Learning Curve:
A Simulation-based Approach to Dynamic Pricing

Joan Morris DiMicco, Amy Greenwald, Pattie Maes

MIT Media Laboratory
20 Ames St
Cambridge, MA 02139 USA
{joanie, pattie}@media.mit.edu

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

By employing dynamic pricing, sellers have the potential to increase their revenue by selling their goods at prices customized to the buyers' demand, the market environment, and the seller's supply at the moment of the transaction. As dynamic pricing becomes a necessary competitive maneuver, and as market mechanisms become more complex, there is a growing need for software agents to be used to automate the task of implementing implementation of instantaneous prices changes. But prior to using dynamic pricing agents, sellers need to understand the implications of agent pricing strategies on their marketplaces. The following article presents the Learning Curve Simulator, a market simulator designed for analyzing agent pricing strategies in markets under finite time horizons and fluctuation buyer demand. Through an in-depth description of the simulator's capabilities and an example of strategy analysis, we demonstrate the strength of a simulation-based approach to understanding agent pricing strategies.

Keywords: dynamic pricing, software agents, simulation, electronic commerce, market design.

1. Introduction

Today, when a ballpark sells baseball tickets, the park charges the same price for the tickets throughout the season. Yet the demand for tickets changes over time, depending on the length of time before the game, the team's success over the season, and unpredictable factors such as the weather. In a best-case scenario, a park sells all of its seats for every game at an optimal fixed ticket price. In a more realistic scenario, some days the park has empty seats and on other days the park is filled with buyers willing to pay more. Nonetheless, today ballparks leave the practice of dynamic pricing to scalpers.

Dynamic pricing, defined as the changing of prices in a marketplace, can be implemented in several different ways. Price discrimination, or personalized pricing, is an intriguing area of dynamic pricing in which sellers charge different segments of customers different prices. While this area is rich with potential, it also has greater risks of customer rejection, as exhibited when Amazon.com experimented with charging customers different prices [1]. In contrast to this approach to dynamic pricing, this body of work focuses on the changing prices over time in a market that makes no assumptions or attempts to segment the buyer population into sub-groups. This perspective on dynamic pricing focuses on how a seller can take advantage of the fluctuations in cumulative buyer demand over time, taking into account a finite time horizon. In this article, we refer to this type of changing of prices over time as dynamic pricing.

Cost is perhaps the greatest factor precluding the widespread use of dynamic pricing by ballparks and other markets, because in traditional markets, it is expensive to continuously re-price goods. But in digital markets, the costs associated with making frequent, instantaneous price changes are greatly diminished [2].

A remaining obstacle that hinders widespread dynamic pricing is the difficulty in understanding the complexities price changes introduce into a market. Now that sellers can easily implement frequent adjustments to price, how should they do so? What are the most effective dynamic pricing strategies, and how do they behave in specific markets? To answer these questions, we propose that sellers analyze dynamic pricing algorithms using a market simulator that is capable of simulating many different market scenarios with realistic models of buyer behavior. Using a market simulator, a seller can model its market's characteristics and the behavior of its customers, to develop a pricing strategy that can capture more profit than fixed-price policies.

To illustrate our proposed approach, we present the Learning Curve Simulator, a platform for running dynamic pricing algorithms in simulated markets. Through an analysis of different pricing strategies under varying market conditions, we demonstrate how, by observing market conditions, a seller can take advantage of fluctuations in buyer demand to earn more revenue and sell more inventory.

Our investigation of dynamic pricing strategies focuses on an extremely common market type, which we call a finite market -- a market with a finite time horizon, seller inventory, and buyer population. In markets under a finite time horizon, such as event tickets, airline tickets, hotel rooms, perishable goods, and seasonal retail, a clear benefit to dynamic pricing is that one can ensure all inventory is sold. Also inherent to the finite nature of these markets is an increased importance of fluctuations in consumer demand. By taking advantage of these demand changes, seller can carge higher prices at different points in time. Given these factors, it seems likely that in the near future, dynamic pricing will become a common competitive maneuver, particularly in markets under a finite time horizon.

Facing the need to liquidate inventory, sellers in finite markets often choose to sell remaining inventory in a side market where it is referred to as "distressed inventory." An example of this market on-line is FairMarket's AutoMarkdown [3]. AutoMarkdown runs as a multi-unit Dutch auction in which items are initially offered at high prices and then offered at progressively lower prices, down to a specified minimum, or until all inventory is sold. While AutoMarkdown's pricing strategy is basic and does not respond to demand in the marketplace, it is a good example of how dynamic pricing can achieve a finite market's seller's goal of selling all of its inventory.

We will present strategies in this article designed for a finite market where the interplay of time, inventory, and revenue determine the seller's success. While more sophisticated than the pricing strategy of AutoMarkdown, our strategy algorithms are still basic in that they make no assumptions about the behavior of the buyers or the type of buyers in the marketplace. Through incremental adjustments in price, these strategies are designed to adapt and learn the behavior of the marketplace, responding to any type of change. While any price changing strategy can be termed a "dynamic pricing strategy," we also refer to these strategies as "adaptive" because of their ability to observe and adapt to market conditions.

In the following section, we discuss the theoretical underpinnings for this research, with a presentation of related work done in the area of dynamic pricing. Next, we present the design and implementation of the Learning Curve Simulator, from the perspective of the user interface. In the next section, Strategy Analysis, we use the simulator as a tool for evaluating two pricing strategies. Our hope is that in addition to demonstrating the power of a simulation-based approach to strategy analysis, our specific strategies will lay the groundwork for designing more complex algorithms to be deployed in real-world markets. We conclude with an outline of future directions for this work.

2. E-Markets and Dynamic Pricing

Electronic markets have dramatically reduced the cost of making changes to price [2], so for the first time sellers are able to realistically make immediate and timely adjustments to price. As evidence of this, several on-line businesses today make automated adjustments in price, as much as every hour. An example of one such on-line business is Buy.com. As described by [2], Buy.com uses software agents to search competitor's web sites for competing prices, and in response, Buy.com lowers its price to match or beat these prices. Their pricing strategy is based on the assumption that their customers are extremely price sensitive and will choose to purchase from the seller offering the lowest price. Not surprisingly, Buy.com has managed to garner enormous sales, but their profits are extremely low, or even negative.

The example of Buy.com highlights two things. First, automated dynamic pricing is a feasible option for companies today. Second, an overly simplistic or incorrect model of buyer behavior can produce undesirable results. Today's economy is ready for dynamic pricing on a more complex scale: more complex in its understanding of buyer behavior and its pricing algorithms. With these changes, sellers stand to increase profits through dynamic price adjustment.

2.1 Today's Example: Revenue Management

The airline industry provides a more sophisticated example of dynamic pricing in today's economy. The airlines use the technique of revenue management to dynamically adjust prices over time by adjusting the number of seats available in each pre-defined fare class, or booking class [4-6]. Commercial revenue management systems forecast demand, monitor booking activities and, in response, adjust the number of tickets available at each fare level. This method is extremely profitable for the airlines and practiced in other industries such as hotel rooms, cruises, and rental cars. Its success is based on these industries' ability to segment their buyers into different groups with different levels of willingness to pay. Two distinctions are made between revenue management and dynamic pricing. Biller et al. clarify that buyer segmentation is not a necessary aspect of dynamic pricing [7]. Additionally, the techniques of revenue management require sellers to make sophisticated assumptions and predictions about the behavior of the marketplace. This limitation was addressed by Gallego and van Ryzin in their discussion of the need to merge the ideas of revenue management with dynamic adjustment of prices [8], where pricing is determined in response to consumer demand. As the revenue management industry exists today, the prices in each fare class are fixed, yet these price levels influence the market. For example, when the lowest fare class is sold out, the demand for the second-lowest fare class increases. In their work, Gallego & van Ryzin propose a model for blending revenue management, with price adjustments based on observed demand, and suggest that this model of price adjustment be applied to new industries, such as the fashion and retail industries. In spite of these differences, the airline industry's adjustment of prices over time still demonstrates the potential of earning more revenue by charging "the right customer, the right price, at the right time."

2.2 Buyers in Electronic Markets

While methods exist for using historical data to predict market behavior [5], the potential problem with using previous data to make assumptions about the future, is the risk of being wrong. For example, there is increasing evidence that while the search costs of finding products on the Internet are lower than in the off-line world, there is not a corresponding increase in buyers' sensitivity to prices as previously predicted [9]. Even with tools such as shopbots performing the task of locating goods and comparing prices, buyers seldom purchase from the lowest priced seller, revealing that they have a more complex utility function for that product or seller. Additionally, when buyers have more information about a product, as is the case in an electronic market, they become even less price sensitive [10]. Another interesting observation of on-line markets is that price dispersion, traditionally thought to be caused by high search costs, can still be high in an environment of low search costs, presumably because buyers have preferences for certain products and sellers [11].

The new purchasing environment created by electronic markets has revealed new and somewhat unpredicted buyer behavior. Initial attempts at providing buyers with shopping assistance (shopbots) and initial use of software agents to adjust prices (Buy.com) both assumed that buyers were extremely price sensitive. Because this has been shown to not be the case, there is a need for more complex tools for buyers [12,13] and for sellers. We propose the Learning Curve Simulator as a tool that will allow sellers to deploy dynamic pricing in an electronic marketplace filled with complex buyers.

2.3 Theoretical Studies

Earlier work of Gallego & van Ryzin [14] presented a theoretical model for calculating optimal prices for finite markets. This model addressed the challenge of dynamic pricing in markets with a finite time horizon and inventory, but from a theoretical standpoint. They examined a deterministic version of the problem of pricing under finite time horizons by making the assumption that consumers' demand curves do not change over time. Under these conditions, they concluded that the optimal pricing strategy is "jittery" and requires constant price adjustments, something they considered to be infeasible at their date of publication (1994). They concluded that a fixed-price strategy works "surprisingly well" when the demand curve is known. A "nearly optimal solution" is to have a fixed set of tiered prices that the seller oscillates between, and this is proposed as a more feasible solution than the optimal solution (of continual, incremental price adjustment).

These results can be easily duplicated in the Learning Curve Simulator. When the demand curve is known, a best fixed price can be selected to nearly optimize revenue, even under cases of changing demand curves over time. But what our analysis of pricing strategies emphasizes is that one cannot assume perfect knowledge of the demand curve, something to which Gallego and van Ryzin concede is more realistic.

In a recent analysis of the automotive industry [7], Biller et al. designed a theoretical model for applying dynamic pricing to a marketplace with unknown changing demand levels. They demonstrate that under fluctuating demand there is always a dynamic pricing strategy which is more successful than a fixed-price strategy. The degree of success increases with the amount of variance in reservation price within the buyer population and the number of times the seller adjusts prices. While their model [15] does not account for constraints in inventory as our does, these findings are similar to the conclusion we made using the Learning Curve Simulator.

2.4 Simulation-based Approach

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, several researchers have made significant contributions to the problem of dynamic pricing.

Researchers at IBM have made significant headway [16-18] in examining buyer and seller agent-driven markets through simulation, focusing on markets of information goods. Their analysis of agent-driven markets highlights some of the potential pitfalls of automated dynamic pricing, such as price wars. In their analysis, they introduced four different agent pricing strategies: game theoretic, derivative following, myopically optimal (dynamic programming), and Q-learning (reinforcement learning). Their specific algorithm for the derivative following strategy was adapted for finite markets and will be analyzed in the Learning Curve Simulator. Their work has provided a strong background for this investigation of successful strategy development.

Brooks et al. [19] also performed analysis of pricing agents in a simulated market environment and discussed the trade-offs between "exploitation" and "exploration" pricing techniques on the part of the seller. They conclude that when a pricing agent is interested in maximizing revenue over a longer period than the immediate purchase period, a simple learning algorithm works best for markets with high levels of uncertainty. While Brooks examines markets of information goods with no constraints on time or inventory, their use of a simulator to demonstrate the strength of different strategies provides a useful guideline for our analysis.

2.5 Our Approach

The McKinsey Quarterly [1], an industry publication on business strategy, recommends sellers pursue dynamic pricing on-line by running different pricing experiments. They state that by making small adjustments in price, sellers can discover the demand levels of their buyers. Despite the abundance of the theoretical models and pricing strategy formulations found in the academic literature, for the real-world seller, making predictions of buyer demand and implementing pricing strategies is far from straightforward, as highlighted by McKinsey's simplistic, yet practical, recommendation. By using a simulator prior to conducting pricing experiments, sellers could develop an intuitive understanding of the theoretical findings and use this knowledge to develop a more sophisticated strategy implementation. Therefore, we propose that sellers use a visual simulation tool, such as the Learning Curve Simulator, to study complex pricing strategies.

The Learning Curve Simulator, as a tool for sellers, addresses the complexities of on-line buyer behavior by providing a rich set of behavior parameters. To express the dispersion within a group of buyers, the simulator allows for the user to indicate a variance in reservation price for the chosen buyer/price distribution curve. Price sensitivity is expressed with a selection of the percentage of buyers whom comparison shop. Preference for a particular type of good or seller is expressed in an option to select a seller as "preferred." Although these parameters do not describe a complete or exact model of real-world markets, this is a more expressive set of variables than any previous set of simulation-based work for dynamic pricing analysis.

3. The Learning Curve Simulator

3.1 Simulator Interaction Design

Our presentation of the Learning Curve Simulator begins with an explanation of the graphical user interface. By presenting the interface first, we will demonstrate how a user typically interacts with the system as well as provide an explanation of each simulator parameter, with a justification of our design rational.

The simulator's graphical interface is a Java Swing application, which can run as either a client application or a web applet. It simulates a market based on user-supplied parameters describing the Market Scenario, the Buyer Behavior, and the Seller Strategies. The Learning Curve Simulator's interface is shown in Figures 1 through 4. These screen shots illustrate the steps a user takes to set up a model of her/his market and run simulations.

Figure 1 shows the initial screen of the simulator. At this screen the user selects from a defined scenario to pre-fill the following input screens or chooses to build a custom market scenario. The first three selections are based on the real-world markets of airline tickets, a grocer selling produce, and a ballpark selling tickets. The remaining selections are designed to illustrate certain strategic results.

3.1.1 The Simulation Cycle

At the moment the user hits the "Run Simulator" button, the simulator initializes the market by creating the sellers, calculating their initial prices and releasing them into the marketplace. (A later version of the simulator could allow for later entry of sellers.) Next, the simulator sequentially runs through each "day," or time period, of the market. Each day, a random number of buyers enter the market, based on a uniform distribution of buyer entrance over the entire market. These buyers stay in the market until either they have purchased a good or their lifetime has expired. On a single day, each buyer, in random sequence, searches through the available sellers, in random sequence, and compares the seller's price with the price it is willing to pay, referred to as its reservation price. In the case of a non-comparison shopper, if the seller's price is less, a transaction occurs and the buyer leaves the market. If the seller's price is greater than the buyer's reservation price, the buyer continues looking for another seller. Comparison shoppers perform the same search, but collect a list of seller prices each day and purchase from the seller offering the price furthest below its reservation price. The day ends when each buyer has either purchased or completed its search through all the sellers.

At the beginning of the next day, a new reservation price for each new and returning buyer is calculated based on the user-defined buyer behavior parameters. It is also at this point that each seller updates its price based on its designated pricing strategy. If the seller is using an adaptive pricing strategy, it examines different results from the market, such as how many goods it sold or how much revenue it earned in the previous day and uses this information to calculate a new price. And then the next day of transactions begins. In this manner, the market progresses until the last day, stopping early only if there are no remaining buyers or goods in the market.

3.1.2 Market Scenario

Figure 2 shows the Market Scenario inputs, the first series of simulator inputs. The Market Scenario is used to set the parameters of the finite market: the number of days, buyers, sellers, and goods. It also sets the market mechanism, buyer population segmentation, the costs of the market (cost of production and marginal cost per good), and the initial price offered by the sellers.

The number of days defines the number of periods the sellers can change their prices and the number of instances buyers can enter the marketplace. The number of goods per seller, as compared with the number of buyers, determines which parameter constrains the market: buyers or goods. The choice of constraining parameter affects the outcome of different strategies as will be shown in the analysis section.

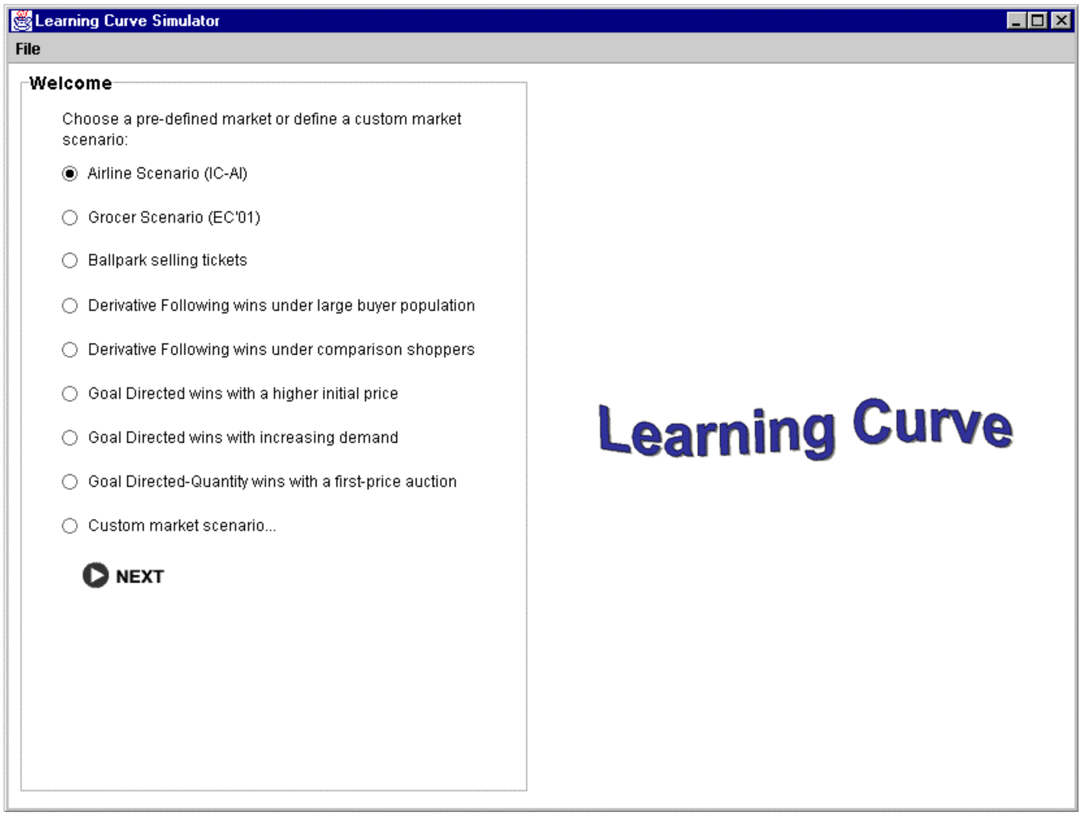

Figure 1: Learning Curve Simulator – Choose from a pre-defined scenario

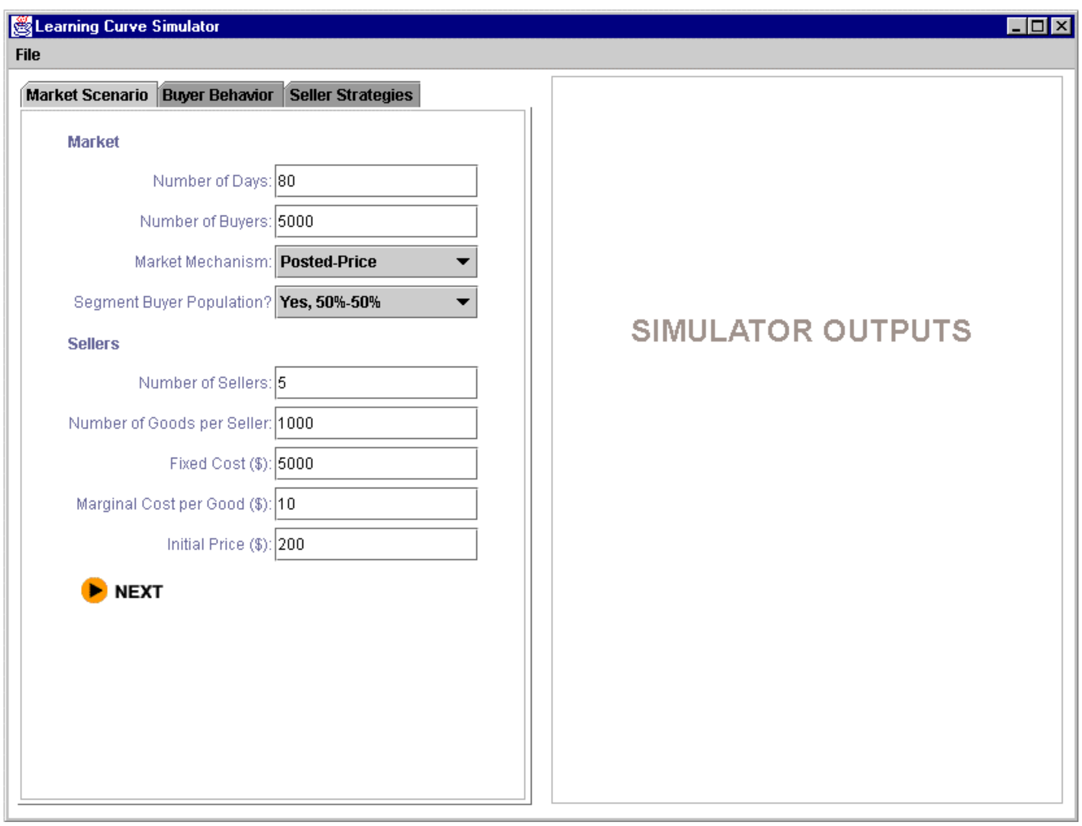

Figure 2: Learning Curve Simulator – Defining a market scenario

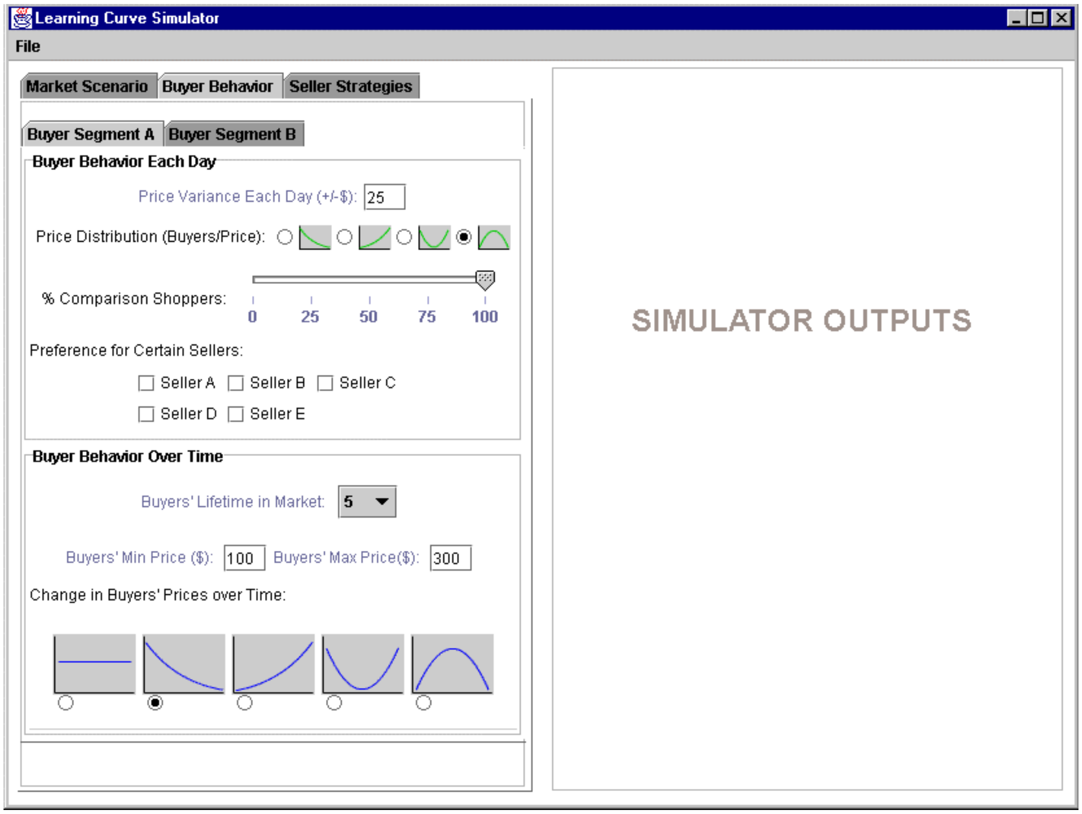

Figure 3: Learning Curve Simulator – Defining the behavior of buyer population.

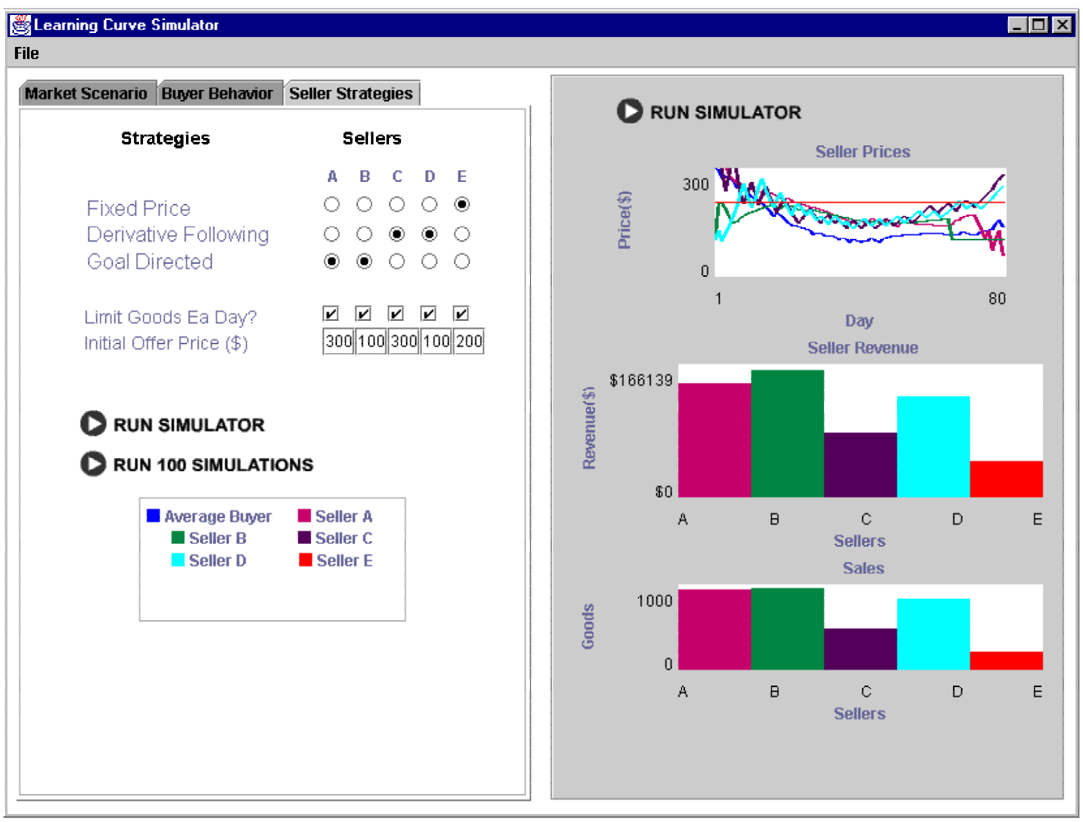

Figure 4: Learning Curve Simulator – Choosing the pricing strategies and viewing simulator results

The buyer population can be segmented into two groups, either into a 50/50 or 75/25 percentage split. By segmenting the buyers, the user can define two distinct types of behavior which will be joined into one population for the market simulation. The purpose of segmenting the population is to allow for users to express different sub-groups within customer populations. We limited the number of population segmentations and the choice of segmentation ratios to simplify the process of analyzing the effect of segmentation on strategy success.

The sellers' costs are defined as the sum of production costs and the marginal cost per good. Many finite markets, such as a ballpark, have a marginal cost of zero per good, so the major cost of the market is the initial cost of production. Although a simplistic assumption, the costs for each seller in the simulator are considered to be identical. Because it is assumed that margin costs are low (i.e. negligible) and because there is no distinction made between each seller's costs, the results of the simulation are reported in terms of revenue (price * units sold), not profit (revenue – costs).

The "initial price" input value is the price offered by each of the sellers on the first day of the market. This value can be adjusted on a per seller basis on the Seller Strategies screen.

3.1.3 Buyer Behavior

After outlining the Market Scenario, the user next defines the behavior of the buyers in the market, both in terms of their behavior on a per day basis and their behavior over time. These parameters are shown in the screenshots in Figure 3.

For each buyer segment, the dispersion among the buyers' reservation prices each day is defined by the variance and daily buyer/price distribution. The variance sets the range for the spread along the chosen distribution curve. The distribution curves model different types of demand curves: the common decreasing curve, an increasing curve which could apply to a luxury item where more buyers are willing to pay more for the good, a double peaked curve which applies to markets with two-tiers of buyers (such as price-insensitive business travelers and price-sensitive leisure travelers), and a mid-peaking curve which applies to a market in which there is a commonly understood average value for the item (known as a "common value" good in the auction literature [20]). While these are rough interpretations of existing markets, we offer these varied distributions to allow for more expressive descriptions of buyer behavior.

Although on-line shopping environments drastically decrease buyers' search costs, as discussed earlier, this has not resulted in buyers always purchasing the lowest priced product [11]. The simulator therefore does not assume the price sensitivity of buyers, but instead allows the user to select the percentage of buyers whom compare prices across sellers. Buyers who do compare prices check the prices of all sellers and buy from the seller with the highest percentage discount below their reservation price. When a buyer does not comparison shop, it incrementally checks sellers' prices until one is found below its reservation price, at which point it makes a purchase.

The final parameter determining the daily behavior of buyers is the designation of certain sellers in the market as "preferred." A preference for a seller can express real-world differentiation among products and sellers, due to higher quality, better product features, or brand loyalty. When a seller is selected as preferred, buyers are willing to pay 20% more for that seller's products. While this percentage mark-up is configurable in the back-end of the simulator, the interface was designed with a fixed percentage over the other sellers to simplify the user's interaction with the simulator.

Over the course of the market, the collective behavior of the buyers is defined by four variables: the lifetime, the minimum and maximum prices, and the valuation curve, each shown in the bottom half of the Buyer Behavior screen in Figure 3. The lifetime parameter indicates how "patient" the buyers are: how many days they are willing to wait in the market, continuously looking for the right price. If the buyer is still looking at the end of its lifetime, it leaves the market without purchasing. Indirectly, the lifetime of buyers determines the number of buyers in the market each day. On a single day, the number of buyers in the market depends on the number of pre-existing buyers who return to the market from yesterday and the number of entering buyers. Pre-existing buyers return if they were unable to purchase yesterday and their specified lifetime has not expired. The number of entering buyers is determined by a uniform distribution of all buyers of the course of the market. Thus, with some likelihood, each day has new buyers entering the market and, if demand was not satisfied the previous day, has pre-existing buyers returning.

The valuation curve choice determines how the buyers' average reservation prices, or valuation, changes over time. To express changes in market demand, the user can choose among flat, decreasing, increasing, mid-dipping, or mid-peaking demand curves. These curves model the buyers' response to limited supply, limited time, and or external market events. A future enhancement to the simulator could also include curves which model sudden and drastic changes to demand triggered by extreme events. To bound the values on the valuation curve, the user specifies a minimum and maximum reservation price. As a determined by these inputs, the buyers' valuation on a single day is a significant factor in how many sales a seller makes. As we shall see in the strategy analysis, the more successful sellers are the ones that can effectively follow the changes in the buyers' valuation over time.

3.1.4 Seller Strategies

The final step to setting up the market is to specify the pricing strategy each seller uses, shown in the left pane of the final screenshot, Figure 4. The simulator is designed to allow multiple strategies to work within the same market, so a user can compare how a strategy performs when competing with other strategies in the marketplace. For simplicity of comparison, a maximum of four strategies can be presented at one time in the simulator, and only three are shown in Figure 4.

The user can adjust each strategy by changing the initial price offered by the seller and by choosing to limit the number of goods sold in a single day for each seller. Changing the initial price effects the first day of sales, and of course, every day after in the case of a Fixed-Price strategy. Some of the strategies use this initial price in the pricing calculation, so this initial price also affects the behavior of these strategies over time. Sales can be limited each day to represent real-world market limitations to selling an entire inventory in a single day. When the user chooses to limit the sales, that seller can only sell three times the ratio of goods to days. In practice, limiting sales constricts the behavior of the sellers, producing less drastic changes in prices because there are less drastic discrepancies in sales between days. The factor three was chosen to by experimentation: lower factors were found to constrain the ability of the sellers to observe demand in the market, producing less profitable price changes, and higher factors were found to have no constraining effect.

3.1.5 Simulator Output

After the simulator runs, the results are presented in the right pane of the interface, as shown in Figure 4. These results summarize the market in terms of pricing, revenue, and sales. Additional output detailing each day and each transaction is saved to a tabbed-delimited file on the user's machine. If the user clicks 'Run 100 Simulations,' after 100 identical simulations run, an output file is created for each simulation, and a summary file is generated reporting the final revenue and sales of each seller per simulation.

The top chart in Figure 4 shows the pricing behavior of each seller on each day in relation to the average reservation price of the buyers. The next two charts report the revenues and sales of each seller. Revenue is the sum of the sale prices of each good sold. The total sales amount is the amount of inventory sold per seller. The success of the individual strategies is measured by the amount of revenue and sales and the pricing chart is used to understand how the sellers priced their goods and achieved their revenue and sales results. As shown in these results, it is straightforward to see which strategy earned the most revenue and sold more inventory, which makes the pricing chart the most interesting to watch between simulations. Future enhancements of the interface could also chart the number of buyers in the market each day and the profit per seller. This data is currently available in a tab-delimited text file generated at each simulator run.

3.2 Simulator Strategies

While the Learning Curve Simulator is designed to accommodate any dynamic pricing strategy, our initial analysis of dynamic pricing focuses on two strategies which we refer to as 'adaptive.' Adaptive strategies make no assumptions about the behavior or even existence of other market players, but instead observe changes in the seller's sales in relation to adjustments in price. In response to these basic observations, adaptive strategies implement increment price changes. Presented here are two such strategies, the Goal-Directed and Derivative-Following strategies. These strategies, as described below, were originally presented in [21].

3.2.1 Goal-Directed

The Goal-Directed (GD) strategy adjusts its price by attempting to reach the goal of selling the entire inventory by the last day of the market, and not before. By lowering prices when sales are low and raising prices when sales are high, this strategy paces its sales over the market, with the plan of selling to the highest paying buyers on each individual day. Equation 1 presents this strategy calculation.

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

$$expGoodsSold_i = i * \left(initialInventory \middle/ totalNumDays \right)$$

$$scale_i = totalNumDays \Big/ 2*\left(totalNumDays - i\right)$$

Equation 1: Goal-Directed Calculation

The strategy calculates the price for a particular day ($price_{i+1}$) by adjusting the price the seller offered at the beginning of the market ($price_0$) by the ratio of the number of goods sold so far in the market by the number of goods expected to have sold in the market by day i. This adjustment is scaled by a ratio of the days in the market to the number of days remaining in the market.
The GD calculation has been modified from our previous work [22] with the addition of this scaling factor ($scale_i$), which improves the strategy's ability to make price adjustments at the end of the market. By incorporating in knowledge of the progress through the market, the strategy now has the ability to make dramatic price changes during the last days, when sales are most important. As presented in [21,22] and as will be demonstrated below, the GD strategy performs best under high variance among the buyer population and when sales are less critical during the first days of the market.

3.2.2 Derivative-Following

The Derivative-Following (DF) strategy adjusts its price by looking at the amount of revenue earned on the previous day as a result of the previous day's price change. If yesterday's price change produced more revenue per good than the previous day, then the strategy makes a similar change in price. If the previous change produced less revenue per good, then the strategy makes an opposing price change. Revenue per good is equivalent to the sale price, except in the case when no goods are sold, so following this calculation, the seller will always sell at the highest price that generates sales.

$$price_{i+1} = price_i + (change_{i+1} * yestSuccess * yestChange)$$

$$change_{i+1} = price_i * \left(\beta + \left(\frac{totalNumDays - i}{(totalNumDays + i) * \alpha} \right) \right)$$

$$yestSuccess = \begin{cases} +1 & \text{if } revenue_i > revenue_{i-1} \\ -1 & \text{if } revenue_i < revenue_{i-1} \end{cases}$$

$$yestChange = \begin{cases} +1 & \text{if } change_i > 0 \\ -1 & \text{if } change_i < 0 \end{cases}$$

Equation 2: Derivative-Following Calculation

This strategy calculation, in Equation 2, computes the price for a particular day ($price_{i+1}$) by adjusting yesterday's price by a percentage change ($change_{i+1}$), scaled by a ratio based on the progress through the market. This scaling ratio takes into account the day of the market, much like the scaling factor in the GD strategy. The value of beta (β) is 0.05, chosen to ensure a minimum percentage change in price each day. The value of alpha (α) is 5.0, and counterbalances beta to ensure that the changes in price are not too large at the beginning of the market. Originally based on the strategy analyzed by Kephart, et al. in [17], this strategy has been adapted for a finite market by incorporating the scaling ratio. The change is either positive or negative depending on a combination of yesterday's change in price (yestChange) and yesterday's revenue (yestSuccess). As will be shown in the analysis section, the DF strategy performs best in the initial days of the market and reacts most strongly to competitive factors. When a market has a high percentage of comparison shoppers, DF sellers generate price wars, particularly when competing with other DF sellers.

3.3 Simulator Performance

The Learning Curve Simulator is built as a single-threaded Java 1.3 application, designed to run on a single Java Virtual Machine (JVM). The code of the simulator is designed around a three-tiered architecture. The lowest level tier is a general market framework consisting of Java interfaces for general market simulations. The mid-level tier is the specific implementation of the "Learning Curve" market. And the top-level tier defines the graphical user interface as described in the previous pages. This architecture design simplifies the process of incorporating new features into the Learning Curve Simulator, as well as allows for the possibility of creating very different simulators. If a new type of market simulator were to be implemented, the lowest level marketplace tier would serve as a starting point and the designer would only need to implement the existing Java interfaces within the general marketplace tier and additional classes as deemed necessary. Additionally, this multi-tiered structure allows for the option of running the interface on a separate JVM from the lower tiers of the simulator.

The speed of each simulation run depends on the number of buyers, and as more buyers are added to the simulator, the simulation time increases linearly. As one example data point, a simulation with 4000 buyers runs in approximately three seconds. The same simulation with 40,000 buyers runs in approximately 30 seconds. In our initial exploration of a simulation-based approach to dynamic pricing, simulation speed was not a primary goal. If future enhancements to the simulator require speed improvements, we will incorporate multi-threading.

4. Strategy Analysis

We present here an analysis of the Goal-Directed (GD) and Derivative-Following (DF) strategies under a finite set of changing buyer behavior parameters, presenting the conditions we found to be most influential over the success of each strategy. This analysis originally appeared in [21].

The following pages present an analysis of the two strategies, first under monopoly conditions (e.g. ballpark tickets) and next under competitive conditions (e.g. airline tickets). In every trial we present, the market has 100 days and each seller has 1000 goods. For each market scenario, we test the strategies under four different buyer valuation/time curves. Initially, we examine the success of the strategies under different populations of buyers (number of buyers and variance among buyers) and then we look at how competition affects the behavior of the strategies, under comparison-shopping and with preferences for certain sellers over others. Table 1 presents the values used in each trial simulation and the values shown in italics varied between trials.

Tables 2 through 8 detail the output of each simulation trial. For each of the pricing graphs shown, the vertical axis represents price – both the price offered by the seller and the price the average buyer was willing to pay – and the horizontal axis plots time across the market. On each graph, the vertical axis ranges from \$0 to \$350 and the horizontal axis ranges from 0 to 99 days. The darkest curve is always the average buyer reservation price and the lighter curves are the prices offered by the sellers. The revenue and sales results below each graph report the averaged results over 100 simulations ± one standard deviation. The purpose of running multiple simulation trials was to ensure the results represented a stable strategy outcome. For more precise revenue and sales results, we recommend running simulations with greater than 1000 trials.

Simulator Inputs:	Input Values
Market Scenario:	
Number of Days	100
Number of Buyers	***Four times as many as the number of goods (4000) or Equal to the number of goods (1000 or 2000)***
Number of Sellers	***1 (monopoly) or 2 (competition)***
Number of Goods	1000/seller
Market Mechanism	Posted-Price
Buyer Behavior:	
Daily Price Distribution	Normal distribution
Price Variance Per Day	***$0 or ±$50***
Percentage Comparison Shoppers	***0% or 100%***
Preference for Certain Sellers	***No seller preference or one seller preference***
Lifetime	***1 or 5 days***
Buyer Valuation over Time	***Increasing, decreasing, mid-peaking, and mid-dipping curves***
Minimum/Maximum Buyer Prices / Time	Minimum: $100 Maximum: $300
Seller Behavior:	
Seller Strategy	***GD or DF***
Initial Price	$200
Available Inventory per Day	3*(initial inventory/days)

Table 1: Simulator Input Values used in our Analysis
The parameter values in italics varied between different trial simulations.

4.1 Monopoly

To provide a baseline for analysis, Table 2 contains the results of eight simulations with one seller in the market, zero variance within the buyers' daily price distribution, and many, long-term buyers in the market. The graphs illustrate the characteristic behavior of the GD and DF strategies under each of the buyer valuation curves. In these trials, the standard deviations are zero because there is no randomness to the results when there are an unlimited number of buyers in the market with no variation between them.

Shown in the left column of Table 2, the GD strategy follows each buyer valuation curve very closely after a brief oscillation period. If the seller still has inventory to sell on the last days of the market, the GD strategy results in another period of drastic price oscillation in order to sell the remaining inventory. While the strategy succeeds in finding and following the demand curve, this is not always the best approach to the market. For example, in the case of constantly decreasing valuation over time, the GD seller paces its sales to include sales on the worst days of the market. Reflecting this poor behavior, this is the only case in which the GD strategy earned less revenue than the DF strategy.

The DF strategy also successfully follows each buyer valuation curve, but in a pattern of over- and under-shooting, shown in the right column of Table 2. When there is no variance in a large buyer population, the DF strategy sells its entire inventory at the halfway point through the market, and depending on the valuation curve, this is often not to the strategy's benefit. Only in the case of decreasing buyer valuation over time, where it is to the seller's advantage to sell during the first half of the market, did the DF strategy out perform the GD strategy.

By adding variance to the buyers' reservation prices, the strategies' ability to adapt dynamically can be demonstrated (Table 3). In the sample pricing graph shown, both strategies adjust their pricing curves to be higher than the average buyer price, thereby capturing the buyers who are willing to pay the highest prices each day. Again, the DF strategy prevails on the decreasing valuation curve because it does not sell goods at the last, i.e. worst, days of the market, unlike the GD strategy. Comparing these results to the initial case with no buyer price variance, both strategies produce significantly more revenue for the sellers under each valuation curve because they are able to raise their prices to meet the demand of the buyers willing to pay higher prices on a single day

Valuation / Time	Goal-Directed Strategy **After an initial oscillation period, the GD strategy follows the buyer valuation curve, pacing its sales through the entire market. Incorporated into the GD pricing calculation is an ability for the strategy to perform more drastic price adjustments at the beginning and ending of the market.**		Derivative-Following Strategy **The DF strategy follows the buyer valuation curve by over and under-shooting each period. When the market is saturated with buyers, this enables the seller to sell out of inventory half way through the market (as shown by the curve's disappearance).**	
Increasing	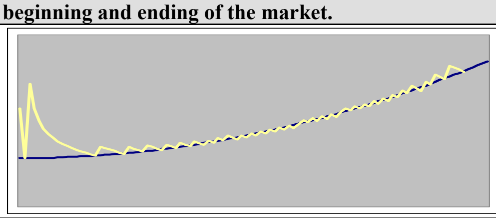		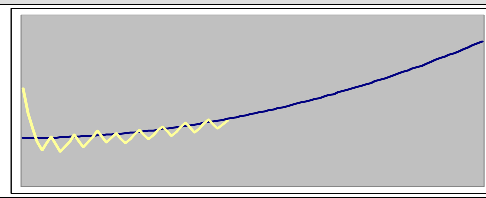	
	Revenue: $168,320 ± 0	**Sales: 1000 ± 0**	**Revenue: $101,910 ± 0**	**Sales: 1000 ± 0**
Decreasing	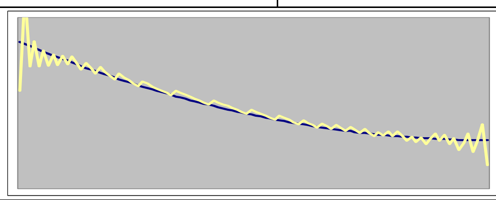		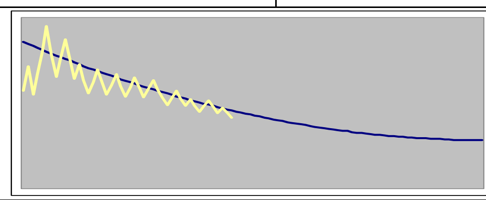	
	Revenue: $148,170 ± 0	**Sales: 990 ± 0**	**Revenue: $203,530 ± 0**	**Sales: 1000 ± 0**
Mid-Peaking	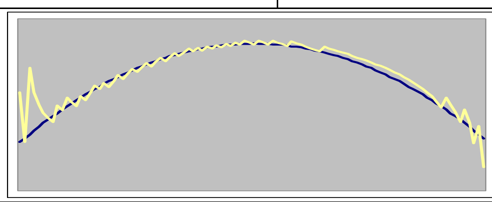		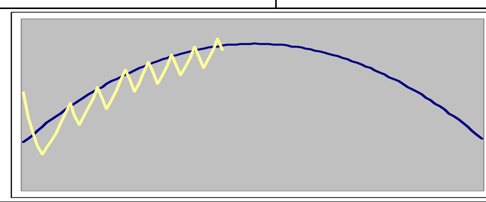	
	Revenue: $226,350 ± 0	**Sales: 1000 ± 0**	**Revenue: $196,760 ± 0**	**Sales: 1000 ± 0**
Mid-Dipping	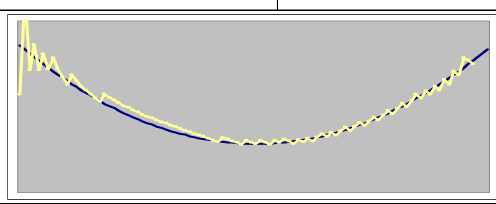		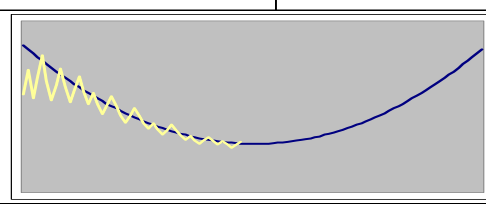	
	Revenue: $159,940 ± 0	**Sales: 1000 ± 0**	**Revenue: $155,710 ± 0**	**Sales: 1000 ± 0**

Table 2: Simulation results under Monopoly conditions with No Variance and Many, Long-term Buyers
The darkest curve is the average buyer reservation price on each day (valuation/time). The lighter curve is the price offered by the seller on a particular day.

Table 4 illustrates the strategies' behaviors when there are a limited number of buyers in the market. To express this limitation the simulator was run with the same number of buyers in the market as goods (1000) and the buyers each have a lifetime of one day, limiting the number of opportunities a seller has to make a sale. As the results show, under most curves, the GD strategy sells a significantly larger amount of inventory than the DF strategy, but this does not always lead to higher total revenue. The sample pricing graph demonstrates the behavior of the two strategies under the mid-peaking valuation curve. The GD strategy falls far below the buyer valuation curve when sales are slow, and near the end of the market drops the price down to $1 in an attempt to sell the remaining inventory. While it does manage to sell inventory, it does not do so at the best price. Conversely, the DF strategy follows the curve closely as it has during the previous trials and manages to maximize revenue per good over the course of the market. Shown in the mid-peak valuation curve, the DF strategy has achieved almost perfect matching of the valuation curve. Examining the revenue results, the DF strategy produces more revenue than the GD strategy except in the case of mid-peaking where the GD strategy managed to sell almost its entire inventory at a mediocre price, while the DF strategy only sold two-thirds of its inventory.

	Goal-Directed Strategy With High Variance		Derivative-Following Strategy With High Variance	
Sample Pricing Graph	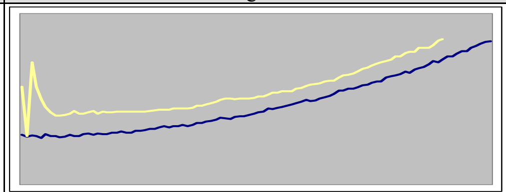		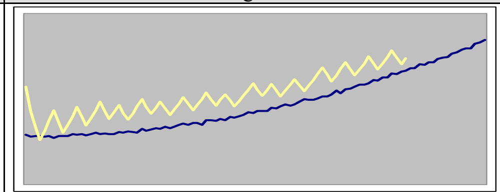	
Valuation Curve:	*Revenue:*	*Sales:*	*Revenue:*	*Sales:*
Increasing	**$199,680 ± 149**	**1000 ± 0**	**$149,036 ± 1089**	**1000 ± 0**
Decreasing	**$208,673 ± 847**	**994 ± 7**	**$228,689 ± 1078**	**1000 ± 0**
Mid-Peaking	**$275,052 ± 601**	**991 ± 2**	**$243,633 ± 1228**	**1000 ± 0**
Mid-Dipping	**$202,006 ± 198**	**1000 ± 0**	**$189,358 ± 739**	**1000 ± 0**

Table 3: Monopoly with High Variance and Many, Long-term Buyers
The darkest curve is the average buyer reservation price on each day (valuation/time). The lighter curve is the price offered by the seller on a particular day

	Goal-Directed Strategy With Few Buyers		Derivative-Following Strategy With Few Buyers	
Sample Pricing Graph	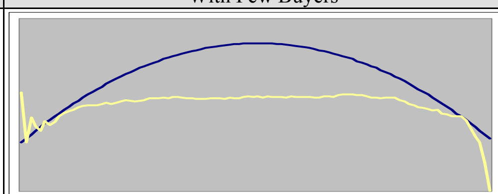		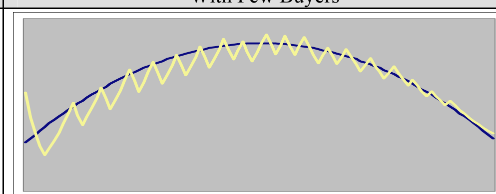	
Valuation Curve:	*Revenue:*	*Sales:*	*Revenue:*	*Sales:*
Increasing	**$79,176 ± 4598**	**814 ± 16**	**$123,112 ± 2690**	**790 ± 13**
Decreasing	**$107,441 ± 2642**	**811 ± 13**	**$111,492 ± 2759**	**710 ± 15**
Mid-Peaking	**$162,147 ± 5530**	**955 ± 7**	**$144,724 ± 3497**	**641 ±14**
Mid-Dipping	**$66,936 ± 2788**	**740 ± 16**	**$120,720 ± 2398**	**782 ± 12**

Table 4: Monopoly with No Variance and Few, Short-term Buyers
The darkest curve is the average buyer reservation price on each day (valuation/time). The lighter curve is the price offered by the seller on a particular day.

	Goal-Directed Strategy With High Variance & Few Buyers		Derivative-Following Strategy With High Variance & Few Buyers	
Sample Pricing Graph	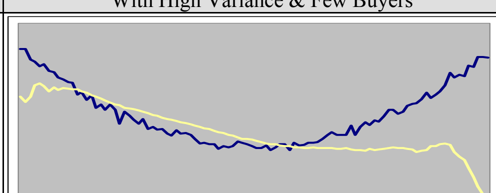		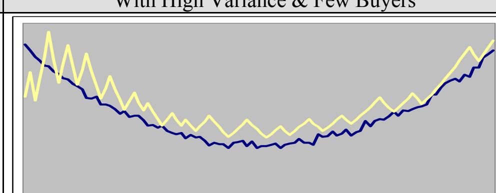	
Valuation Curve:	*Revenue:*	*Sales:*	*Revenue:*	*Sales:*
Increasing	**$141,958 ± 4619**	**999 ± 3**	**$189,363 ± 4114**	**977 ± 22**
Decreasing	**$127,302 ± 2107**	**889 ± 13**	**$67,333 ± 4180**	**328 ± 21**
Mid-Peaking	**$207,286 ± 2036**	**972 ± 5**	**$85,747 ± 5860**	**335 ± 24**
Mid-Dipping	**$102,601 ± 4409**	**907 ± 17**	**$75,253 ± 4669**	**372 ± 25**

Table 5: Monopoly with High Variance and Few, Short-term Buyers
The darkest curve is the average buyer reservation price on each day (valuation/time). The lighter curve is the price offered by the seller on a particular day.

When the market is severely limited in the number of buyers, the contrasting approaches of the strategies demonstrate strengths and weaknesses. The GD over compensates for the shortage of buyers and sacrifices daily revenue for daily sales. If it can manage to sell its entire inventory, then the total revenue can make up for the sacrifice. The DF strategy, by focusing on revenue per good, consistently makes sales on each day of the market at the highest possible price which can eliminate lower-paying buyers. When it is able to sell a large percentage of its inventory, the total resulting revenue is high.

When high variance is coupled with a small buyer population, the results are quite interesting. What is most notable about the results in Table 5 is that the DF strategy sells only a third of its goods under all valuation curves except the increasing curve. Examining the DF pricing curve, the pricing behavior looks very similar to the pricing under a higher variance (shown in Table 3), falling just above the average buyer curve. DF does not adjust for the limited number of buyers, and this lack of adjustment costs the seller the majority of its potential sales.

Contrast this result with the performance of the GD strategy. Referring to the sample pricing graph in Table 5, the GD strategy is able to sell at a relatively high price just before midway through the market because of the higher variance in buyer valuations. Then, when sales slip in the second half of the market, the GD strategy adjusts to a low price, and finally drastically drops the price to $1 at the end of the market. Both in sales and total revenue, the GD strategy performs extremely well. Although on average, it is selling at a lower price than the DF strategy, selling over 90% of its revenue produces significantly higher revenue.

4.2 Competition

In a competitive marketplace, the adaptive pricing strategies react to the other strategies in the marketplace in addition to the buyers' demand. While competitive markets frequently lead to comparison-shopping, we initially present here a market scenario in which none of the buyers compare prices across sellers or treat the sellers differently. Then we present the effects of comparison-shopping and seller-preference to demonstrate their unique affect on the strategies. As in the monopoly setting, each of the pricing graphs in the following tables are based on a 100 day simulation with the buyer valuation ranging from $100 to $300, depending on the valuation/time curve. In each of the competitive simulations, there were 2000 buyers, the same number of total goods in the marketplace.

Table 6 presents three different competitive pairings: Goal-Directed vs. Fixed-Price, Derivative-Following vs. Fixed-Price, and Goal-Directed vs. Derivative-Following. Logically, the success of a fixed-price seller depends on the fixed price it chooses. When used as a pricing policy, a "fixed-price strategy" should be optimized based on the predicted behavior of the market [8,14]. We are not examining the success of fixed-price strategies here, so we have simply chosen the fixed-price to be $200, the average valuation over time, across all the valuation curves. We present the fixed-price seller as a way of demonstrating the interplay between the adaptive and fixed-price policies.

In the left column of Table 6, when the Fixed-Price seller is able to sell goods (when its price is below the buyer valuation curve), the GD strategy stops adjusting its price and appears to mimic the Fixed-Price seller, particularly under the increasing and decreasing valuation curves. The reason the GD strategy stops changing its price is that when the Fixed-Price seller enters the market, the sales are split between the two sellers, and in this case with 2000 buyers (1000 per seller), the GD strategy sells the exact amount it aims to sell each day, making it unnecessary to change the price. If there were more or less buyers in the market, the GD strategy would result in a flat price curve at a higher or lower price point, respectively. Having a Fixed-Price seller in the market prevents the GD strategy from finding the highest price the buyers are willing to pay, yet in spite of this drawback, under every curve, the GD strategy produces a high amount of revenue and sells almost its entire inventory.

When the DF strategy is paired with a Fixed-Price seller, in the center column of Table 6, it has difficulty finding the buyer demand curve because of the low number of buyers and thus resorts to more frequent, higher oscillations in price. When the Fixed-Price seller is not making any sales, the DF strategy closely follows the buyer curve. This results in the DF strategy selling a much higher percentage of its goods, but at much lower prices than the Fixed-Priced seller. Under some curves this results in higher revenue for DF than for a Fixed-Price seller.

When DF and GD strategies are combined into the same marketplace, they do not respond to each other in a dramatic way. In fact, the individual strategies in the right column of Table 6 look much like when these strategies exist in a monopoly such as in Table 5. Each strategy is responding to the lack of buyers in the marketplace – the GD strategy starts to drop prices as sales drop off and the DF strategy keeps raising the price until the revenue ends and then dramatically lowers the price again

Valuation / Time	Goal-Directed Strategy vs. Fixed-Price Strategy		Derivative-Following Strategy vs. Fixed-Price Strategy		Goal-Directed Strategy vs. Derivative-Following Strategy	
Increasing	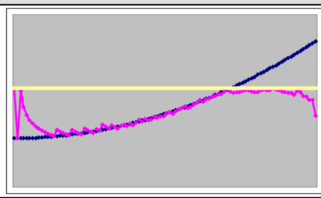		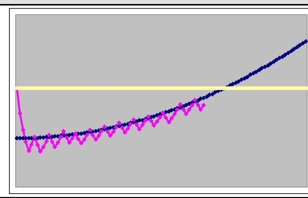		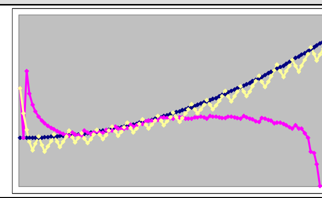	
	Goal-Directed **$157,307 ± 5596** **Sales: 990 ± 10**	*Fixed Price* **$58,912 ± 3273** **Sales: 295 ± 16**	*Deriv-Following* **$117,551 ± 1311** **Sales: 1000 ± 0**	*Fixed Price* **$116,258 ± 4365** **Sales: 581 ± 22**	*Goal-Directed* **$113,551 ± 8815** **Sales: 926 ± 19**	*Deriv-Following* **$139,024 ±3499** **Sales: 964 ± 22**
Decreasing	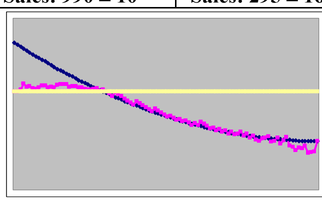		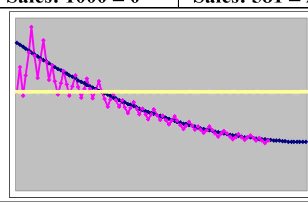		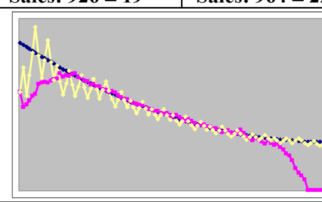	
	Goal-Directed **$134,204 ± 3069** **Sales: 976 ± 9**	*Fixed Price* **$62,086 ± 3813** **Sales: 310 ± 19**	*Deriv-Following* **$154,366 ± 1758** **Sales: 1000 ± 0**	*Fixed Price* **$75,882 ± 3265** **Sales: 379 ± 16**	*Goal-Directed* **$121,028 ± 3399** **Sales: 884 ± 17**	*Deriv-Following* **$145,205 ± 5725** **Sales: 922 ± 35**
Mid-Peaking	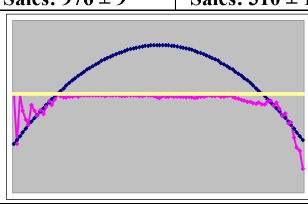		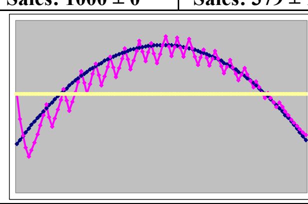		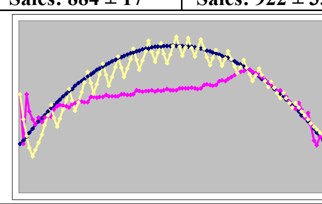	
	Goal-Directed **$178,410 ±4218** **Sales: 980 ± 6**	*Fixed Price* **$143,326 ± 4440** **Sales: 717 ± 22**	*Deriv-Following* **$160,391 ± 5307** **Sales: 765 ± 23**	*Fixed Price* **$183000 ± 4691** **Sales: 915 ± 23**	*Goal-Directed* **$190,593 ± 3397** **Sales: 980 ± 6**	*Deriv Following* **$164,719 ± 4025** **Sales: 752 ± 19**
Mid-Dipping	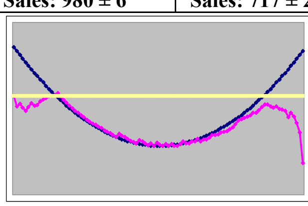		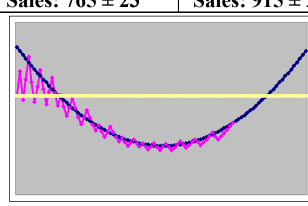		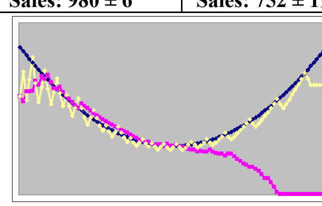	
	Goal-Directed **$144,565 ± 5263** **Sales: 987 ± 11**	*Fixed Price* **$58,794 ± 11** **Sales: 294 ± 16**	*Deriv-Following* **$128,381 ±1143** **Sales: 1000 ± 0**	*Fixed Price* **$92,126 ± 4029** **Sales: 461 ± 20**	*Goal-Directed* **$77,342 ± 3890** **Sales: 809 ± 20**	*Deriv-Following* **$146,845 ±2850** **Sales: 988 ± 17**

Table 6: Competition with No Variance and Few Buyers

The darkest curve is the average price that the buyers are willing to pay on each day (valuation/time). The lighter curves are the prices offered by the sellers on a particular day. In the right column, the medium colored curve is the GD strategy and the lightest curve is the DF strategy.

When a population of comparison shoppers is added to the marketplace, there is much more interaction between the two strategies. Table 7 illustrates the competitive effects of pairing two Goal-Directed strategies, two Derivative-Following strategies, and one Goal-Directed strategy with one Derivative-Following strategy when 100% of the buyer population compares the prices of the two sellers and purchases from the lowest priced seller. When we ran this trial with 75%, 50% and 25% comparison shoppers, the results linearly approached those with no comparison-shopping.
Across the results, the amount of revenue earned by each seller has been dramatically reduced from the same trials with no comparison-shopping (compare the right columns of Tables 6 and 7). Examining the results of the two GD strategies, they behave much as they did in a monopoly setting with limited buyers (Table 4), except they do not respond to the high variance in the buyer population. The center column shows the two DF strategies, and as shown most dramatically by the sample pricing graph, when they are paired together, they produce a price war. When one GD competes with one DF, there is a modified price war, where prices do not drop as dramatically, but are still forced down by the DF strategy. The DF strategy sells approximately the same amount of inventory as GD, yet earns more revenue than the GD strategy under all valuation curves and increases its revenue as compared to the DF-DF competition. This occurs because the DF strategy does not limit the amount of inventory it sells at the beginning of the market when prices are higher, while the GD strategy spreads out its sales, including selling on the last days of the price war when prices approach zero.

	GD vs. GD With Comparison-Shopping		DF vs. DF With Comparison-Shopping		GD vs. DF With Comparison-Shopping	
Valuation Curve:	*GD Revenue:*	*GD Revenue:*	*DF Revenue:*	*DF Revenue:*	*GD Revenue:*	*DF Revenue:*
Increasing	**\$57,881 ± 2220**	**\$57,881 ± 2220**	**\$40,532 ± 8211**	**\$40,532 ± 8211**	**\$35,639 ± 2831**	**\$58,713 ± 1856**
Decreasing	**\$87,058 ± 1875**	**\$87,058 ± 1875**	**\$86,512 ± 6549**	**\$86,512 ± 6549**	**\$71,826 ± 3564**	**\$117,151±4074**
Mid-Peaking	**\$143,472 ± 2837**	**\$143,472 ± 2837**	**\$53,273 ± 28,092**	**\$53,273 ± 28,092**	**\$57,763 ± 4968**	**\$96,786 ± 3833**
Mid-Dipping	**\$63,595 ± 1664**	**\$63,595 ± 1664**	**\$63,595 ± 1664**	**\$63,595 ± 1664**	**\$50,765 ± 3939**	**\$80,820 ± 3158**
Sample Pricing Graph	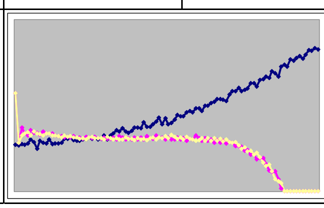		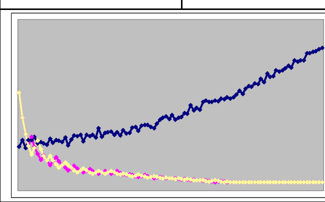		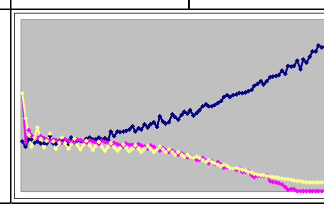	

Table 7: Competition under Comparison Shopping and High Variance
The darkest curve is the average price that the buyers are willing to pay on each day (valuation/time). The lighter curves are the prices offered by the sellers on a particular day. In the right column, the medium colored curve is the GD strategy and the lightest curve is the DF strategy.

	Goal-Directed vs. Derivative-Following With Preference for GD		Goal-Directed vs. Derivative-Following With Preference for DF	
Valuation Curve:	*GD Revenue:*	*DF Revenue:*	*GD Revenue:*	*DF Revenue:*
Mid-Peaking	**\$208,822 ± 5102**	**\$157,476 ± 4674**	**\$190,360 ± 4126**	**\$212,647 ± 4422**
Sample Pricing Graph	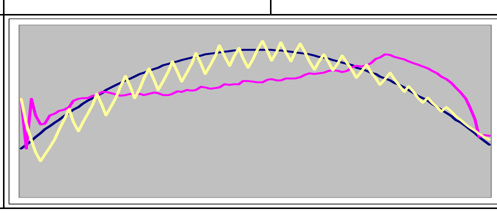		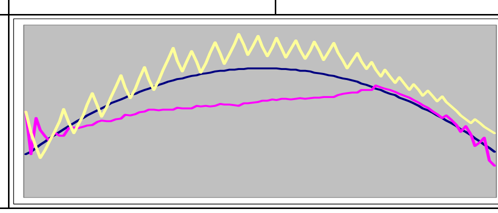	

Table 8: Competition under a Buyer Preference for Different Sellers
The darkest curve is the average price that the buyers are willing to pay on each day (valuation/time). The medium colored curve is the GD strategy and the lightest curve is the DF strategy.

The last strategy comparison we present shows the effects of buyers having a preference for one seller's product over another. When buyers have a preference for a certain seller, the population of buyers considers that seller's product to be more valuable, perhaps because of brand, quality, or reputation. In our simulator, this is modeled by boosting up the reservation price a buyer has for that seller by a fixed percentage, in this case 20%. Table 8 shows the competition between the GD and the DF when there is a preference for one of the sellers. What we observe is that both strategies are able to charge higher prices at certain points in the market, but the GD strategy is forced to lower its price during the middle portion of the market to ensure it made enough sales. Under both trials, the sellers sold approximately 70-80% of their inventory. While the preferred seller earns more revenue under the different trials, the earnings spread between the two sellers is not nearly as large when there is a preference for the DF seller.

4.3 Strategy Analysis Conclusions

While the Goal-Directed and Derivative-Following strategies are computationally straightforward, they are surprisingly robust under extremely different market conditions. Under every case we presented, excluding the situation of 100% comparison-shopping, the strategies managed to adjust prices in the direction of learning the changing demand in the marketplace, without knowing the true buyer demand, competitors' prices, even the number of other agents in the marketplace. We have summarized below the behavior if these two strategies as a way of pointing towards general guidelines for choosing and designing adaptive pricing strategies:

- The Goal-Directed strategy consistently sells all or the majority of its inventory given any combination of buyer behaviors and competition, at the expense of drastically over- and under-shooting the buyer valuation curve early and late in the market. Thus the Goal-Directed strategy is best for slower moving markets where the first and last days do not require fine-tuned price adjustments. It is also ideal for markets in which inventory liquidation is essential.
- The Derivative-Following strategy consistently sells at the highest price possible on any single day. When there is a relative peak in demand during the first days of the market and an abundance of buyers, Derivative-Following performs very well. If buyer demand peaks at some later time, Derivative-Following does not space out its sales so as to ensure that it capitalizes on that later peak. Thus the Derivative-Following strategy excels in a market with an abundance of buyers and a peak in demand early in the market.
- In a monopoly, the shape of the valuation/time curve has an enormous effect on the success of an individual strategy. Variance among buyer reservation prices and few numbers of buyers requires adaptive strategies to be more agile. When designing an optimal strategy for a monopoly setting, knowledge of the typical valuation curve and the buyer population should be incorporated into the pricing algorithm.
- If buyers are extremely price sensitive (100% comparison-shoppers), adaptive strategies can easily breakdown into price wars. In particular, the Derivative-Following strategy generates a price war between itself and other adaptive strategies.
- When there is product or seller differentiation (a willingness to pay more for certain seller's products), a carefully designed adaptive strategy can narrow or widen the discrepancy between the sellers' earnings.
- As dynamic pricing is deployed in real-world markets, it is important to understand the interplay of different pricing strategies. [23] compared two simple pricing strategies, price matching and price cutting, and combined them into one simulated market setting, demonstrating that both strategies were weakened in a mixed strategy marketplace. Our strategies, while neither price matching or cutting, produced mixed results. When there was no comparison-shopping, the DF and GD strategies did not significantly affect each other's behavior or success because these algorithms are not tied to competitor prices. But in the market with comparison-shoppers, the two strategies began to affect each other. The presence of a DF strategy hurt the success of the GD strategy while the presence of the GD strategy improved the success of the DF strategy over when it competed with another DF strategy.

The Goal-Directed and Derivative-Following strategies demonstrate just two approaches to dynamic pricing within finite markets. Additionally, the analysis conclusions we have summarized above are a small set of possible simulation observations. The richness of the simulator inputs enable modeling of more realistic markets than we have shown, and in these non-tractable market situations, our strategies can be just as easily observed and evaluated. We hope through our highlighting of the strategies' key characteristics and our demonstration of strategy analysis, we have provided a roadmap for how a simulation-based approach to market analysis and strategy development could be implemented by real-world sellers.

5. Future Directions

Returning to our original example of a ballpark selling baseball tickets, today when scalpers sell tickets outside the park they are reselling tickets purchased through the park's fixed-price policies. Scalpers adjust their prices on as much as a per ticket basis, responding to changes in the time left before the game, weather changes, and the size of crowd heading from the parking lot toward the park. The mere existence of the scalped ticket market is evidence that dynamic pricing is profitable. So why aren't ballparks adjusting their prices? We believe the biggest challenge to changing prices is making instantaneous strategic changes in price. We propose a ballpark, or similar seller, use a market simulator to model their market and analyze which pricing strategy is best for their marketplace. Through the use of a simulator, a seller can be informed and prepared before implementing an automatic pricing strategy in a market.

One of the goals of this research is to develop a tool that a ballpark, or similar seller in a finite market, could use to explore and understand the conditions for which an adaptive or other dynamic pricing strategy works. By working with the Learning Curve Simulator, a ballpark could model its market and test different strategies, to determine an optimal pricing strategy for its specific market conditions. Once an optimal strategy had been determined, a ballpark could take its algorithm and further customize it for the real-world market and eventually deploy the strategy to perform automated price changes in the baseball ticket market.

There are several open issues in the deployment of dynamic pricing, for which the Learning Curve Simulator can contribute towards solving. The following sections highlight some of these issues in electronic markets and how a simulation-based approach can facilitate their solution.

5.1 Strategy Development

The adaptive pricing strategies implemented in this body of work illustrate one type of approach to designing pricing strategies. There are many potential approaches to strategy development and a simulator can serve as a platform for testing such strategies.

An effective technique for optimal pricing is dynamic programming [24] which, like revenue management, makes assumptions about the marketplace to forecast and make optimal decisions, taking into account time and inventory constraints. By considering the problem of pricing in a market to be a multi-armed bandit allocation problem [25] and simplifying the strategy decision to a finite number of decision variables, a strategy could be developed and tested in the Learning Curve Simulator to find an optimal pricing solution for each market scenario. Although, as discussed earlier, a drawback to this approach is the number of required market behavior assumptions, such as the shape of the buyer valuation/time curve. Another drawback is that to deploy an optimal solution, the calculation is often times too computationally intense for a real-world setting [5]. But these drawbacks do not preclude the benefit of understanding how dynamic programming strategies perform in a market and the Learning Curve Simulator can provide the mechanism to do that.

5.2 Buyer Strategies

In addition to the behaviors implemented in the simulator, what can not be ignored is the intelligence of buyers to adapt their buying behavior based on observing prices change. An exciting direction to take this body of work is to incorporate buying strategies in the simulator to evaluate their impact on pricing strategies. Just as leisure travelers purposefully purchase airline tickets more than twenty-one days in advance to receive a discounted fare, when sellers implement dynamic pricing into new markets, discount-seeking buyers will decipher pricing rules to locate a lower price. One can easily imagine buyers employing strategic agents which are much more sophisticated than today's pricebots, able to aggregate and time purchases such that customers receive comprehensive discounts. The effect of this type of buyer behavior on a seller's success could be evaluated in a simulator that accurately modeled the different ways in which buyers could respond to dynamic pricing.

5.3 Market Types

This article focused on finite markets with posted-prices, yet this constraint on the analysis does not limit the impact of the simulator as a tool for understanding alternative markets and dynamic pricing strategies. The lessons learned from finite markets can be extended to markets with non-perishable goods, such as the automotive industry [7,15] and to alternative market mechanisms such as auctions. Pricing strategies can be designed to have knowledge of production and distribution decisions and how changing prices can improve the entire supply chain.

Auctions have become an extremely popular market mechanism for selling products, and while tempting to enter auction markets to sell goods, sellers should proceed with caution when deploying dynamic pricing strategies in these markets. Consumers behave differently in markets in which they name their own prices and this affects the sale price of the item [26]. Before developing a pricing strategy for an auction, a seller should gather an understanding of how their customers will behave within the chosen auction type. The Learning Curve Simulator could serve as a platform for modeling this buyer behavior and studying the effects of this behavior on different auction pricing strategies.

5.4 Conclusion

Dynamic pricing will likely become a common competitive maneuver in the near future and because of this sellers need to be equipped with an understanding of how different pricing strategies will play out in their marketplaces. This article proposes a way of approaching the problem of pricing strategy implementation. We believe dynamic pricing is a powerful method for increasing revenue in an electronic marketplace, but it is a non-trivial problem to implement effective pricing strategies. In the business strategy magazine Darwin Online, the difficulty and risks of dynamic pricing are summarized with a warning to sellers: "poorly implemented pricing schemes create the potential for competitive price wars and lowered profitability for all" [27]. The Learning Curve Simulator is designed to alleviate these risks of dynamic pricing by providing a mechanism and approach for understanding dynamic markets and analyzing pricing strategies.

Acknowledgements

We would like to acknowledge the valuable feedback and assistance provided by the anonymous reviewers, the members of the MIT Media Lab's Software Agents Group, Dan Ariely, and Katarzyna Dratewska. This work was supported in part by the Digital Life Consortium and the E-Markets SIG.

References

[1] W. L. Baker, E. Lin, M. V. Marn and C. C. Zawada. 'Getting Prices Right on the Web,' *The McKinsey Quarterly* No. 2, 2001.

[2] M. Smith, J. Bailey and E. Brynjolfsson. 'Understanding Digital Markets: Review and Assessment.' In E. Brynjolfsson and B. Kahin (eds.), *Understanding the Digital Economy*. Cambridge, MA: MIT Press. 2000.

[3] P. Keskinocak and S. Tayur. 'Quantitative Analysis for Internet-Enabled Supply Chains,' *Interfaces, Special issue on Operations Research in the e-Business Era* Vol. 31, March/April, 2001.

[4] E. A. Boyd. 'Airline Alliance Revenue Management,' *OR/MS Today* October, 1998.

[5] J. I. McGill and G. van Ryzin. 'Revenue Management: Research Overview and Prospects,' *Transportation Science* Vol. 33, No. 2, 1999, pp. 233-256.

[6] B. C. Smith, 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* Vol. 31, March/April, 2001.

[7] S. Biller, L. M. A. Chan, D. Simchi-Levi and J. Swann. 'Dynamic Pricing and the Direct-to-Customer Model in the Automotive Industry.' *GM Research & Development Center*, 2000.

[8] G. Gallego and G. van Ryzin. 'A Multiproduct Dynamic Pricing Problem and Its Applications to Network Yield Management,' *Operations Research* Vol. 45, No. 1, 1997, pp. 24-41.

[9] K. Clay, R. Krishnan and E. Wolff. 'Pricing Strategies on the Web: Evidence from the Online Book Industry.' *Proceedings of the 2nd ACM Conference on Electronic Commerce (EC '00)*, Minneapolis, Minnesota, 2000, pp. 44-55.

[10] Z. Carmon and D. Ariely. 'Focusing On The Forgone: How Value Can Appear So Different To Buyers And Sellers,' *Journal of Consumer Research* Vol. 27, No. 3, 2000, pp. 360-370.

[11] E. Brynjolfsson and M. Smith. 'The Great Equalizer? Consumer Choice at Internet Shopbots (Working Paper).' *Massachusetts Institute of Technology*, 2000.

[12] R. Guttman and P. Maes. 'Agent-mediated Integrative Negotiation for Retail Electronic Commerce.' *Proceedings of the Workshop on Agent Mediated Electronic Trading (AMET'98)*, Minneapolis, MN, 1998.

[13] J. Morris and P. Maes. 'Sardine: An Agent-facilitated Airline Ticket Bidding System.' *Proceedings of the Fourth International Conference on Autonomous Agents (Agents 2000)*, Barcelona, Catalonia, Spain, 2000.

[14] G. Gallego and G. van Ryzin. 'Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizons,' *Management Science* Vol. 40, No. 8, 1994, pp. 999-1020.

[15] L. M. A. Chan, D. Simchi-Levi and J. Swann. 'Flexible Pricing Strategies to Improve Supply Chain Performance (Working Paper).' *Northwestern University*, 2000.

[16] J. O. Kephart and A. Greenwald. 'Shopbot Economics.' *Proceedings of the Third International Conference on Autonomous Agents (Agents'99)*, Seattle, WA, 1999, pp. 378-379.

[17] J. Kephart, J. Hanson and A. Greenwald. 'Dynamic Pricing by Software Agents,' *Computer Networks* Vol. 32, No. 6, 2000, pp. 731-752.

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

[19] C. H. Brooks, 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)*, Denver, CO, 1999.

[20] A. Ockenfels and A. E. Roth. 'Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction,' *American Economic Review (forthcoming)* 2002.

[21] J. M. DiMicco, A. Greenwald and P. Maes. 'Dynamic Pricing Strategies under a Finite Time Horizon.' *Proceedings of the ACM Conference on Electronic Commerce (EC '01)*, Tampa, FL, 2001.

[22] J. Morris, P. Maes and A. Greenwald. 'Learning Curve: Analysis of an Agent Pricing Strategy Under Varying Conditions.' *Proceedings of the International Conference on Artificial Intelligence (IC-AI)*, Las Vegas, NV, 2001.

[23] C. A. Deck and B. J. Wilson. 'Interactions of Automated Pricing Algorithms: An Experimental Investigation.' *Proceedings of the 2nd ACM Conference on Electronic Commerce (EC '00)*, Minneapolis, MN, 2000, pp. 77-85.

[24] R. E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, 1957.

[25] B. Leloup and L. Deveaux. 'Dynamic Pricing on the Internet: Theory and Simulations,' *Journal of Electronic Commerce Research, Special Issue on Electronic Market Design* Vol. 1, No. 3, 2001, pp. 265-276.

[26] D. Ariely and I. Simonson. 'The Psychology of On-line Auctions (Working Paper).' *Massachusetts Institute of Technology*, 2001.

[27] S. Kalin. 'How Low Can You Go?,' *Darwin Online* April, 2001.