Dynamic Pricing Strategies under a Finite Time Horizon

Joan Morris DiMicco
MIT Media Laboratory
20 Ames St
Cambridge, MA 02139 USA
+1 617 252 1603
joanie@media.mit.edu

Amy Greenwald
Brown University
Box 1910
Providence, RI 02912 USA
+1 401 863 7678
amygreen@cs.brown.edu

Pattie Maes
MIT Media Laboratory
20 Ames St
Cambridge, MA 02139 USA
+1 617 252 1603
pattie@media.mit.edu

ABSTRACT
In the near future, dynamic pricing will be a common competitive maneuver. In this age of digital markets, sellers in electronic marketplaces can implement automated and frequent adjustments to prices and can easily imagine how this will increase their revenue by selling to buyers "at the right time, at the right price." But at present, most sellers do not have an adequate understanding of the performance of dynamic pricing algorithms in their marketplaces. This paper addresses this concern by analyzing the performance of two adaptive pricing algorithms. We study the behavior of these algorithms within the Learning Curve Simulator, a platform for analyzing dynamic pricing strategies in finite markets assuming various buyer behaviors. The goals of our research are twofold: (i) to explore the use of simulation as a tool to aid in the development of dynamic pricing strategies; and (ii) to explicitly identify the market conditions under which our example strategies, Goal-Directed and Derivative-Following, are successful.

General Terms
Algorithms, Measurement, Economics, Experimentation.

Keywords
Agent simulation, dynamic pricing, electronic markets, buyer behavior, pricing strategies.

1. INTRODUCTION
Today, when a ballpark sells baseball tickets, it 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 additional unpredictable factors such as the weather. In a best-case scenario, a park sells all 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 fans willing to pay more. Nonetheless, today ballparks leave the practice of dynamic pricing to scalpers.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
EC'01, October 14-17, 2001, Tampa Florida, USA.
Copyright 2000 ACM 1-58113-387-1/01/00010…$5.00.

Cost has been perhaps the greatest factor precluding the widespread use of dynamic pricing by ballparks, theaters, and retail shops. 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 [11]. Moreover, in markets under a finite time horizon, such as ballpark and theater tickets, a clear benefit to changing prices over time is the ability to clear inventory before the market ends. Thus, it seems likely that in the near future, dynamic pricing will become a common competitive maneuver.

A remaining obstacle that still precludes widespread dynamic pricing is the lack of understanding of the inner-workings of dynamic pricing models. Now that sellers can easily implement automated algorithms that make frequent adjustments to price, how should they do so? What are the most effective dynamic pricing strategies, and how should prices change with changing market conditions? 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 ballpark could model the characteristics of its market and the behavior of the team's fans, to develop a pricing strategy that would capture more revenue than an existing fixed-price policy.

To illustrate our approach, we analyze two pricing strategies within the Learning Curve Simulator, our platform for running dynamic pricing algorithms in simulated markets. Our investigation focuses on adjusting prices over time in what we call "finite markets" – markets with a finite time horizon, seller inventory, and buyer population. In this investigation, we are not exploring price discrimination, the adjustment of prices between individual buyers. Our strategies would apply to markets such as event tickets, airlines, hotels, perishable goods, and seasonal retail. The adaptive pricing strategies we present, Goal-Directed and Derivative-Following, demonstrate two approaches to dynamic pricing within finite markets. We hope these strategies will lay the groundwork for designing more complex strategies designed to be deployed in a real-world market.

2. LEARNING CURVE SIMULATOR
The Learning Curve Simulator models market scenarios, such as a ballpark selling game tickets, through a rich set of inputs describing the market and the behavior of buyers over time. The simulator, a Java 1.3 application, accepts three categories of inputs to describe the marketplace: the Market Scenario,

Simulator Inputs:	Description
Market Scenario:	
Number of Days	Number of periods in the market. Each seller can change its price at the end of a day.
Number of Buyers	The size of the buyer population over the entire market.
Number of Sellers	Number of sellers.
Number of Goods	Initial inventory for each seller.
Market Mechanism	Posted-Price or First-Price Auction. See [7] for discussion of the auction implementation.
Buyer Behavior:	
Daily Price Distribution	The demand distribution of buyers on a single day. Available choices are normal distribution, positive slope, negative slope, or segmented into a high and low grouping.
Price Variance Per Day	The buyers' reservation prices vary ± the variance in a single day. The variance determines the range for the daily price distribution.
Percentage Comparison Shoppers	The percentage of the buyer population (0-100%) who compare each seller's offer price and purchase from the seller with the greatest % discount below its reservation price for that seller.
Preference for Certain Sellers	The entire buyer population can have a preference for one or more of the sellers, which is represented by a higher reservation price for that individual seller. This is a method for expressing product and seller differentiation.
Lifetime	Number of days a single buyer will be in market, actively looking for seller. Regardless of lifetime, once a buyer purchases, it leaves the market.
Buyer Valuation over Time	Over the course of the market, the buyers' demand curve will change, and the valuation/time curve expresses how the demand will change over time. The shape of the curve can be flat, increasing, decreasing, mid-peaking, or mid-dipping over time.
Minimum/Maximum Buyer Prices over Time	The range of prices for the buyer valuation curve. These values are the minimum and maximum reservation prices over the market.
Seller Behavior:	
Seller Strategies	The different pricing strategies sellers use in the market, either Goal-Directed or Derivative-Following. See [7] for a discussion of other implemented strategies.
Initial Prices	The different prices sellers offer on the first day of the market, before adjusting price through the chosen strategy.
Available Inventory per Day	Amount of inventory a seller can sell in one day. This can be limited to represent shelving costs and to prevent 100% inventory sell-off in a single day.

Table 1: Learning Curve Simulator Inputs

the Buyer Behavior, and the Seller Behavior. Table 1 presents the simulator inputs in each of these categories. The simulator creates populations of buyers and sellers that search for each other on each day of the market. When a buyer and seller match on price, they perform a transaction and the buyer leaves the market. At the end of each day, each seller computes its price for the next day. At the end of the market, the success of a pricing strategy is determined by the total revenue earned, a function of the amount of inventory sold by a seller. We treat revenue, not profit, as the success metric because in finite markets the incremental costs of inventory production are often minimal. More information on the simulator's behavior and modeling can be found in [7-9].

3. ADAPTIVE STRATEGIES

The Goal-Directed and Derivative-Following strategies perform dynamic pricing by making incremental, exploratory adjustments to price each day in an attempt to learn the demand in the marketplace. The Learning Curve Simulator is designed to accommodate any dynamic pricing strategy, so these two represent our initial strategy implementations as methods for understanding the dynamics of finite markets. The key characteristics of these strategies are their relative computational simplicity and the lack of assumptions made about the behavior of competitors or buyers.

The *Goal-Directed* (GD) strategy adjusts its price by attempting to reach the goal of selling its 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 goal of selling to the highest paying buyers on each individual day. Equation 1 presents this strategy calculation.

The GD calculation has been modified from our previous work [9] with the addition of a scaling factor ($scale_i$ in Equation 1). This scaling improves the strategy's ability to make price adjustments at the end of the market. Previously, the strategy adjusted its price dramatically in the first third of the market but was unable to make large adjustments to price in the last days 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 [9] 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.

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, 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 equals the sale price, except in the case when no goods are sold, so following this calculation, the seller sells at the highest price each day that still generates sales.

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

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

$$scale_i = daysInMarket / 2*(daysInMarket - i)$$

Equation 1: Goal-Directed Calculation

$$price_{i+1} = price_i \pm change_{i+1}$$

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

Equation 2: Derivative-Following Calculation

This strategy calculation, shown in Equation 2, is an adjustment of the strategy analyzed by Kephart, et al in [5]. We tailored the DF's performance for a finite market taking into account the day of the market. Much like the scaling factor in the GD strategy, instead of adjusting the price each day by a fixed percentage (β in Equation 2), the change is scaled by a ratio based on the progress through the market. As will be shown in the analysis sections, 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.

4. STRATEGY ANALYSIS

We present an analysis of the GD and DF strategies under a small set of changing buyer behavior parameters, presenting the conditions that we found to be most influential over the success of each strategy in finite markets. Based on the input parameters detailed in Table 1, Table 2 presents the values used in each trial simulation. The values shown in italics varied between trials.

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.

For each of the pricing graphs in Tables 3-9, the vertical axis represents price – both the price offered by the seller and the price the average buyer is 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.

Simulator Inputs:	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 2: Simulator Input Values used in our Analysis

The parameter values in italics varied between different trial simulations.

4.1 Monopoly: One Seller in the Market

To provide a baseline for analysis, Table 3 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 and there is no variation between these buyers.

Shown in the left column of Table 3, 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.

Valuation / Time	Goal-Directed Strategy		Derivative-Following 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.		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	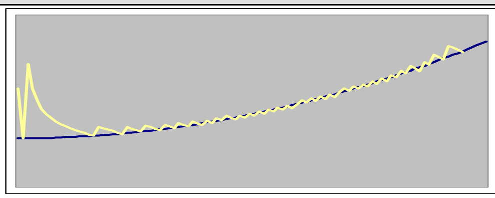		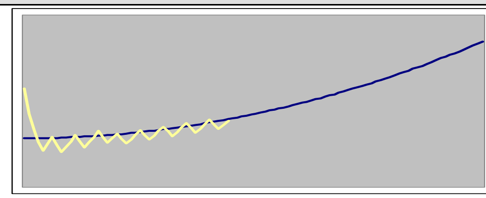	
	Revenue: $168,320 ± 0	Sales: 1000 ± 0	Revenue: $101,910 ± 0	Sales: 1000 ± 0
Decreasing	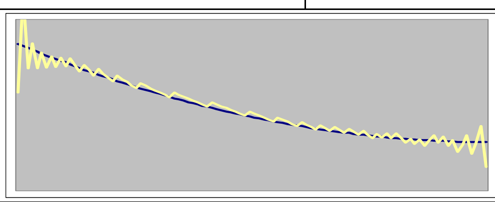		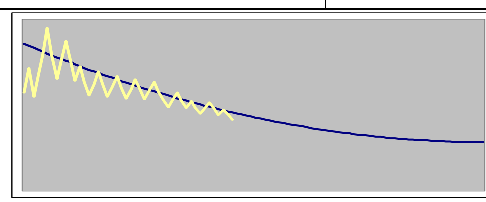	
	Revenue: $148,170 ± 0	Sales: 990 ± 0	Revenue: $203,530 ± 0	Sales: 1000 ± 0
Mid-Peaking	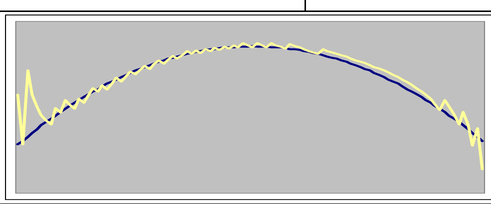		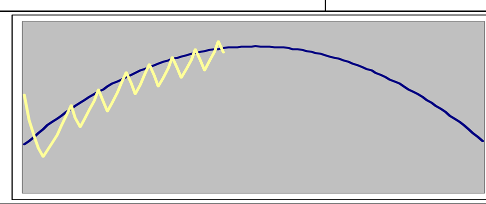	
	Revenue: $226,350 ± 0	Sales: 1000 ± 0	Revenue: $196,760 ± 0	Sales: 1000 ± 0
Mid-Dipping	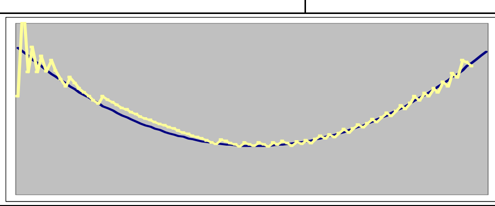		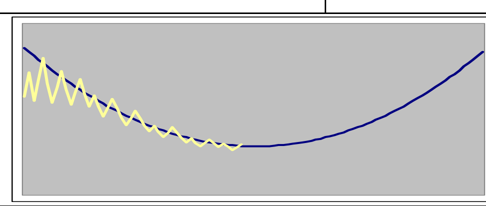	
	Revenue: $159,940 ± 0	Sales: 1000 ± 0	Revenue: $155,710 ± 0	Sales: 1000 ± 0

Table 3: 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.

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 3. 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.

The effect of variance within the buyer population is demonstrated in Table 4. 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.

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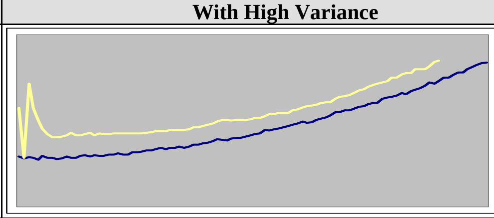		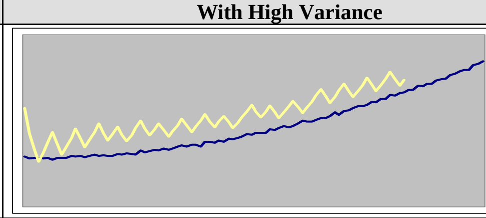	
Valuation Curve:	***Revenue:***	***Sales:***	***Revenue:***	***Sales:***
Increasing	$199,680 ± 149	1000 ± 0	$149,036 ± 1089	1000 ± 0
Decreasing	$208,673 ± 847	994 ± 7	$228,689 ± 1078	1000 ± 0
Mid-Peaking	$275,052 ± 601	991 ± 2	$243,633 ± 1228	1000 ± 0
Mid-Dipping	$202,006 ± 198	1000 ± 0	$189,358 ± 739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

Table 4: 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.

Table 5 presents the simulation results when there are 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 seat 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.

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 6 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 4), falling just above the average buyer curve. DF does 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, 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 keeps 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.

	Goal-Directed Strategy With Few Buyers		Derivative-Following Strategy With Few Buyers	
Sample Pricing Graph	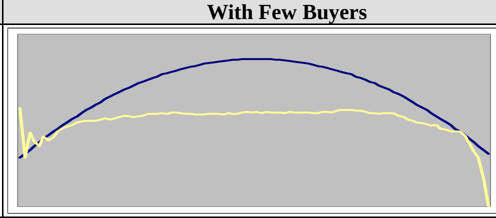		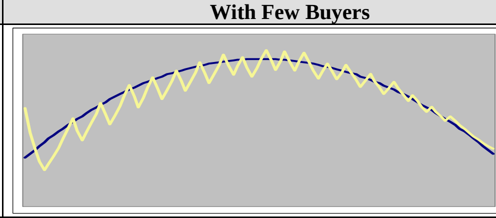	
Valuation Curve:	***Revenue:***	***Sales:***	***Revenue:***	***Sales:***
Increasing	$79,176 ± 4598	814 ± 16	$123112 ± 2690	790 ± 13
Decreasing	$107,441 ± 2642	811 ± 13	$111,492 ± 2759	710 ± 15
Mid-Peaking	$162,147 ± 5530	955 ± 7	$144,724 ± 3497	641 ±14
Mid-Dipping	$66,936 ± 2788	740 ± 16	$120,720 ± 2398	782 ± 12

Table 5: 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	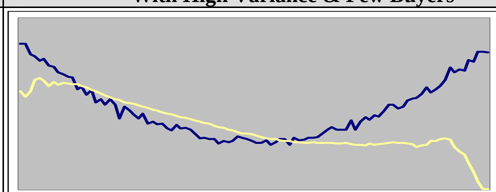		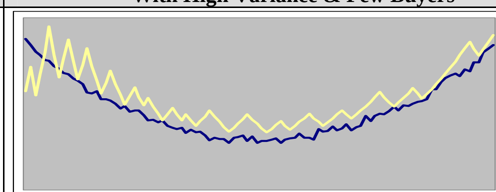	
Valuation Curve:	***Revenue:***	***Sales:***	***Revenue:***	***Sales:***
Increasing	$141,958 ± 4619	999 ± 3	$189,363 ± 4114	977 ± 22
Decreasing	$127,302 ± 2107	889 ± 13	$67,333 ± 4180	328 ± 21
Mid-Peaking	$207,286 ± 2036	972 ± 5	$85,747 ± 5860	335 ± 24
Mid-Dipping	$102,601 ± 4409	907 ± 17	$75,253 ± 4669	372 ± 25

Table 6: 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.

4.2 Competition: Two Sellers in the Market

In a competitive marketplace, the adaptive pricing strategies react to the other strategies in the marketplace, in addition to the buyers' demand. Initially we present a market scenario in which none of the buyers compare prices across sellers or treat the sellers differently. Next we present the effects of comparison-shopping and seller-preference. 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 7 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 [3, 4]. 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 7, 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.

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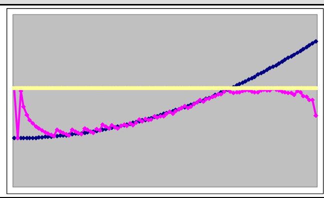		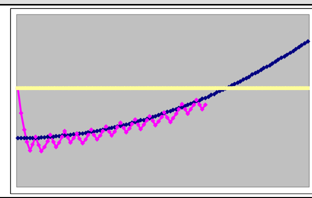		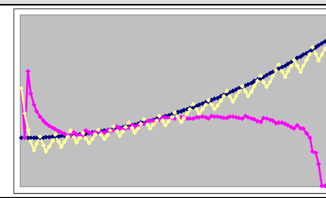	
	Goal-Directed: $157,307 ± 5596 Sales: 990 ± 10	*Fixed Price*: $58,912 ± 3273 Sales: 295 ± 16	*Deriv-Following*: $117,551 ± 1311 Sales: 1000 ± 0	*Fixed Price:* $116,258 ± 4365 Sales: 581 ± 22	*Goal-Directed:* $113,551 ± 8815 Sales: 926 ± 19	*Deriv-Following*: $139,024 ± 3499 Sales: 964 ± 22
Decreasing	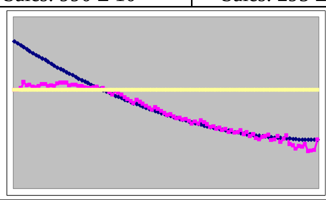		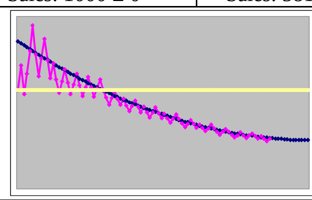		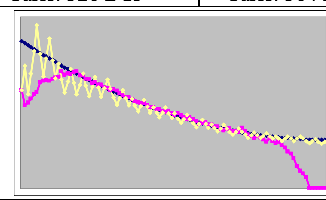	
	Goal-Directed: $134,204 ± 3069 Sales: 976 ± 9	*Fixed Price*: $62,086 ± 3813 Sales: 310 ± 19	*Deriv-Following*: $154,366 ± 1758 Sales: 1000 ± 0	*Fixed Price*: $75,882 ± 3265 Sales: 379 ± 16	*Goal-Directed*: $121,028 ± 3399 Sales: 884 ± 17	*Deriv-Following*: $145,205 ± 5725 Sales: 922 ± 35
Mid-Peaking	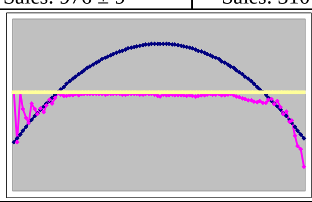		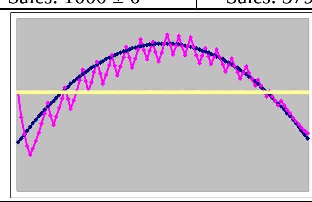		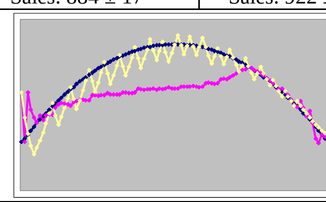	
	Goal-Directed: $178,410 ± 4218 Sales: 980 ± 6	*Fixed Price*: $143,326 ± 4440 Sales: 717 ± 22	*Deriv-Following*: $160,391 ± 5307 Sales: 765 ± 23	*Fixed Price*: $183000 ± 4691 Sales: 915 ± 23	*Goal-Directed*: $190,593 ± 3397 Sales: 980 ± 6	*Deriv-Following*: $164,719 ± 4025 Sales: 752 ± 19
Mid-Dipping	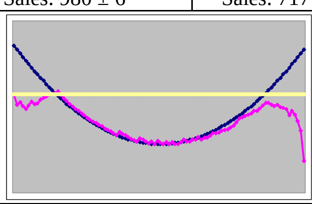		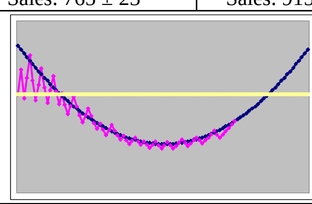		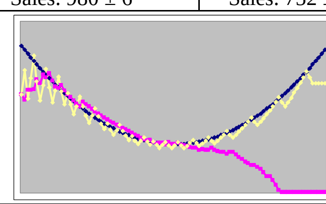	
	Goal-Directed: $144,565 ± 5263 Sales: 987 ± 11	*Fixed Price*: $58,794 ± 11 Sales: 294 ± 16	*Deriv-Following*: $128,381 ± 1143 Sales: 1000 ± 0	*Fixed Price*: $92,126 ± 4029 Sales: 461 ± 20	*Goal-Directed*: $77,342 ± 3890 Sales: 809 ± 20	*Deriv-Following*: $146,845 ± 2850 Sales: 988 ± 17

Table 7: 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 DF strategy is paired with a Fixed-Price seller, in the center column of Table 7, 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 7 look much like when these strategies exist in a monopoly such as in Table 6. 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.

When a population of comparison shoppers is added to the marketplace, there is much more interaction between the two strategies. Table 8 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. The revenue in the right columns of tables 7 and 8 demonstrates this. Examining the results of the two GD strategies, they behave much as they did in a monopoly setting with limited buyers (Table 5), 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.

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 9 shows 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.

	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	$40532 ± 8211	$40532 ± 8211	$35,639 ± 2831	$58,713 ± 1856
Decreasing	$87,058 ± 1875	$87,058 ± 1875	$86512 ± 6549	$86512 ± 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	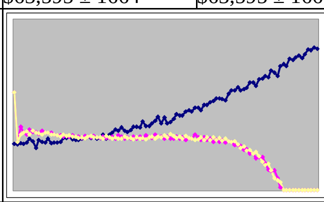		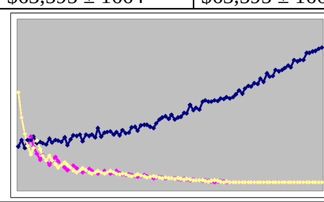		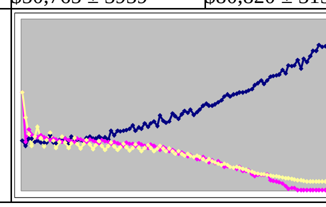	

Table 8: Competition under Comparison Shopping and High Variance

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	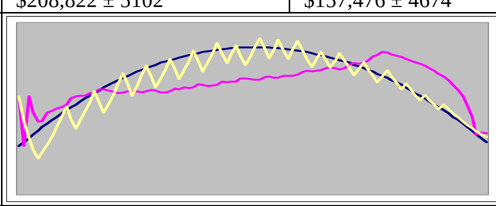		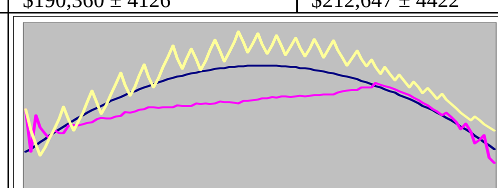	

Table 9: Competition under a Buyer Preference for Different Sellers

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.

4.3 Strategy Analysis Conclusions

While these strategies are computationally straight-forward, 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 or competitors' prices. These strategies point towards some 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 GD strategy is best for slower moving markets where the first and last days do not require fine-tuned price adjustments.
- 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 there is an abundance of buyers, DF performs very well. If buyer demand peaks at some later time, DF does not space out its sales so as to ensure that it sells a large number of goods. Thus the DF strategy excels in a market with an abundance of buyers and a relative 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 about 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. Deck, et al. in [2] 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 (Table 8), 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.

5. RELATED WORK

Finite markets, markets with a finite time horizon, seller inventory, and buyer population, are extremely common yet there is surprisingly little research on how to apply dynamic pricing to this type of market. The field of operations research, which is referred to as revenue or yield management when applied to real-world markets, examines the problem of charging different buyers different prices over time. Traditionally yield management focuses on controlling the number of goods available for sale within discrete pricing levels, such as in the airline industry [1, 10]. Airline yield management systems forecast demand, monitor booking activities and, in response, adjust the number of tickets available in each fare class on an hourly basis. Yield management is effective and practiced by multiple industries, but requires sellers to make assumptions and predictions about the behavior of the marketplace, which is often difficult to do when initially implementing a dynamically priced system.

Previous theoretical studies of pricing strategies in finite markets make conclusions about optimal pricing strategies, but the drawback of these theoretical approaches is the ability to apply results to real-world situations. Gallego & van Ryzin [4], for example, studied this problem with an assumed, static demand curve throughout the market, and Keskinocak & Tayur [6] presented a pricing strategy for a situation with two time periods and two bidders in a multi-unit Dutch auction. The benefit of our simulator is its ability to model diverse and complex scenarios, rather than only simplified cases. By producing tangible, numerical results, the Learning Curve Simulator has enabled us to explore the possibilities and potential for pricing within these complex markets.

Outside of the specific area of finite markets, there are several researchers studying the application of software agents to dynamic pricing. Kephart, Hanson and Greenwald [5] have built a simulated agent marketplace and developed several agent pricing strategies. Their work provided a background for this investigation on successful strategy development. They introduced a successful Derivative-Following strategy, which we used here in a modified form. Like our DF sellers, their myopically optimal agents engaged in price-wars when pitted against one another in a price-sensitive market.

6. CONCLUSION

Returning to the scenario of baseball tickets, it is interesting to consider which dynamic pricing strategy a ballpark should apply to its market. Based on the market conditions of a ballpark (monopoly, high variance among the buyers, and a low marginal cost per seat in the park), we would recommend using a strategy similar to the Goal-Directed strategy. The Goal-Directed strategy’s strength is its focus on selling the entire inventory, sometimes at lower prices, which is a good approach under low marginal costs. The GD strategy also adjusts easily under high buyer variance, as shown in Table 4. Under the conditions in Table 4, the GD strategy performs well under each valuation curve. The valuation curve for baseball tickets could have an unexpected shape, and only actual market data could provide us an accurate curve estimate. If we assume the baseball ticket valuation curve does not continuously decrease over time, then selling all the inventory at the beginning of the market is to the ballpark's disadvantage, which a Fixed-Price policy or Derivative-Following strategy does not protect against.

While our exact algorithm for the Goal-Directed strategy has not been optimized for the baseball ticket market, the process of modeling a market and determining which adaptive strategy is most successful is a useful exercise. The Learning Curve Simulator provides a mechanism for analyzing pricing strategies, making the process of understanding and modeling a market a straightforward task rather than a highly elusive problem.

The future direction of our work will be to both evaluate additional dynamic pricing strategies and to further develop the model of buyer behavior in the simulator. Dynamic programming and Q-learning techniques will be applied to develop new strategy algorithms and the Learning Curve Simulator will continue to serve as the platform for evaluation.

To enrich the behavior of the simulated buyers, we will add the ability to segment the buyer population into different sub-groups that behave independently. This will allow for modeling of mixed populations such as a group of brand loyal customers combined with a group of committed comparison shoppers. The evaluation will determine if our adaptive pricing strategies are able to focus their pricing on the loyal subgroups, only price cutting when necessary.

Another research direction will be to extend the simulator to include buyer strategies, which adapt based on observed changes in the sellers' prices. This is an area where market simulation will be particularly powerful because of the lack of existing market data on how buyers respond to the introduction of dynamic pricing. By incorporating buyer strategies into the market simulator, the buyer valuation curves will be derived from the events within the marketplace.

7. REFERENCES

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

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

[3] Gallego, G. and van Ryzin, G. "A Multiproduct Dynamic Pricing Problem and Its Applications to Network Yield Management." *Operations Research*. vol. 45 (1). 24-41.

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

[5] Kephart, J., Hanson, J. and Greenwald, A. "Dynamic Pricing by Software Agents." *Computer Networks*. vol. 32 (6). 731-752.

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

[7] Morris, J. "A Simulation-based Approach to Dynamic Pricing." Master's Thesis in Media Arts & Sciences, MIT, Cambridge, May 2001.

[8] Morris, J. and Maes, P. "Understanding Dynamic Pricing Agents." Software Demos, Fourth International Conference on Autonomous Agents (Agents 2001), Montreal, Canada (June 2001).

[9] Morris, J., Maes, P. and Greenwald, A. "Learning Curve: Analysis of an Agent Pricing Strategy Under Varying Conditions." Proceedings of the 2001 International Conference on Artificial Intelligence (IC-AI'2001), pp. 1135-1141, Las Vegas, NV (June 2001).

[10] Smith, B.C., Gunther, D.P., Rao, B.V. and Ratliff, R.M. "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).

[11] Smith, M., Bailey, J. and Brynjolfsson, E. "Understanding Digital Markets: Review and Assessment." in Brynjolfsson, E. and Kahin, B. eds. *Understanding the Digital Economy*, MIT Press, Cambridge, MA, 2000.