# Multi-attribute Dynamic Pricing for Online Markets using Intelligent Agents

Prithviraj(Raj) Dasgupta and Yoshitsugu Hashimoto
Computer Science Department
University of Nebraska, Omaha, NE 68182, USA.
{pdasgupta, yhashimoto}@mail.unomaha.edu

## Abstract

*Intelligent agents called pricebots provide a convienient mechanism for implementing automated dynamic pricing algorithms for sellers in an online economy. Pricebots enable an online seller to dynamically calculate a competitive price for a product in response to variations in market parameters such as competitors' prices and consumers' purchase preferences. Previous research on pricebot mediated pricing makes certain simplifying assumptions of online markets such as providing sellers with complete knowledge of market parameters to facililate calculations by the dynamic pricing algorithm, and, considering product price as the only attribute that determines consumers' purchase decision. In this paper, we address the problem of dynamic pricing in a competitive online economy where a product is differentiated by buyers and sellers on multiple attributes, and, sellers possess limited knowledge about market parameters. A seller uses a collaborative filtering algorithm to determine temporal consumers' purchase preferences followed by a dynamic pricing algorithm to determine a competitive price for the product. Simulation results using our market model show that collaborative filtering enabled dynamic pricing techniques compare favorably against other dynamic pricing algorithms. Collaborative filtering enables sellers to rapidly identify temporal customer preferences and improve sellers' profits.*

**Keywords**: Agent mediated e-commerce, dynamic pricing, collaborative filtering.

## 1. Introduction

With the advent of the Internet, e-commerce has improved traditional business processes by increasing the accessibility between buyers and sellers, and, automating various trading processes. Buyers using online e-commerce services such as comparison shopping and merchant rating[3, 17] are able to make an informed purchase decision after comparing offers from different sellers from all over the world. This has resulted in increased competition between sellers offering similar products. Online sellers have responded by using automated pricing techniques that dynamically update the advertised price of a product. Intelligent agents called *pricebots*[15] provide a suitable paradigm for rapidly and accurately updating the price of product for a seller by implementing a dynamic pricing algorithm. Although dynamic pricing had been implemented by some online merchants such as Amazon Inc.[1], it is yet to be adopted widely in e-commerce. The principal drawback of the dynamic pricing mechanism employed[4] was that it offered identical products to different consumers at different prices based on the consumers' purchase preferences.

Surveys of consumers who purchase products online[8, 20] reveal that online buyers are frequently willing to pay an elevated price for enhanced values on particular product attributes such as delivery time, seller reputation, and after-sales service. Different consumers have been reported to prefer different product attributes and these preferences vary over time depending on exogenous factors such as sales promotion, aggressive advertising and even time of the year. Therefore, it is important for an online seller to differentiate a product along multiple attributes, and, determine a potential buyer's purchase preferences over the different product attributes so that it can tailor its offer to meet the buyer's requirements and improve its profits. In this paper, we describe a software agent enabled dynamic pricing algorithm for online sellers which uses collaborative filtering to determine buyers' purchase preferences, and then, calculates a profit maximizing price for the seller using a dynamic pricing algorithm.

## 2. Multi-attribute Dynamic Pricing Using Intelligent Agents

Automated dynamic pricing for online economies has been analyzed and implemented through simulated market models in[7, 9, 10, 15]. A seller in an automated market employs the services of a pricebot that dynamically calculates a profit maximizing price of a product for the seller


AAMAS'04, July 19-23, 2004, New York, New York, USA.

in response to fluctuations in market parameters such as prices and profits of competing sellers and buyers' reservation prices. The updated product price is posted by the seller in the market at regular intervals to continue to attract buyers while maintaining a competitive edge. Most of these systems consider product price as the principal determinant of a consumer's purchase decision. However, micro-economic literature and online consumer surveys[8, 14, 13] suggest that a consumer's purchase decision is determined by multiple product attributes including price, delivery time, seller reputation, product quality and after-sales service. Therefore, it makes sense to model an economy where a product is differentiated by buyers and sellers over multiple attributes. Kephart et al[15] have also shown that a market model in which buyers discriminate between sellers based only on product price is susceptible to price-wars between sellers which prevent them from converging on an equilibrium. However, in real markets perennial price wars are infrequent because consumer preferences over different product attributes are temporal[13]. Consequently, in addition to determining buyer attribute preferences, sellers must also continuously update those preferences to remain competitive in a market with dynamic consumer preferences.

Economic analyses of e-markets[12, 15] assume that every seller has complete knowledge of market parameters such as buyer reservation prices, competing sellers' prices and profits, and equilibrium points that are then used by the seller's pricing algorithm. In real-life, a seller has to explicitly request other competing sellers for their price information. Rapid fluctuation of market prices in an online economy can frequently leave a seller with outdated competitor price information that can cause the seller's dynamic pricing algorithm to function incorrectly. Also, it is difficult, if not impossible for sellers to obtain prior information about buyers' parameters. Therefore, in this paper, we have not assumed prior knowledge about market parameters to be available with online sellers. A seller in our model determines and updates the temporally varying buyers' purchase preferences over different product attributes using a collaborative filtering algorithm. The information about the buyers' attribute preferences is then used by the seller's dynamic pricing algorithm to calculate a competitive price of the product that is posted by the seller in the market.

## 3. Model

Real-life internet economies involve complex interactions between several buyers, sellers and possibly brokers that facilitate trading. We have made certain simplifying assumptions of an online economy that simplify analysis while retaining the essential features of the market. Our Internet market model is based on the shopbot economy model of Kephart and Greenwald[15]. Our market model consists of $S$ sellers who compete to provide $B$ buyers ($B \gg S$) with a single indivisible commodity. Sellers behave as profit maximizers. The goods being sold are consumables and every seller has sufficient supply of the good to last the lifetime of the buyers. Buyers return to the market repeatedly to re-purchase the good. Examples of such markets include commodities used daily such as groceries and even renewable services such as telephone or Internet services. A product is characterized by buyers and sellers on multiple attributes. A seller offers a slightly different price for the product along each of its attributes. As shown in Figure 1, a buyer first requests a quote from the sellers for the price along the buyer's preferred product attribute and selects the seller that makes the best offer. The buyer's preferred attribute is not revealed to a seller when the buyer makes a quote request. Therefore, the objective of a profit maximizing seller is to determine a buyer's preferred attribute in response to the buyer's quote request. The seller can then calculate a competitive price of the product along the buyer's preferred attribute and make an attractive offer to the buyer.

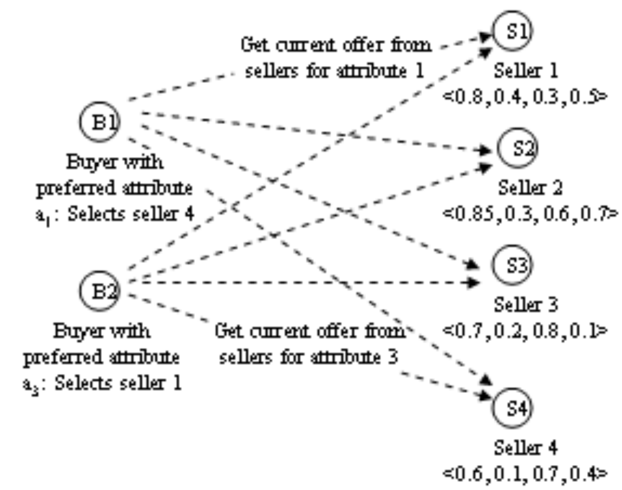


**Figure 1. A hypothetical market showing two buyers with preferred attributes as $a_1$ and $a_3$ respectively making a quote request to four sellers and selecting the seller that offers the best price for the product on their respective attributes. The tuple $< p_{a_1}, p_{a_2}, p_{a_3}, p_{a_4} >$ below each seller denotes the normalized price offered by each seller on the different product attributes.**

**Sellers.** A seller $S_j$ enters the market with an initial posted price $p^0_{a_i,S_j}$ for a unit of the good under attribute $a_i$. Every seller has a unit production cost $p_{co}$ below which it is not willing to sell the good. The price charged by a seller $S_j$ during interval $t$ along product attribute $a_i$ is denoted by $p^t_{a_i,S_j}$. This price is updated by the pricebot at intervals $\tau_{S_j}$ using the dynamic pricing algorithm. All prices are normalized to ensure uniformity. Different sellers update their product prices asynchronously and each seller adopts its own pricing strategy.

**Buyers.** Buyers make quote requests to sellers at a constant rate $\rho_b$ and determine the offers made by the different sellers. We assume that the buyers are aware of the existence of all the sellers in the market. Since we analyze pricing algorithms for sellers, seller discovery is not treated as a major issue in the model. In the real world, online buyers can employ comparison shopping services[17] to discover sellers. Every buyer has one of the product's attributes as its preferred attribute. A buyer's preferred product attribute is drawn from a discrete distribution $f_{pa}$ and is allowed to vary temporally. Every buyer $b$ has a reservation price $p_{r,b}$ above which it is not willing to purchase the product. Buyer reservation prices are drawn from a fixed continuous distribution $f_{pr}$. Sellers are unaware of both the distributions $f_{pa}$ and $f_{pr}$. The utility of the product to a buyer $b$ with preferred attribute $a_i$, which makes a quote request from seller $S_j$ during interval $t$ is given by $U_{b,i,S_j} = p_{r,b} - p^t_{a_i,S_j}$. The purchase decision is made by buyer $b$ by comparing offers made by all sellers on its preferred attribute $a_i$ and selecting seller $S_k$ whose offer gives the maximum utility $U_{b,i,S_k}$. Buyer $b$ then pays seller $S_k$ the price of the product and the seller delivers the good. Payment and product delivery are not discussed any further here as we concentrate on seller pricing strategies.

To make a competitive offer, a seller should be able to determine the preferred attribute of a buyer in response to the buyer's purchase request and then calculate a competitive price on that attribute. To achieve this a seller has to correctly estimate the distribution of preferred attributes $f_{pa}$ and the distribution of reservation prices of buyers $f_{pr}$. In the next section, we analyze seller profits in a mathematical model of the market and then we describe the algorithms that are used by sellers to estimate buyers' preferences and dynamically update the posted prices over the different product attributes.

### 3.1. Analysis

Since we assume sellers are profit maximizers, the objective of every seller in our model is to determine a price for every attribute of the product which maximizes the seller's profit. However, the pricing function of a seller cannot be stationary as other competing sellers would revise their prices to improve their offers and attract buyers away from the seller. Therefore, every seller must update the prices it charges on different product attributes at intervals of $\tau_{S_j}$ in response to competitors' pricing strategies, and also, in response to changes in buyers' preferred attributes. We now analyze the pricing problem for sellers in our model. The market described in Section 3 is defined by the following parameters:

| | |
|---|---|
| $B$ | Number of buyers in the market |
| $S_N$ | Number sellers in the market |
| $\nu$ | Number of attributes of a product. All products have the same attributes. |
| $a_i$ | The i-th product attribute, $i = 1...N$ |
| $f_{pr}$ | Continuous distribution of number of buyers over reservation prices. This distribution is time invariant. |
| $p_{r,b}$ | Buyer $b$'s reservation price. |
| $f_{pa}$ | Discrete distribution of buyer preferences over product attributes. This distribution varies over time. |
| $\rho_b$ | Arrival rate of quote requests from buyers to a seller |
| $p^0_{a_i,S_j}$ | Seller $S_j$'s market entry price for a unit of the item on attribute $a_i$. |
| $p_{co}$ | Unit production cost for a seller. This cost is is assumed to be uniform over all attributes. |
| $\tau_{S_j}$ | Price update interval for seller $S_j$ |
| $p^t_{a_i,S_j}$ | Price offered by seller $S_j$ to a buyer with preferred attribute $a_i$ during interval $t$ |

The fraction of buyers that make quote requests to a seller $S_j$ during one update interval is given by $\beta_{S_j} = \rho_b \tau_{S_j}/B$. We analyze two separate scenarios: i) When there is only one seller in a monopolistic market, and, ii) When there are multiple sellers in a competitive market.

**Case I** $S_N = 1$ *(Single monopolist seller).* Let us denote the single seller as $S_1$.t The number of buyers with preference on product attribute $a_i$ is given by $f_{pa}(a_i)$. The number of purchase requests to a seller in one update interval from these buyers is $\beta_{S_1} \times f_{pa}(a_i)$. The unit profit to the seller is given by $(p^t_{a_i,S_1} - p_{co})$ and the number of buyers for whom the reservation price is not exceeded when the seller offers a price $p^t_{a_i,S_1}$ is given by $\int_{p_{co}}^{p} f_{pr}(p^t_{a_i,S_1})$. The profit to the seller from the buyers with preference on attribute $a_i$ is then given by:

$$\pi^t_{S_1,i} = \beta_{S_1} f_{pa}(a_i)(p^t_{a_i,S_1} - p_{co}) \int_{p_{co}}^{p} f_{pr}(p^t_{a_i,S_1})dp \quad (1)$$

and, the total profit to the seller is given by:

$$\pi^t_{S_1} = \beta_{S_1} \sum_{a_i} f_{pa}(a_i)(p^t_{a_i,S_1} - p_{co}) \int_{p_{co}}^{p} f_{pr}(p^t_{a_i,S_1})dp \quad (2)$$

Since the seller is a profit maximizer, the maximum profit to the seller can be obtained by setting $\delta\pi^t_{S_1}/\delta p^t_{a_i,S_1} = 0$ in Equation 2. $f_{pa}$ becomes a constant in the differential. This means that a monopolist seller's pricing strategy is independent of the buyers' attribute preferences. Further analysis of Equation 2, assuming that the seller uses a normal distribution to model $f_{pr}$ shows that a monopolist seller can charge a fixed profit-maximizing price that attracts the most number of buyers.

**Case II** $S_N$ *competing sellers.* We consider the analysis from the point of view of a single seller $S_1$ that is competing with all other sellers $S_2, ...S_N$ in the market. Similar to equation 2 we can write the profit to seller $S_1$ on attribute $a_i$ as:

$$\pi^t_{S_1,i} = \beta_{S_1} f_{pa}(a_i)(p^t_{a_i,S_1} - p_{co}) \int_{p_{co}}^{p} f_{pr}(p^t_{a_i,S_1})dp \quad (3)$$

provided $U_{b,i,S_1} > U_{b,i,S_j}$, and, $\pi_{S_1,i} = 0$, if $U_{b,S_1,i} < U_{b,S_j,i}$ where $j = 2...N$.

Let us analyze each of the terms in Equation 3. Determining the price that maximizes the profit $\pi^t_{S_1,i}$ is an optimization problem provided the seller estimates the buyer distributions $f_{pa}$ and $f_{pr}$ correctly. $f_{pa}(a_i)$ gives the fraction of buyers whose preferred attribute is $a_i$. When $U_{b,i,S_1} < U_{b,i,S_j}$, where $j = 2...N$, $f_{pa}(a_i) = 0$ because no buyers purchase from $S_1$ when it is not making the best offer on the buyer's preferred attribute. The distribution $f_{pa}$ is not available with the sellers and the seller must estimate it correctly to maximize profits. If the seller underestimates $f_{pa}(a_i)$ then it ends up losing revenue from the buyers that could not be correctly identified with preferred attribute $a_i$. On the other hand, if it overestimates $f_{pa}(a_i)$, it means that buyers whose preferred attribute is not $a_i$ have been incorrectly identified with preferred attribute $a_i$. The seller then ends up losing revenue from these buyers whose preferred attribute was incorrectly identified. Therefore, it is important for a profit maximizing seller to correctly estimate the distribution $f_{pa}$. Here, we have used a probabilistic prediction algorithm based on collaborative filtering that enables a seller estimate $f_{pa}$ efficiently to solve this problem in Section 4. The term $(p^t_{a_i,S_1} - p_{co}) \int f_{pr}(p^t_{a_i,S_1})dp$ in Equation 3 gives the total profit to the seller on attribute $a_i$ independent of the buyers' attribute preference distribution $f_{pa}$. Since both these terms include the current price of the item, they are dependent variables. We optimize the product of these two terms using our dynamic pricing algorithm in Section 5.

## 4. Dynamically Determining Consumer Preferences

As shown in Figure 1, a buyer compares the prices offered by the different sellers along its preferred product attribute. To make a competitive offer in response to a buyer's purchase request, a seller should be able to identify the buyer's preferred attribute and then use a dynamic pricing algorithm to offer a competitive price to the buyer on that attribute. For this, the seller should be able to estimate $f_{pa}$ - the distribution of buyer preferences over the product attributes and then use it to predict the preferred attribute of a buyer in response to the buyer's purchase request. Here, we have used collaborative filtering (CF) to enable a seller predict a buyer's preferred attribute. Collaborative filtering algorithms[5, 16, 21] collect user opinions or preferences on items of interest. A correlation method is then used to predict and recommend possible items to new or returning users based on the similarity of their interests with those of the users in the collected information. In CF algorithms the preferences of users for different items is assumed to be invariant over time. For example, in a CF-based recommender system for customers of books, a buyer who has purchased several books on physics, is recommended books on or related to physics with a high probability. Stationary user preferences allows CF algorithms to analyze the collected data in multiple dimensions (memory-based and model-based) for more robust predictions. Many CF systems also use the preference history of returning users to predict future preferences of the user with reasonable accuracy. In contrast, in our market model, customers' attribute preferences vary over time due to exogenous factors. For example, a buyer whose usual preferred attribute is price could change its preferred attribute to delivery time when an urgent delivery is required. Therefore, a seller's attribute prediction algorithm for a potential buyer has to adaptively respond to possible changes in buyer preferences. Our attribute prediction algorithm described in the next section achieves this by dynamically updating the model of the buyer attribute preferences constructed by the seller.

### 4.1. Attribute Prediction Algorithm based on Collaborative Filtering

We illustrate the algorithm for a single seller $S_1$ assuming that it is competing with the rest of the sellers in the market. The same algorithm can be used by other sellers as well. A seller constructs one buyer cluster for each product attribute. Let us suppose there are $C = \nu$ clusters maintained by the seller. A buyers with the preferred attribute $a_i$ is placed into cluster $c_i$ with a probability $w_{i,t}$ during interval $t$. The vector of probabilities over all clusters is given by $W_t = < w_{i,t} >$. These probabilities get updated dynamically in response to changes in the buyer preferences. When a buyer makes a purchase request, the prediction algorithm takes the history of $w_i$-s and outputs the predicted cluster(preferred attribute) for the buyer. When a buyer is placed in a cluster by the prediction algorithm, the current value of $W$, the buyer-id, the cluster in which the buyer is


AAMAS'04, July 19-23, 2004, New York, New York, USA.

placed and the purchase decision of the buyer are recorded by the seller in a purchase history table $P$. The $j$-th row of this table is then given by $P_j =< W_j, b_j, c_j, d_j >$.

At the end of an update interval $t$, let $p_{ci,t}$ denote the number of positive purchase decisions, and $n_{ci,t}$ denote the number of negative purchase decisions for the buyers in cluster $c_i$. If $p_{ci,t} > 0$, it means that seller $S_1$ has made the most attractive offer on attribute $a_i$ during interval $t$. At the same time, if $n_{ci,t} > 0$, it means that seller $S_1$ had incorrectly identified some buyers from another cluster into $c_i$ that made the negative purchase decisions. On the other hand, if $p_{ci,t} = 0$ it means that another seller had made a better offer than the current seller during the last update interval. If $n_{ci,t} > 0$ and/or $p_{ci,t} = 0$, the seller needs to revise the $w_i$-s during the next interval so that it can reduce the number of buyers that have been incorrectly identified. The update equation we have used for $w_i$ uses the difference between the number of buyers that have correctly and incorrectly identified. This ensures that the update amount is reinforced by the number of correct identifications and penalized by the number of incorrect identifications.

$$w_{i,t+1} = w_{i,t}^{\frac{p_{ci,t}}{p} - \frac{n_{ci,t}}{n}} \tag{4}$$

where $p = \sum_{c_i} p_{ci,t}$ and $n = \sum_{c_i} n_{ci,t}$

The preferred attribute for a buyer that has made a purchase request during interval $t + 1$ is predicted as follows:

Let $W_{t+1}$ denote the value of $W$ during interval $t + 1$ determined using Equation 4. First we find the cosine similarity between the different values of $W$ from the purchase history table:

$$sim_{j,t+1} = \frac{W_{t+1}.W_j}{|W_{t+1}| * |W_j|} \tag{5}$$

where "." denotes the dot product of the two vectors.

The cluster prediction weight $r_{c_i,t+1}$ for predicting cluster $c_i$ during interval $t + 1$ is calculated as the sum of similarity values associated with cluster $c_i$ divided by the size of (number of buyers in) cluster $c_i$ as shown below:

$$r_{c_i,t+1} = \frac{\sum_{c_i} sim_{i,t+1}}{|c_i|} \tag{6}$$

In equations 5 the similarities between the historical and current values of the $W$-vectors are determined, and, these values are used to compute a probability in 6. In accordance to the CF algorithm, these equations ensure that clusters (attributes) that are similar to each other get assigned a high probablity value while dissimilar clusters (attributes) get assigned a low probability value.

Let $R_{t+1}$ represent the normalized vector of the $r_{c_i}$-s during interval $t + 1$. The predicted cluster for a buyer that makes a quote request during interval $t+1$ is then given by: *Predict cluster $c_i$ where $c_i$ is selected probabilistically according to the probability density function given by $R_{t+1}$.*

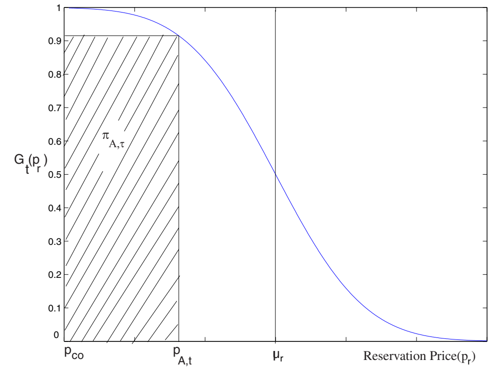


**Figure 2. Distribution function $G_t(p_r)$ gives the number of buyers willing to purchase the good at the end of interval $t$, as a function of the reservation prices of buyers. $p_r$ denotes the reservation prices of buyers for one product attribute.**

## 5. Dynamic Pricing Algorithm

The objective of the dynamic pricing algorithm of a seller's pricebot should be to determine the optimum price $p^t_{a_i,S_1}$ for every interval $t$ that maximizes the expression $(p^t_{a_i,S_1} - p_{co}) \int_{p_{co}}^{p} f_{pr}(p^t_{a_i,S_1})dp$ where, $\int_{p_{co}}^{p} f_{pr}(p^t_{a_i,S_1})dp$ represents the number of buyers purchasing during interval $t$. The dynamic pricing algorithm described here assumes that the buyer reservation price distribution $f_{pr}$ follows a standard distribution whose parameters are unknown and must be determined by the pricebot. We give an analysis for a normal distribution with mean $\mu_r$ and standard deviation $\sigma_r$ that are estimated by the seller. As before, we use seller $S_1$ for illustrating our algorithm. The calculations are shown for attribute $a_i$; similar calculations are done by the seller for each product attribute.

The values $\mu_r$ and $\sigma_r$ must be estimated with reasonable accuracy so that a seller can determine its profit maximizing price. As every buyer has a different value for its reservation price $p_{r,b}$, charging a low price to the buyers results in some buyers getting the good at a bargain price, thereby reducing the seller's profits. On the other hand, a high price might exceed the reservation price for many buyers and once again result in reduced profits to the seller.

As shown in Figure 2, $G_t(p_r) = 1 - F_t(p_r) = \int_{p^t_{a_i,S_1}}^{\infty} f_{pr}(p^t_{a_i,S_1})dp$ represents the number of buyers that are still willing to purchase the good when its price is $p^t_{a_i,S_1}$. The profit obtained from these buyers is then represented by the area of the rectangle between the current price $p^t_{a_i,S_1}$, the number of buyers $G(p^t_{a_i,S_1})$ purchasing at that price, and the two axes. The objective of the pricebot is to determine the price $p^{t+1}_{a_i,S_1}$ for the next inter-

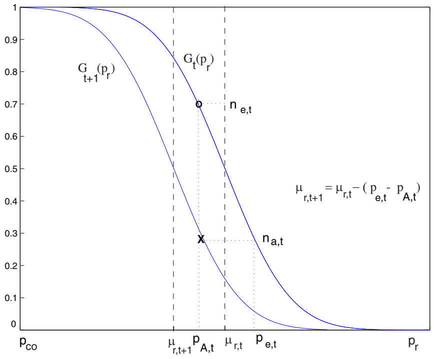


**Figure 3. The pricebot refines its estimate of $\mu_r$ using the actual and observed values of the number of customers purchasing at price $p^t_{a_i,S_1}$.**

val so that the area of this rectangle is maximized.

The pricebot starts with an initial estimate for the values $\mu_r$ and $\sigma_r$ and refines its estimate at the end of every interval as more sales data arrive, as follows:

1. Set the initial assumption of the mean of the normal distribution of the buyers' reservation price $\mu_{r,0}$ equal to the market entry price of the seller, i.e., $\mu_{r,0} = p^0_{a_i,S_1}$
2. At the end of interval $t$, determine the actual number of buyers that purchased during that interval. Let $n_{a_i,t}$ denote this number, as shown in Figure 3.
3. Determine the number of buyers that should have purchased at $p^t_{a_i,S_1}$ from the pricebot's approximation of the distribution of the buyers' reservation prices $n_{e,t} = G_t(p^t_{a_i,S_1})$, as shown in Figure 3.
4. Refine $\mu_{r,t}$ according to the equation $\mu_{r,t+1} = \mu_{r,t} - (p_{e,t} - p^t_{a_i,S_1})$ where $p_{e,t} = G^{-1}_t(n_{a,t})$, as shown in Figure 3.
5. Recalculate $G_{t+1}(p^t_{a_i,S_1})$ with the refined value $\mu_{r,t+1}$ and determine $p^{t+1}_{a_i,S_1}$ as the price that yields

$$max\{\frac{p^t_{a_i,S_1} - p_{co}}{\sigma_r} \times G_{t+1}(p^t_{a_i,S_1})\}$$

At the end of every interval, the pricebot uses steps 2 through 4 above to refine its estimate of the distribution of buyers' reservation prices and to determine $p^{t+1}_{a_i,S_1}$ for the next interval, as shown in Figure 3. As the pricebot adjusts the price being charged to buyers at successive intervals, some buyers cease to purchase as soon as $p^t_{a_i,S_1}$ exceeds their reservation price. Therefore, the value of $n_{a_i,t}$ keeps changing. Also, at each successive iteration, the pricebot's estimation of $\mu_{r,t}$ is refined. The process converges as $\mu_{r,t}$ approaches the actual value of $\mu_r$.

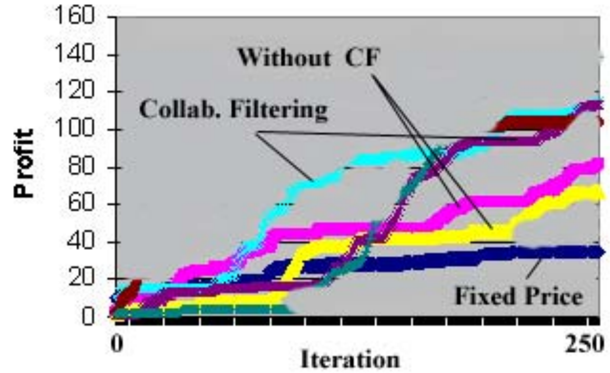


**Figure 4. Profit profile of 5 sellers using different strategies for attribute 'customer satisfaction'.**

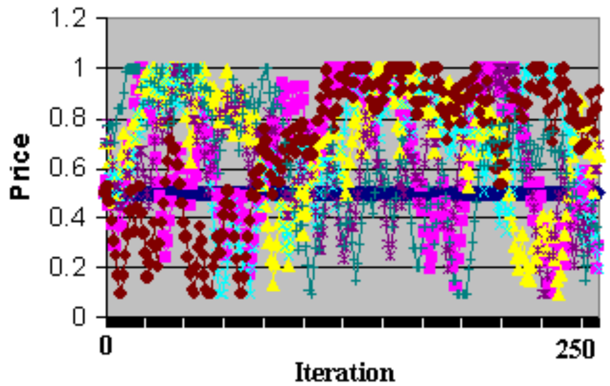


**Figure 5. Price competition between 5 sellers using different strategies for attribute 'customer satisfaction'.**

## 6. Simulation Results

The parameters used for our simulations were $B = 1000$, $S_N = 5$ ,$\nu = 5$, and $p_{co} = 0.1$. The market entry price $p^0_{a_i,S_j}$ for a seller for attribute $a_i$ was drawn randomly from the uniform distribution $U[p_{co}, 1.0]$. The price adjustment interval for sellers was taken as $\tau_s = 20$ quote requests from buyers. We selected the number of quote-requests received from buyers as the unit for measuring time to equalize differences between sellers with different response times. $f_{pr}$ was drawn from the normal distribution $N(0.7, 0.3)$. $f_{pa}$ was intialized with data from Reseller Ratings Inc.[20] - a publicly available Website that contains data on consumer responses for purchases made from various online merchants over different product attributes. The selected attributes and the initial distribution of buyers over those attributes were the following: (i) product price (30.4 percent), (ii) delivery time (29.6 percent) , (iii) customer satisfaction (30.4 percent), (iv) techinical support (5.6 percent), and product replacement (after-sales service) (4.8 percent). $f_{pa}$ was varied by randomly selecting an attribute $a_k$ and incrementing $f_{pa}(a_k)$ by 10 percent of total number of buyers $(0.1 \times 1000)$. $f_{pa}(a_i) \forall a_i \neq a_k$ was simultaneously decremented uniformly so that the total number of buyers in the market remained constant. Sellers use the following strategies to update their posted prices:

1. Fixed price seller that charges a constant price.
2. Dynamic pricing algorithm only, without CF based prediction.
3. Dynamic pricing algorithm and CF based prediction.

Figure 4 shows the profit profile of five sellers that use different pricing strategies for the product attribute 'customer satisfaction'. The corresponding price profile shown in Figure 5 illustrates that market prices fluctuate consistently due to the competition between sellers. The fixed price seller ends up with the least profit (about 6 percent of the total profit in the market) as it can only attract buyers when other sellers under-cut each other while competing and end up charging a price lower than the fixed price seller. The two sellers that use the dynamic pricing algorithm only, without CF based prediction to identify buyer preferences, obtain about 18 percent each of the total profit in the market. The two sellers using the dynamic pricing algorithm with CF based prediction perform the best in this setting and each obtains about 29 percent of the total market profit. The CF algorithm enables these sellers to rapidly estimate the changes in the value of $f_{pa}$ and thereby obtain more profits according to Equation 3.

An interesting result was obtained for attributes such as technical support and product replacement as shown in Figure 6 : the sellers with CF based prediction were outperformed by sellers that did not use the CF algorithm. This phenomenon can be explained by the fact that the CF algorithm uses similarities among attributes to predict the preferred attribute. If the buyer's preferred attribute changes to a 'similar' attribute with approximately the same value in $f_{pa}$, $W_t$ does not change significantly and the similar attribute gets predicted with a high probability. However, if the buyer's preferred attribute changes to a 'dissimilar' attribute that has a siginificantly different value in $f_{pa}$, CF predicts the dissimilar attribute with a low probability. The attribute gets selected after several iterations during which the probability associated with the attribute slowly

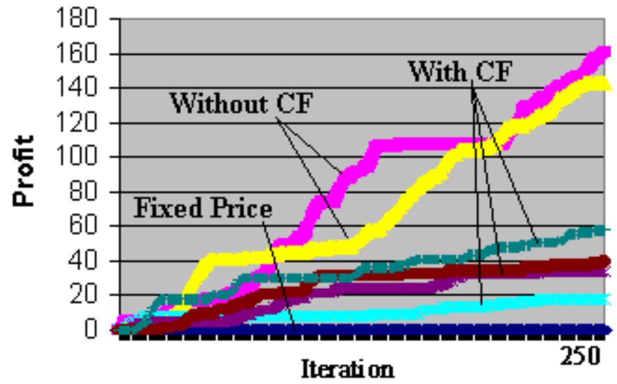


**Figure 6. Profit profile of sellers using different strategies for attribute 'product replacement'.**

increases. When a customer changes its preferred attribute from an attribute with a high value in the initial $f_{pa}$ distribution (e.g.: customer satisfaction) to an attribute with a low value in the initial distribution (e.g: product replacement), the CF algorithm takes several iterations to update the probability values and predict the change in the preferred attribute. This scenario occurs in our simulations because buyers starting off with preferred attribute as price, customer satisfaction or delivery time, change their preferred attribute to product replacement or technical support due to the random variations simulated in $f_{pa}$. Therefore, in these cases, sellers using the CF algorithm are outperformed by sellers that are not using it.

## 7. Related Work

Over the past few years online dynamic pricing has stimulated considerable interest in commercial and research communities. Increased profits and rapidly clearing inventories resulting from efficient pricing has encouraged the development of software pricing tools including Azerity[2] and Live Exchange[18]. Several analytical models have been developed for dynamic pricing in online economies. Kephart et al[15, 23] have viewed the price setting problem as a one-shot game. In these analyses, information about market parameters such as the distribution of buyer reservation prices, and competitor's prices and profits are assumed to be available with a seller. In contrast, the dynamic pricing algorithm described here uses collaborative filtering to determine buyer preferences and does not use prior knowledge about market parameters. Brooks et al[6] have applied neural networks and genetic algorithms to enable a monopolist seller learn buyers' pricing schedule and optimize profits in a monopolist market. In contrast, here, we have modeled a competitive market where profit dynamics of a seller are affected by buyer preferences as well as prices and profits of competing sellers.

Exchange of products differentiated on multiple attributes has been studied in an auction setting in[11, 19, 22]. In auction setting buyer preferences are assumed to be stationary and sellers compete to offer the best bid. Also, sellers bid sequentially and the objective of a seller is to determine an optimum bid while observing the bids that arrive before it. In our model, sellers make their prices adjustments completely asynchronously. Our algorithm does not involve any ordering of price adjustments between sellers and can operate successfully without relying on competitors' price or profit information.

## 8. Conclusions and Future Work

The collaborative filtering algorithm described in this paper is our first attempt at enabling online sellers to determine buyer preferences over multiple product attributes and update the posted product prices efficiently in a competitive market. We are currently investigating more powerful learning techniques such as Q-learning and multi-objective evolutionary algorithms (MOEA) as a mechanism to enable sellers search the profit landscape more efficiently. Implementing complex techniques involve a tradeoff between rapidity and accuracy of the learning algorithm and we suspect that a naive, but fast learning technique might ultimately compare favorably against a complex and accurate, but slow learning mechanism in a dynamic environment like a competitive online market.

An interesting scenario emerges when buyers' purchase preferences are dependent on the prices being charged by sellers in the market. In such a scenario, a seller can attempt not only to learn the temporally varying buyers' purchase preference distribution but also the function that determines the variation in that distribution. Probabilistic algorithms such as Hidden Markov Models and moving target functions that estimate the dependence between temporally varying functions could be possibly applied in such an environment. Although dynamic pricing has been tested by online sellers it is yet to be adopted widely in e-commerce. We envisage that the directions described in this paper will encourage techniques and mechanisms for making online dynamic pricing an enabling technology for e-commerce in the future.

## References


[1] Amazon Inc., URL http://www.amazon.com

[2] Azerity Inc. URL http://www.azerity.com/

[3] BizRate Inc., URL http://www.bizrate.com

[4] G. Bodhi et al, "Amazon.com and Dynamic Pricing," URL http://faculty-gsb.stanford.edu/groseclose/Papers/amazon.pdf

[5] J. Breese, D. Heckerman, C. Kadie, "Emperical Analysis of Predictve Algorithms for Collaborative Filtering," Proc. of the 14th Conference on Uncertainty in Artificial Intelligence, 1998, pp. 43-52.

[6] C. Brooks, S. Fay, R. Das, J. MacKie-Mason, J. Kephart, E. Durfee, "Automated strategy searches in an electronic goods market: learning and complex price schedules," Proc. of 1st ACM Conference on E-Commerce 1999, pp. 31-40.

[7] C. Brooks, R. Gazzale, J. MacKie-Mason, E. Durfee, "Improving learning performance by applying economic knowledge," Proc. of the 3rd ACM Conference on Electronic Commerce 2003, pp. 252-253.

[8] J. Brown and A. Goolsbee, "Does the Internet Make Markets More Competitive?," NBER Working Papers 7996, National Bureau of Economic Research, 2000.

[9] P. Dasgupta and R. Das, "Dynamic Service Pricing for Brokers in a Multi-Agent Economy", Proc. of the Third International Conference for Multi-Agent Systems(ICMAS), Boston, MA, July 4-9, 2000, pp. 375-376.

[10] P. Dasgupta, P. M. Melliar-Smith, "Dynamic Consumer Profiling and Tiered Pricing Using Software Agents", Journal of Electronic Commerce Research, Special Issue on Agents for E-business Automation, Kluwer Academic Publishers, The Netherlands, vol.3, no 3-4, July-October 2003, pp. 277-296.

[11] E. David, R. Azoulay-Schwartz and S. Kraus, "Bidders' strategy for multi-attribute sequential english auction with a deadline," Proc. of the 2nd International Conference on Autonomous Agents and Multi-Agent Systems, 2003, pp. 457-464.

[12] A. Greenwald and J. Kephart, "Probabilistic Pricebots," Proc. of the Fifth International Conference on Autonomous Agents, 2001, pp. 560-567.

[13] J. Hirschleifer and J. Hirschleifer, "Price Theory and Applications," Addison-Wesley, 1998.

[14] R. Keeney and H. Raiffa, "Decisions with Multiple Objectives : Preferences and Value Trade-offs," Cambridge University Press, 2002.

[15] J. Kephart, J. Hanson and A. Greenwald, "Dynamic Pricing by Software Agents," Computer Networks, vol 32, no 6, 2000, pp. 731-752.

[16] J. Kleinberg and M. Sandler, "Convergent Algorithms for Collaborative Filtering," Proc. of the 4th ACM Conference on E-commerce, 2003, pp. 1-10.

[17] Mysimon Inc., URL http://www.mysimon.com

[18] Moai Inc., URL http://www.moai.com

[19] D. Parkes, "Iterative Multiattribute Vickrey Auctions," (submitted for publication).

[20] ResellerRatings Inc., URL http://www.resellerratings.com

[21] B. Sarwar, G. Karypis, J. Konstan, and J. Reidl, "Item-based collaborative filtering recommendation algorithms," Proc. of Tenth International WWW conference, 2001, pp. 285-295.

[22] T. Sandholm, S. Suri, A. Gilpin, and D. Levine, "Winner determination in combinatorial auction generalizations," Proceeding of the 1st International Conference on Autonomous Agents and Multi-Agent Systems, 2002, pp. 69-76.

[23] G. Tesauro and J. Kephart, "Pricing in agent economies using multi-agent Q-learning," Autonomous Agents and Multi-Agent Systems, vol 5, no 3, 2002, pp. 289-304.





AAMAS'04, July 19-23, 2004, New York, New York, USA.