public's collective assigned value for the good. Depending on the simulator input value of private or public valuation, the distribution of buyer prices is a uniform (uncorrelated) or normal (correlated) distribution, with range defined by the variance value.

Second, for different types of markets, buyers are willing to search for the right price for different lengths of time. This is modeled in the simulator with the buyer lifetime variable. For each simulation, a lifetime value is selected that indicates how many days a single buyer will search for a seller offering the good at an acceptable price before leaving the market.

3.4 The Simulator Cycle

Given these market inputs, the simulator sequentially runs through each day of the market. On a single day, each active buyer, in random sequence, searches through the available sellers, in random sequence, and compares the seller's price with its own reserve price. If the seller's price is less, a transaction occurs and the buyer leaves the market. If the seller's price is more, the buyer continues looking. The day ends when each buyer has completed its search through the sellers. At the end of the day, a new reserve price for each buyer is calculated based on the valuation/time curve and variance values. Each seller updates its price based on its strategy calculation. If the seller is using the Goal-Directed strategy, the seller examines how many goods it has sold and what day it is, and makes a price adjustment. If the seller is using a Fixed-Price strategy, there is no change to the price. In this manner, the market progresses until the last day, stopping only if there are no more buyers or no more goods in the market.

More information on the inner workings of the simulator can be found in [3] and screenshots of the simulator's Java Swing interface can be found on-line at [4].

Table 1: Simulator Inputs

Simulator Inputs:	**Description**	**Sample Values**
Market Scenario:		
Number of Days	Number of periods in which the seller can implement a price change	30
Number of Buyers	Number of buyers within entire market	250
Number of Sellers	Number of sellers competing	2
Number of Goods	Number of goods per seller	100
Market Mechanism	Posted-Price only. (In later analyses, auctions could be included as a possible mechanism.)	Posted-Price
Buyer Behavior:		
Lifetime	Number of days each buyer is in market	1
Price Variance Per Day	The buyers' reservation prices vary ± the variance in a single day.	$100
Private or Public Valuation	Uniform or normal distribution	Private
Min/Max of Buyers' prices	The minimum and maximum average price desired by the buyers, over time	$200/$600
Valuation over Time Curve	The buyers' valuation/time curve can be either flat, increasing, decreasing, mid-peak, or mid-dip	All (flat, increasing, decreasing, mid-peak, and mid-dip)
Seller Strategies:		
Seller Strategy	Either Goal-Directed or Fixed-Price	1 Goal-Directed, Fixed-Price
Initial Price	The price the seller offers the first day of the market. In the case of a Fixed-Price Seller, this will be the price offered on all days of the market.	$400
Available Inventory per Day	Amount of inventory a seller can sell in one day	9 goods per day (3* inventory/days)

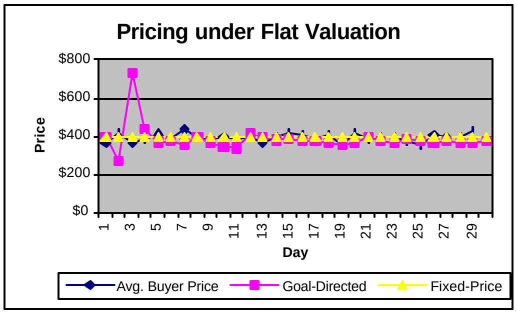

2A: Flat Valuation Curve

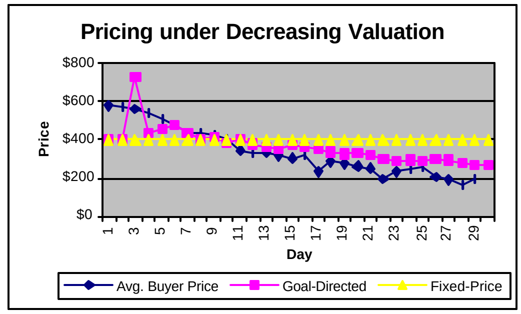

2B: Decreasing Valuation Curve

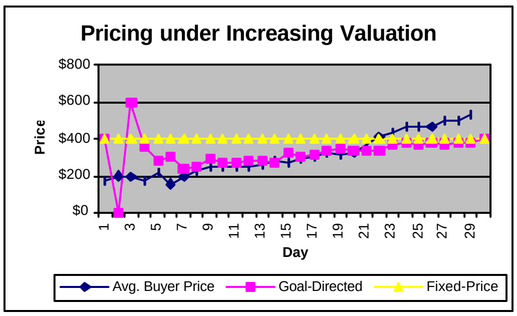

2C: Increasing Valuation Curve

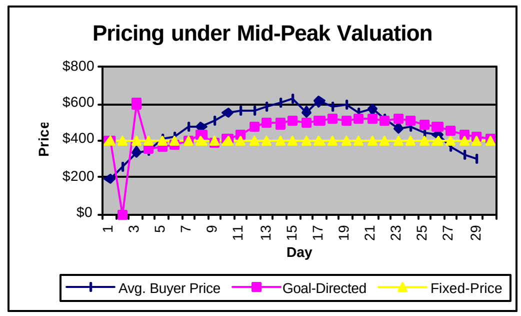

2D: Mid-Peak Valuation Curve

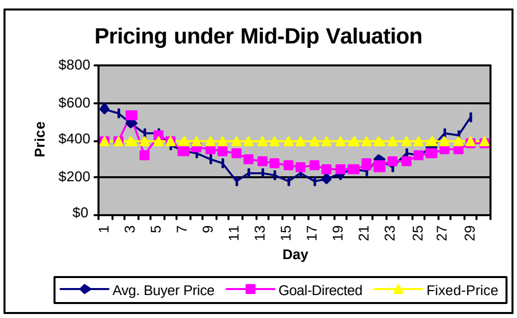

2E: Mid-Dip Valuation Curve

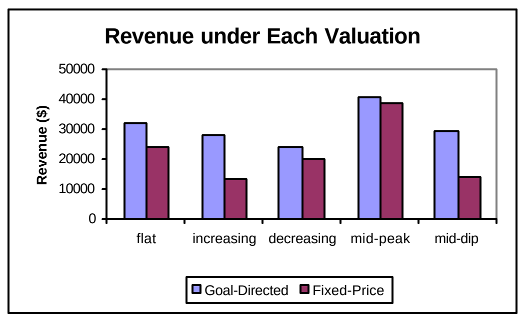

2F: Revenue per Valuation Curve

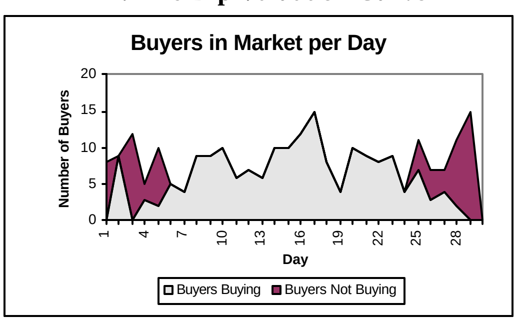

2G: Buyers in Market Each Day

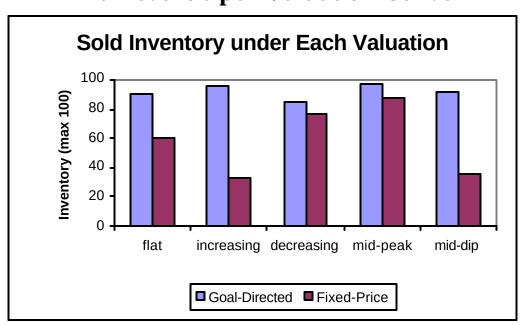

2H: Inventory Sold per Valuation Curve

Figure 2: Sample Simulation Results

Figures 2A-2E show the pricing behavior of a Goal-Directed (GD) seller and a Fixed-Price (FP) seller under different valuation conditions. In each case, the FP seller offers $400 while the GD seller adjusts price each day based on the amount of inventory it has sold at each point in the market. The revenue each seller earns under each valuation condition is shown in Figure 2F. As shown, the GD seller captures more revenue under each valuation curve, even under a flat valuation curve. In the specific case of flat valuation, the GD strategy prevailed by adapting to the high variance among the buyer population (on a single day, price ranging between $300 and $500). In each trial, there were 250 buyers, each appearing in the market for one day, and 100 goods per seller. Figure 2G shows when the buyers appeared in the market and when they made purchases, for the case of mid-peak valuation. Figure 2H shows that the GD strategy consistently sells nearly its entire inventory, which results in higher revenue despite the often lower sale prices.

4. Strategy Analysis

We tested the Goal-Directed strategy under many different market conditions in order to understand its strengths, varying the shape of the valuation/time curve, the amount and type of competition in the market, and the size and behavior of the buyer population on a per day basis. Based on these numerous simulations, we have made several conclusions about the effectiveness of the Goal-Directed strategy.

Before outlining these results, we present here a sample set of simulation results to illustrate our analysis process. Figure 2 shows the results of five simulation trials generated from the input values listed in the right column of Table 1. These simulations each had two sellers, one using the Goal-Directed strategy and the other using a Fixed-Price strategy. The initial price for both sellers was the average of the buyer price range, $400. Figures 2A-2E show the pricing behavior of the sellers under the five different valuation/time curves: flat, increasing, decreasing, mid-dip, and mid-peak. These charts illustrate the characteristic pricing pattern of the Goal-Directed strategy. The first days of the market are characterized by extreme over and undershooting of the buyers' average price as the strategy adjusts for over- and under-selling. As the days of the market progress, the price changes become less extreme as the strategy begins to track the buyers' valuation curve. This following of the buyers' valuation curve is what makes the strategy so effective: regardless of the type of buyer behavior presented, the Goal-Directed strategy is able to compensate for the change in buyer behavior and achieved its goal of selling goods by adjusting price. As shown in Figures 2F and 2H, the Goal-Directed strategy sells more inventory and earns more revenue, despite often offering lower prices than the seller using a Fixed-Price strategy. Figure 2G shows the number of buyers in the market for the specific case of the Mid-Peak valuation curve, illustrating that most of the sales occurred during the middle day of the market when valuation peaked, with a few significant sales occurring at the beginning of the market when the Goal-Directed strategy captured critical sales.

5. Analysis Conclusions

The results presented in Figure 2 demonstrate the strengths of the Goal-Directed strategy. We analyzed the Goal-Directed strategy under a spectrum of market, buyer, and seller conditions, which resulted in varied levels of success over a Fixed-Price strategy. Here we present our analysis of when and why the Goal-Directed strategy experiences success over a Fixed-Price strategy.

1. *Initial Price. The further the seller's initial price is from the optimal sale price, the better the Goal-Directed strategy performs.*

Each seller in the simulator offers an initial price for the good, as indicated in the simulator inputs, and a seller using a Fixed-Price strategy will offer this same price every day. In our trials, all the sellers began with the same initial price, under the assumption that the initial price is the sellers' best guess at the optimal price which will sell all the inventory at the highest price. In the example in Figure 2, this initial price was set to $400. If sellers do not make an optimal price decision, due to limited or incorrect knowledge about the buyer population, a Fixed-Price seller will be unable to sell its inventory and the Goal-Directed strategy will prevail by adjusting its price. If a seller can make an accurate prediction on how buyer demand will change over time, then the seller can confidently pick an optimal price and achieve maximum revenues, but in the more common situation where there is incomplete information, a Goal-Directed strategy allows a seller to adjust for mistakes.

2. *Seller Inventory. If a seller is limited in the amount of inventory it can sell each day, the Goal-Directed strategy wins.*

When the amount of inventory that can be sold each day is restricted, whether because of a shelf-stocking fee or an impracticality of selling the entire inventory in one day, then a Fixed-Price strategy is unable to take advantage of a slim window of opportunity to sell all of its goods at a high price. The Goal-Directed strategy on the other hand, paces its sales across the market. Under a condition where a seller can sell everything quickly, then a high fixed-price can work best, but if that is not possible, then the Goal-Directed strategy will ensure that all or the majority of the inventory is sold at the highest price to be gotten on an individual day.

3. *Number and Lifetime of Buyers. In a market with a limited number of buyers with a limited lifetime, a Goal-Directed strategy is able to sell to more buyers than a Fixed-Price strategy.*

When the number of buyers is close to the number of goods available in the marketplace, each seller needs take advantage of each day buyers are available in the market. Because the Goal-Directed strategy focuses on selling goods every day, a Goal-Directed seller is able to sell more inventory than a Fixed-Price seller.

4. *Competition. The more Fixed-Price sellers in the market, the better the Goal-Directed sellers perform, given a limited numbers of buyers in the market.*

Increasing the number of Goal-Directed sellers in the market does not significantly change the performance of the individual Goal-Directed sellers. But when we increased the number of Fixed-Price sellers in the market, the Goal-Directed sellers increased the amount of earned revenue and sold inventory in relation to the Fixed-Price sellers. The Goal-Directed strategy's adaptive prices allow a seller to sell to a higher proportion of buyers than the Fixed-Price seller, effectively stealing sales from the Fixed-Price seller. The more Fixed-Price sellers in the market, the more these results are exaggerated.

5. *Variance in Buyer Price. The higher the variance in price among the buyers, the better the Goal-Directed strategy performs.*

If, on a single day, every single buyer wants to pay the same price, then it would be possible for a Fixed-Price seller to pick the right price and sell the entire inventory. A more realistic situation is that most buyers consider multiple factors in calculating their reservation price and in this situation the Goal-Directed strategy works effectively. When there is a high variance among the buyers, the Goal-Directed strategy adjusts and fine-tunes its price within the spread of buyer prices to sell the highest paying customers each day.

6. *Buyer Valuation/Time Curve. A Fixed-Price strategy performs best with curves that are at a relative high early in the market, but a Goal-Directed strategy performs consistently for all types of curves.*

When the demand is at a high point in the early days of the market, a Fixed-Price strategy can successfully sell all or most of its inventory in the early days, while a Goal-Directed seller is observing and make drastic adjustments to price. Over longer market periods when buyers change their behavior in unexpected ways, the Goal-Directed strategy has the time to learn the valuation/time curve at the beginning and then consistently outperforms a Fixed-Price strategy for the duration of the market. We consider one of the key strengths of the Goal-Directed strategy to be this ability to quickly learn a curve and then, as it follows the curve, decrease the amount of price adjustment each day.

7. *Number of Days. The fewer days in the market, the less effective the Goal-Directed strategy.*

The sample results in Figure 2 contained thirty days, which provided enough time for the Goal-Directed strategy to adapt to changes in buyer behavior. If the market has only seven days, the seller does not have enough time to observe and adjust its price, yet when the market contains more than thirty days, the performance of the Goal-Directed strategy improves further because it has more cycles in which to learn the buyers' behavior. So in a market with a limited number of days, or cycles in which to change prices, a Fixed-Price strategy could be the better strategy.

6. Related Work

Theoretical studies of pricing strategies in finite markets have made conclusions about optimal pricing strategies, but the drawback of these theoretical approaches is the difficulty in applying the results to real-world markets. Gallego & van Ryzin [5], for example, studied this problem with an assumed, flat valuation/time curve (i.e. a static demand curve). The benefit of using the Learning Curve Simulator is its ability to model diverse and complex scenarios, rather than only simplified cases.

Our investigation of agent pricing strategies is unique in its use of a simulator to model a finite market. Several researchers have studied pricing strategies in simulated information-good marketplaces [6-9], in which inventory and time are not constraints. In these markets, the best strategy is the one that competes best in a competitive market. The additional complexities of constraints on time and inventory further illustrate the usefulness of studying and developing strategies in a market simulator.

Of current industries employing dynamic pricing, the airline industry sets the standard for dynamic pricing by using techniques of revenue management to implement automated price changes over time [1, 10]. Commercial revenue management systems forecast demand, monitor booking activities and, in response, adjust the number of tickets available at each pre-defined pricing level, or 'fare class.' This method is effective and practiced in other industries as well, but requires sellers to make assumptions and predictions about the behavior of the marketplace. Our Goal-Directed strategy makes no assumptions or predictions about future demand, but instead learns to adapt to the current demand levels. The Learning Curve Simulator allows multiple types of strategies to be analyzed and compared against one another, providing a method for comparing dynamic pricing approaches.

7. Conclusion

We found that the Goal-Directed strategy, designed for a finite market, works best in a market in which a seller has restrictions on when and how much it can sell. When a limited number of buyers, competitive factors, or restrictions on selling practices constrict the amount of inventory sold each day, the Goal-Directed strategy prevails over a Fixed-Price strategy because it focuses on consistently making daily sales through basic price adjustments.

These conclusions demonstrate the strength of our simulation-based analysis. The tangible results we found can now be used to inform a real-world seller's process of designing a Goal-Directed or similar dynamic pricing strategy for their finite market. The complexities of a finite market make theoretical studies challenging, and our belief is that a simulator with a rich set of market variable conditions allows for straightforward strategy development and analysis. Through this process sellers can gain an understanding of which strategies are best for their markets.

In terms of future work, our analysis results lead us to more questions and inform the immediate direction of our research. In addition to developing more strategies, we plan to add more realistic behavior to the simulated buyers, eventually applying adaptive buying behavior that responds to the sellers' dynamic pricing behavior. The Learning Curve Simulator will provide the platform for this further market modeling and strategy analysis.

References

[1]. Boyd, E.A., *Airline Alliance Revenue Management.* ORMS Today, 1998(October 1998).

[2]. Morris, J., P. Ree, and P. Maes. *Sardine: Dynamic Seller Strategies in an Auction Marketplace*. Proceedings of the *2nd ACM Conference on Electronic Commerce (EC '00)*. 2000. Minneapolis, Minnesota.

[3]. Morris, J., *A Simulation-based Approach to Dynamic Pricing*, Master's Thesis in *Media Arts & Sciences*. 2001, MIT: Cambridge, MA.

[4]. *Learning Curve Simulator Screenshots, http://www.media.mit.edu/~joanie/learningcurve/screenshots/*. 2001.

[5]. Gallego, G. and G. van Ryzin, *Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizons.* Management Science, 1994. 40 (8): p. 999-1020.

[6]. Deck, C.A. and B.J. Wilson. *Interactions of Automated Pricing Algorithms: An Experimental Investigation*. Proceedings of the *2nd ACM Conference on Electronic Commerce (EC '00)*. 2000. Minneapolis, Minnesota.

[7]. Brooks, C.H., S. Fay, R. Das, J. MacKie-Mason, J. Kephart, and E. Durfee, *Automated Strategy Searches in an Electronic Goods Market: Learning and Complex Price Schedules*. Proceedings of the *ACM Conference on Electronic Commerce (EC '99)*. 1999. Denver, CO.

[8]. Greenwald, A., J.O. Kephart, and G.J. Tesauro. *Strategic Pricebot Dynamics*. Proceedings of the *ACM Conference on Electronic Commerce (EC '99)*. 1999. Denver, CO.

[9]. Kephart, J.O. and A. Greenwald. *Shopbot Economics*. Proceedings of the *Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*. 1999.

[10]. Smith, B.C., D.P. Gunther, B.V. Rao, and R.M. Ratliff. *e-Commerce and Operations Research in Airline Planning, Marketing, and Distribution.* Interfaces, Special issue on Operations Research in the e-Business Era, 2001. 31 (March/April 2001).