# Sardine: Dynamic Seller Strategies in an Auction Marketplace

Joan Morris, Peter Ree, Pattie Maes
MIT Media Lab
E15-305, 20 Ames St.
Cambridge, MA 02139
+1 617 253 9603
{joanie, pree, pattie}@media.mit.edu

## ABSTRACT
This paper examines seller strategies for dynamic pricing in an auction-driven marketplace. Specifically, this paper focuses on the airline industry, a field experienced in demand forecasting and dynamic pricing capabilities. Our goal is to discover the relevant factors when a seller dynamically evaluates incoming bids on a finite number of goods. We present two adaptive pricing strategies and evaluate them using a market simulator.

## Keywords
Auctions, dynamic pricing strategies, airline industry, electronic commerce.

## 1. INTRODUCTION
Priceline.com, a popular Internet bidding system, invites consumers to "Name Your Own Price…and Save" on airline tickets, hotel rooms, and groceries [10]. This phrase highlights the benefit of dynamic pricing for the consumer, but current Internet bidding systems do not present a significant benefit to the seller. For example, when a buyer bids for airline tickets on Priceline.com, Priceline does not act as an auction clearinghouse to pass on the winning bid amount to the seller. Instead, when a consumer places a winning bid on a ticket, Priceline purchases the ticket from the airline at a prearranged, discounted fare and sells the ticket to the consumer for the winning bid price. While Priceline presents a radical transaction model for the consumer, it does not begin to capture the potential of dynamic pricing for the seller.

Our research project Sardine challenges the Priceline model by presenting a richer bidding experience for both the buyer and the seller. Yet, while presented within the domain of the airline industry, this work is designed to be applied to any industry domain.

The *Buyer-side* of Sardine examines the consumer problem of constructing meaningful bids for complex products, such as airline tickets [9]. This system's web-based, agent interface assists a consumer in locating and then selecting items to place bids on. After a consumer submits a bid, the consumer's agent sends the bid to the appropriate seller and the seller responses by either accepting or rejecting the bid.

This paper discusses the *Seller-side* of Sardine, which completes the story from the seller's perspective. What happens after the consumer submits a bid to the airline? How does an airline decide to accept or reject a bid? When hundreds of bids are submitted to an airline for the same flight, how does the airline make an intelligent, dynamic decision, maximizing revenue and seat sales? The goal of the *Seller-side* of Sardine is to understand dynamic bidding scenarios and how a seller should approach these questions in a marketplace with a finite number of goods, a finite amount of time, and a changing perceived value of the goods.

As an initial exploration of these research questions, this paper presents a specific scenario for airline ticket bidding, focused on the central question of how an airline can make dynamic decisions to maximize revenue. After the constraints and assumptions of the scenario are discussed, different seller strategies are presented. Through a general marketplace simulator, the airline scenario with these different strategies is tested and our initial results are evaluated. Lastly, related work in the area of pricebots and seller auction strategies is discussed with our plan for future work in this area of market simulation.

## 2. THE AIRLINE SCENARIO
The field of revenue management studies the practice of dynamic pricing in the airline industry. Our market simulator does not attempt to reproduce these established pricing strategies. Instead, this research attempts to understand how to approach pricing strategies when a seller is in a *bidding* marketplace where buyers name their own price and drive the market's prices through their bids.

To do this, we have outlined an auction scenario where buyers place bids on a finite number of goods during a finite number of bidding periods. In airline terminology, this means that travelers have a fixed number of days to bid and purchase a fixed number of seats on a plane flying from location A to location B. The demand for seats changes over time, influencing the number of incoming bids and the value of those bids. The seller's goal is to take advantage of these fluctuations in demand to maximize revenue.

In our scenario, the auction behaves in the following way:

- There are 30 bidding periods (or "days") in which to bid on 100 available airline seats.

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'00*, October 17-20, 2000, Minneapolis, Minnesota.
Copyright 2000 ACM 1-58113-272-7/00/0010…$5.00.

- Each day the airline releases a certain number of seats at a certain reserve price.
- Each day a certain number of buyers send bids to the airline, distributed normally around an average bid price.
- At the end of each day, the airline accepts the bids that are above the reserve price, up to the number of released seats.
- After allocating the winning bids each day, the seller chooses the reserve price and the number of seats to release for the next day.

## 2.1 Determining Winning Bids

The auction functions as a sealed-bid, discriminatory auction [1]. The airline receives a full day's bids before allocating any winning seats. Once the airline receives all the day's bids, it awards seats to the highest bids above that day's reserve price. The number of awarded seats does not exceed the number of seats released for that bidding day. In this allocation method, buyers compete with each other for the available seats on a given day. Each winning bidder pays the price of his or her bid and does not know the amount of any other bids, encouraging each buyer to bid at his or her personal valuation level.

## 2.2 The Market Simulator

A Java-based market simulator called Arena was used to simulate the Sardine auction scenario. Arena is a general simulator of marketplace behavior in an environment filled with buyers, sellers, and negotiation strategies. Arena allows marketplace parameters to be created and altered both before and during a market simulation, allowing users to apply ad hoc strategies to the buyers and sellers during a single simulation [11].

In our scenario, a pared down version of Arena was used, where the buyers, sellers and strategies were created before the simulation initiated. At each bidding period, or day, Arena created a certain quantity of buyers and generated their bids. When all bids had been sent to the seller for evaluation, that day ended and the market variables for the next day were calculated. The simulation terminated when either the product inventory was zero or the simulator had completed all bidding periods.

The simulator took three categories of inputs: *fixed market parameters*, *dynamic market parameters* and *seller control parameters* as outlined in Table 1. The fixed parameters did not change within a simulation. These parameters were the number of seats to be sold, the number of days to sell them in, and the linear slope of consumer demand. The two dynamic market parameters in the simulation were the number of bids on a given day and the average bid value on a given day. These parameters defined the current demand level in the marketplace and their values for any given day were determined by their initial values and the demand slope value for the simulation. The seller control parameters were the two variables left up to the seller, and this was where the seller expressed a pricing strategy. These parameters were the daily reserve price and the number of released seats each day.

**Table 1. Market Simulator Parameters**

| Simulator Parameters | |
|---|---|
| **Fixed market parameters**<br>– No change throughout simulation. | Number of airline seats |
| | Number of bidding days |
| | Demand slope |
| **Dynamic market parameters**<br>– Values changed daily according to demand slope value.<br>– These values were the average of a normal distribution with a fixed standard deviation. | Average bid price (per day) |
| | Average number of bids (per day) |
| **Seller control parameters**<br>– Values changed daily depending on the current strategy simulation. | Number of seats released (per day) |
| | Reserve price (per day) |

The simulator output statistics presented overall market information and per-bidding period information. At the market level, the simulator reported the total revenue (sum of all winning bid prices), the minimum, maximum, and average values of accepted bids, and the percentage of bids above reserve price. For each day (one bidding period), the simulator reported the number of seats sold, the number of released seats not sold, the current inventory level, the reserve price and the average price of the buyers' bids.

## 2.3 Model Assumptions

Now that the auction scenario is outlined in terms of its control variables, it is worth enumerating the built-in assumptions of the system. They are as follows:

1. An individual buyer can only place one bid per day. This means the seller does not consider users placing multiple bids on the same day.
2. All bids are committing, so when the airline accepts a buyer's bid, the buyer is obligated to purchase the ticket. This does not necessarily have to be the case, as presented by Sandholm in [12] and this constraint may be released in future simulations.
3. This model assumes there are no fixed priced seats. The scenario could be thought of as a plane with 200 seats where 100 of the seats are biddable and 100 have fixed prices, but this model does not address how to allocate which seats are dynamically priced versus fixed priced.
4. There are no assumptions made about how consumer demand changes over time. Although consumers' behavior in today's airline market is somewhat predictable, we hypothesize this is because consumers understand the established pricing practices of airlines [6]. For example, leisure travelers know to buy tickets twenty-one days in advance of the date of travel to get a less expensive fare. In this bidding scenario, the seller does not have a pre-set pricing scheme – instead, the consumers drive the pricing. Thus, this system makes no assumptions about the direction or fluctuations of consumer demand.

**Table 2: Simulation Input Values**

| Simulation Parameters | Base Case | Reserve Pricing Strategy | Seat Releasing Strategy |
|---|---|---|---|
| **Number of airline seats** | 100 seats | | |
| **Number of days** | 30 days | | |
| **Demand slope** | Three trials, with values –5, 0, and 5. | | |
| **Average bid price** | Initially $415, linearly decreasing, remaining constant, or increasing, depending on simulation's slope value. | | |
| **Average number of bids** | Initially 40 bids, linearly increasing, remaining constant, or decreasing, depending on the average bid price. | | |
| **Number of seats released** | 10 seats each day | 10 seats each day | Initial value of 1 seat, and then changing according to the seat releasing strategy. |
| **Reserve price** | Four trials, with fixed values of $200, $350, $415, and $450. | Four trials, with *initial* values of $200, $350, $415, and $450, and then changing according to the reserve pricing strategy. | $200 each day |

5. During each day of the auction, consumer demand is represented by two heuristics: the average bid value and the number of incoming bids. We assume that when consumers perceive the market value for a flight to be low, then they place more bids at that low value, distributed normally around a perceived value. The opposite case holds as well, that if consumers perceive the market value to be high, fewer consumers place bids, averaged around a higher perceived value.

   The demand slope parameter determines how the average bid value increases or decreases over the simulation and this in turn determines how the number of incoming bids decreases or increases with the movement of bid values. The number of incoming bids changes in reverse proportion to the change in the average bid value.

6. The seller does not consider costs, such as the marginal cost of each seat, in its pricing strategy. The calculated revenue discussed throughout the paper refers to the total sum of the seat sales, regardless of the airline's costs. Marginal costs could be incorporated later into the system by setting a minimum reserve price.

# 3. AIRLINE STRATEGIES

As stated earlier, the airline can control the outcome of bidding by altering two decision parameters: the daily reserve price and the number of seats released each day. Using these two variables, this study proposes the following two airline decision strategies as methods for increasing revenue.

## 3.1 Reserve Pricing Strategy

In the reserve pricing strategy, the airline attempts to maximize revenue by adjusting the daily reserve price to ensure that all seats sell by the end of the auction. This computationally basic pricing strategy is categorized as *derivative following* [7], because the airline makes its pricing decisions based on observing its own behavior and not any of the market behaviors.

This pricing strategy computes today's reserve price by looking at yesterday's reserve price and the amount of seats sold up until this point in the auction. The specific algorithm for setting a day's reserve price is:

- If total number of seats sold < (number of days past) *

  (total seats/total days in auction),

  => then *decrease* reserve price by (% off-target/2)

- If total number of seats sold > (number of days past) *

  (total seats/total days in auction),

  => then *increase* the reserve price by % off-target/2).

- If total number of seats sold = (number of days past) *

  (total seats/total days in auction),

  => then no change.

*Hypothesis*: This strategy will work best when demand is relatively constant because it does not take into account any demand movement or predictions. In addition, the first day's reserve price should be as accurate as possible or the first days could produce poor sales.

## 3.2 Seat Releasing Strategy

In this strategy, the airline attempts to maximize revenue by releasing the seats for bidding in relation to the current demand level. The airline tracks and tries to predict the movement of the bid prices so that more seats are released and sold during relatively higher demand levels. We describe this strategy as *myopically optimal* [7], because the airline uses information about the buyers' behavior to determine an optimal strategy, while ignoring all other market factors, such as competitors.

To accomplish this, the airline observes the previous two days' demand and guesses which direction demand is moving, up or down, and accordingly releases more or less seats. The reserve price in this strategy serves as a safety net, leaving the potential for extremely low bids to be accepted. To ensure that all seats sell, in the last 2 days of bidding, the airline releases all seats.

The algorithm for releasing seats is as follows:

- Δ = (yesterday's av. bid value) – (day before yesterday's av. bid value) / (yesterday's av. bid value)
- If Δ > 0,

  => then *increase* number of seats released by Δ.
- If Δ < 0,

  => then *decrease* number of seats released by Δ.

Except:

- If day <= 2,

  => then release 1 seat.
- If day >= (second to last day),

  => then release all remaining seats.

*Hypothesis*: If demand goes in constant, steep cycles, the seller can easily predict and take advantage of swings in demand, but unpredictable demand could result in not enough seats being released and sold.

# 4. SIMULATION RESULTS

We implemented the two strategies discussed with the market simulator to evaluate their effectiveness in maximizing revenue. These strategy simulations were compared to a base case where the seller applied no dynamic strategy by keeping the reserve price and the number of released seats constant through the bidding periods. The results of these strategy simulations are presented here with a discussion of the strengths and weaknesses of the different approaches.

Table 2 outlines the input values for the base case simulation and for the two strategy simulations. As shown in the table, the fixed and dynamic market parameters had the same values across all simulations. The market control variables remained static in the base case and changed during the appropriate strategy simulations depending on which control variable was being tested.

To test the seller strategies under varied market demand conditions, the base case and each strategy were tested with different values for the demand slope. Under a constant average bid price, the slope was zero, with a linearly decreasing average bid price, the slope was negative five, and with a linearly increasing average bid price, the slope was five.

## 4.1 Base Case

The base case sets a benchmark for comparing the seller strategies. As shown in Table 2, in the base case, the airline released 10 seats each day, which was high enough to sell all 100 seats under the three demand conditions. This value was purposefully chosen so it was not a constraining variable at the termination of the simulation. The reserve price was fixed across each simulation trial at $200, $350, $415, and $450. Two of these values fell below the initial average bid price, one was the initial price, and one fell above. These base case parameters produced consistent, deterministic revenue results across all trials. It is worth noting that the choice of these values affected the outcome of both the base case and the strategy simulations and if different values were used, the numeric results would vary, but as we will discuss, our conclusions would not change.

Twelve trials were run under the base condition (3 demand conditions * 4 fixed reserve price levels). The sum of all winning bid prices, or the total revenue, for each simulation is show in Figure 1. For the three different reserve price settings $200, $350, and $415, the revenue remained relatively constant over the different demand levels, decreasing as demand decreased. In the case where the reserve price was $450, above the average bid level of $415, when demand increased, this reserve price produced the most revenue, but under decreasing demand, revenue approached zero.

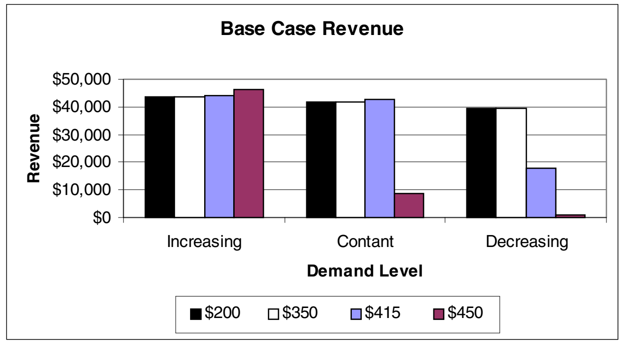


**Figure 1: Base Case Revenue at reserve price levels of $200, $350, $415, and $450.**

Intuitively, the simulation base case demonstrates that revenue is highest when demand is increasing and the reserve price is below the average bid amount. As an overall revenue maximizing position, if an airline is not going to change any control variables during the auction, then it should pick a reserve price below the minimum expected level of demand to ensure the greatest return.

With the base case in hand, we next applied dynamic strategies to the airline scenario in attempt to increase these revenue levels.

## 4.2 Reserve Pricing Strategy

To review, the reserve pricing strategy adjusts the reserve price based on the market's behavior. Specifically, the airline adjusts reserve price in an attempt to stay on target to sell all the seats by the end of 30 days. Figures 2, 3, and 4 show the behavior of the reserve price under the three cases of increasing, constant, and decreasing bid amounts, shown for the case where initial reserve price is $350. The results were similar in the three other initial reserve price simulations, run at the three different demand levels. Through a periodic over-shooting and under-shooting of the reserve price, the airline tracked the level of the bids fairly closely after about the fifth bidding period.

An interesting aspect to these results is while this strategy is derivative following, meaning the seller only monitored how many seats it had sold, the results show that the reserve price closely tracked the demand level, or current bid price. By following the average bid price, the seller captured the higher end of the bid distribution each day, rejecting the lower half of the bid distribution. Compared with not varying the reserve price at all, this increased revenue across all initial reserve price settings, as shown in Figures 5, 6, and 7. The increase in revenue was most significant when the initial reserve price was $450, above the initial market demand level of $415, because in comparison to the base case, the reserve pricing strategy allowed the reserve price to move closer to the bid level.

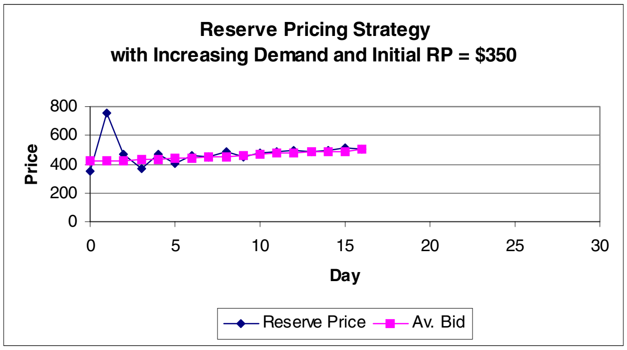


**Figure 2: As demand increased, the airline adjusted the reserve price to the bid level and then followed its increase.**

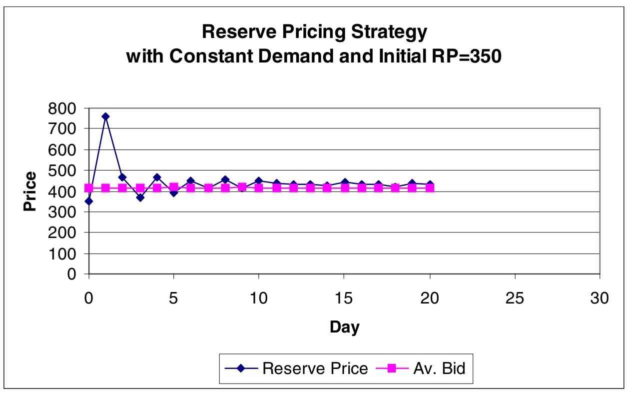


**Figure 3: While demand stayed constant, the reserve price oscillated above and below the average bid price.**

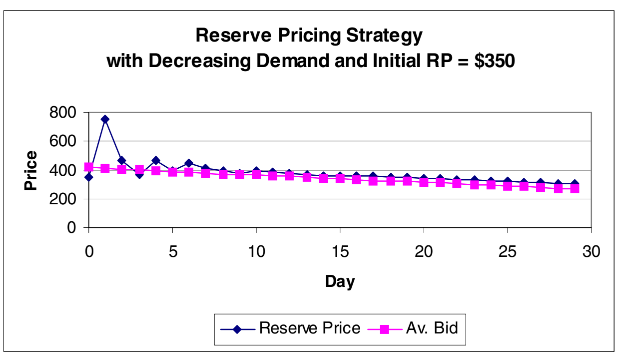


**Figure 4: While demand decreased, the reserve price adjusted and decreased.**

An exception to this increase in revenue was the case of decreasing demand with an initial reserve price of $350, marked by the ‘A’ on Figure 7. This was the only trial across all twelve simulations where the base case out performed the reserve pricing strategy. The specific movement of the reserve price is shown in Figure 4. Not shown in this charts is that the airline did not sell all of the airplane seats in the 30 days using the reserve pricing strategy. So while the strategy succeeded in selling seats from the upper half of the bid price distribution, producing higher revenue on a per seat basis, it did not sell enough of the seats to match the revenue made in the base case.

This exception points out the importance of the choice of initial parameter values. The number of released seats, 10 per day, was not

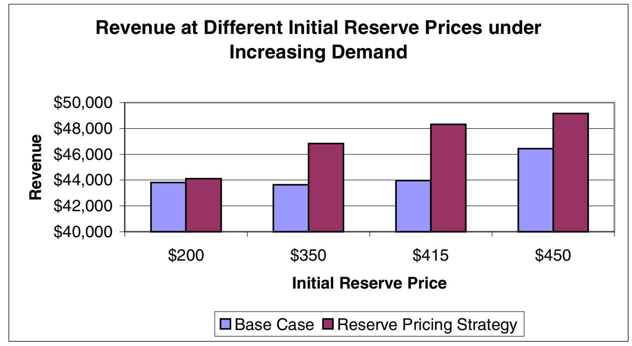


**Figure 5: The reserve pricing strategy compared to the base case, where demand is increasing.**

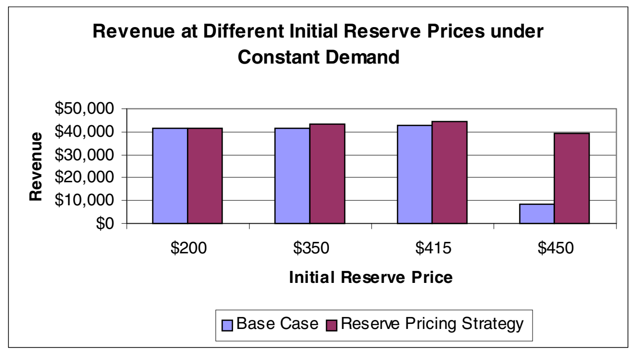


**Figure 6: The reserve pricing strategy compared to the base case, where demand is constant.**

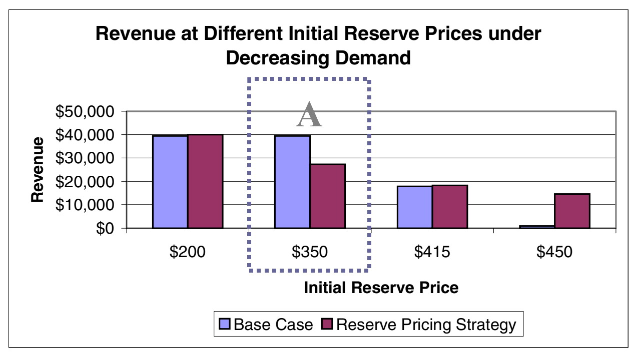


**Figure 7: The reserve pricing strategy compared to the base case, where demand is decreasing.**

enough to produce sales for all 100 seats in this case. As stated earlier, we specifically decided to release 10 seats a day because this was not a constraining variable at the termination of 30 days for the base case trials, but it turns out it was a constraining variable in this strategy simulation. So either a better initial value could be chosen so it did not constrain the results, or this situation could be used as a way of developing improvements to the derivative pricing algorithm. One such improvement could be, if the end of the end of the simulation was approaching, because of either limited time or inventory, the airline could make more aggressive changes in reserve price to ensure that all the seats were sold and that they were sold as close to the last day of bidding as possible.

## 4.3 Seat Releasing Strategy

The second strategy attempts to maximize revenue by dynamically determining how many seats to release depending on the current level of market demand. In this seat releasing strategy, the airline attempts to sell more tickets when demand is at its peak and less seats when demand is at its low. The challenge is knowing whether the current level of demand is a relative high point or low point, especially at the beginning of the bidding periods.

Our current simulator only models linear demand, which exaggerates this challenge. We conducted three trials of the strategy under increasing, constant, and decreasing demand, as stated in Table 2.

The seat releasing algorithm determines how many seats to release each day by looking at the previous two days' average bid prices, as outlined in Section 3.2. The number of seats to release is determined by adjusting the previous day's number of seats by the percentage change in the previous days' demand. There are exception cases at either end of the simulation: at the beginning of bidding, the seller releases one seat on each the first and second bidding days, and then the seller releases all of the seats on the second to last day of bidding, leaving two bidding days to sell any remaining seats. The reserve price in this strategy represents the minimum acceptable price, so is set to $200 in all simulations of this strategy. The number of seats released each day can not be less than one.

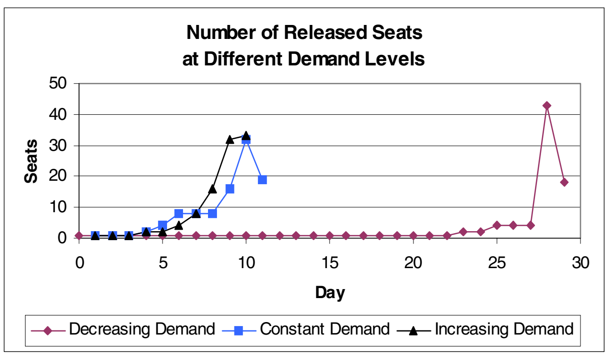


**Figure 8: The number of seats released during each bidding period, at the different demand levels.**

The results of the simulations are shown in Figure 8 and the corresponding revenue compared with the base case is shown in Figure 9. By looking at the revenue results, it is clear this strategy did not significantly increase revenue over the base case.

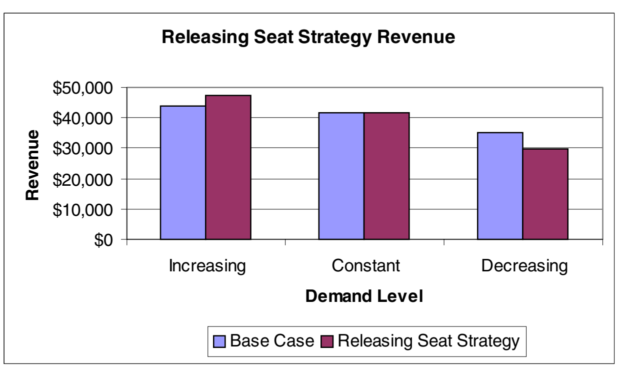


**Figure 9: The releasing seat strategy compared with the base case, where reserve price is set to $200.**

The goal of this strategy was to sell the maximum number of seats at the peak demand and the minimum number of seats at the minimum demand level. In fact, the opposite occurred! In Figure 8, the longest curve shows the number of seats released for bidding each day when demand was decreasing. As a result of the algorithm, the seller released one seat each day of bidding, until the last two days when all seats were released to prevent unsold seats at the end of the simulation. In effect, this resulted in the majority of the seats being sold when demand was at its absolute lowest. The reverse happened with increasing demand. As the average bid price went up, as designed, the seller released more and more seats each day, selling all the seats by the tenth day of bidding, but this was well before the maximum bidding price was reached on day thirty. So again, the seller ended up selling the most seats when demand was at its lowest over the thirty days.

While in theory this strategy seemed like a better approach than the simple reserve pricing algorithm, it failed to produce positive results. Our conclusion is that tracking demand based on short-sighted buyer observances does not provide enough information to make an intelligent strategy. For a seat releasing strategy to work, more knowledge about the marketplace is needed. For example, knowledge about the current bidding period, how many seats are left to be sold, and a complete history of the bid price movement could enhance the algorithm.

## 5. CONCLUSION

Comparison of the derivative-following strategy for adjusting reserve price with the myopically optimal approach for adjusting the number of seats to release leads to some interesting observations.

While not obvious, by adjusting reserve price based on the number of seats sold so far in the simulation, the reserve price was able to track demand very well. While an adjustment of the algorithm could improve its ability to sell all the seats by the end of the bidding, it proved to be an effective method of increasing revenue over the base case.

In contrast, the strategy of adjusting the number of seats released according to the demand level did not succeed in increasing revenue. While information about the bid amounts can indicate the level of demand, information about the state of the marketplace and the seller's sales level is necessary to make a strategic decision. Thus in this airline bidding scenario, the basic myopically optimal approach did not succeed. While the model of demand in the scenario was overly simplified, it served to exaggerate the flaw in the strategy.

## 6. RELATED WORK

The work presented in this paper relates directly to the current work on pricing agents, or pricebots, done by Greenwald et al. [7]. Their work on pricebot dynamics studies the behavior of buying and selling agents in a dynamic marketplace with multiple strategies on the part of both buyers and sellers. In their simulations, sellers with game-theoretic, myopically optimal, derivative following, and Q-learning strategies compete against each other in a marketplace. While these pricebots are not operating in an auction setting, the basic principles of observing one's own behavior, the buyers' behavior, and the other sellers' behavior hold true to both marketplaces. As discussed in the Future Work section, many of these approaches can be applied to the Sardine scenario to explore more seller strategies.

Brooks et al. [5] also performs analysises of pricing agents in a simulated market environment and discusses 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. This principle has held for our auction scenario as well.

Beam et al. [2] present pricing strategies for sellers in a single-item auction and examine these strategies in real auction settings. They discuss extensions for an N-item auction and propose strategies for managing inventory surplus. Their N-item auction has several similar features to our airline auction scenario, but it does not address the issues involved when the perceived value of the good changes over time.

The field of revenue management examines the challenge of revenue maximization in the domain of the airline industry [4]. Commercial revenue management systems use dynamic pricing to balance the supply and demand for airline seats. These systems forecast demand, analyze competitive fares, monitor booking activities and in response, adjust ticket prices on an hourly basis. By offering different fares at different "booking classes" with varied restrictions, airlines attempt to capture sales from each type of consumer (leisure, business, etc…), at their maximum willing price. According to [3], today's revenue management techniques make overly simplistic assumptions about how consumers segment across different booking classes and respond to the different restrictions associated with each class.

These short comings aside, revenue management techniques differ in a significant way from Sardine. Revenue management attempts to *forecast* demand levels so pricing decisions can be made today based on predictions of tomorrow's behavior. The key point of Sardine is that, in an auction setting, the buyers name the market price with their bids, so airlines can make decisions during and in response to consumer behavior.

## 7. FUTURE WORK

The immediate direction of this work is to conceive of more seller strategies and test them with the simulator. Furthermore, by combining strategies, the behavior of the airline may prove to generate higher and more consistent levels of revenue. For example, by combining the two strategies of reserve pricing and seat releasing, the airline could target sales to the higher end of the bid price distribution through reserve price setting, while using the seat releasing strategy to fine-tune the number of seat sales on each day.

Currently, the simulator's model of consumer demand is overly simplified. We will enhance the simulator to allow for fluctuating demand levels to test the robustness of the airline strategies under these conditions. The eventual goal for the simulator is to allow multiple sellers with competing strategies to co-exist within the same auction scenario, much like Kephart's pricebots [8].

While these are basic enhancements to the existing simulation, there are several bigger research questions we would like to address involving the model of bidding and market behavior.

For example, another direction of work is to consider what happens when a seller considers more than just maximizing revenue, but considers long term goals like customer loyalty or airline reputation. If the airline is interested in customer loyalty, then it may accept a lower bid from a customer it is familiar with. Another aspect an airline could consider is the increasing marginal cost as the plane fills to capacity, manifested in longer lines at check-in, lost baggage, and irritable travelers. If an airline places value on ensuring a smooth travel experience for its consumers and employees, then other auction strategies may come in to play.

As stated under the scenario assumptions, the buyer bids in this system are committed, so the airline does not need to consider buyers who do not purchase an accepted bid. In the *Buyer-side* of Sardine, the interface does not require buyers to purchase the plane ticket of an accepted bid. In fact, the buyer is able to submit multiple bids to multiple airlines for the same trip, an "X-OR" combinatorial auction. To incorporate this aspect into the *Seller-side* of Sardine, a method for leveled commitment contracts should be developed, which would implement a penalty of some kind if the buyer chooses to not purchase the plane ticket. While in Sandholm's work [12], the penalty was assumed to be a monetary exchange, within the airline purchasing process, this penalty could take many forms, such as disallowing re-bidding or limiting airline flight benefits such as frequent flier mileage credits.

## 8. REFERENCES


[1] Agorics, "A Survey of Auctions," http://www.agorics.com/new.html (1996).

[2] Beam, C., A. Segev, and J. G. Shanthikumar, "Optimal Design of Internet-based Auctions," Working Paper 98-WP-1034, Fisher Center for Information Technology & Management, Haas School of Business, University of Berkeley, Berkeley, CA (January, 1999).

[3] Botimer, T. C. and P. P. Belobaba, "Airline pricing and fare product differentiation: A new theoretical framework," Journal of the Operations Research Society Vol. 50, No. 11 (November, 1999).

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

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

[6] Francett, B. "Using pricing to change buying patterns," Computerworld, v. 24, n. 43, pp. 116-7 (October, 29 1990).

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

[8] Kephart, J. and A. Greenwald, "Shopbot Economics," Proceedings of Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, (July, 1999).

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

[10] Priceline, http://www.priceline.com (August, 2000).

[11] Ree, P., "Arena: Simulating E-Commerce Strategies." M.Eng. Thesis, MIT, Cambridge, MA (May, 2000).

[12] Sandholm, T. and V. Lesser, "Advantages of a Leveled Commitment Contracting Protocol," Thirteenth National Conference on Artificial Intelligence (AAAI-96), Portland, OR (1996).