Dynamic Pricing with Early Cancellation and Resale

Kwan-Ang An

Thesis submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in fulfillment of the requirements for the degree of

Master of Science

in

Industrial and Systems Engineering

Kyle Y. Lin, Chair

Ebru K. Bish

Joel A. Nachlas

January 24, 2003

Blacksburg, Virginia

Keywords: Revenue Management; Dynamic Pricing: Cancellation and Refund;

Stochastic Demand; Perishable Asset

Dynamic Pricing with
Early Cancellation and Resale

Kwan-Ang An

(ABSTRACT)

We consider a continuous time dynamic pricing model where a seller needs to sell a single item over a finite time horizon. Customers arrive in accordance with a Poisson process. Upon arrival, a customer either purchases the item if the posted price is lower than his/her reservation price, or leaves empty-handed. After purchasing the item, some customers, however, will return the item to the seller at an exponential rate for a full refund. We assume that a returned item is in mint condition and the seller can resell it to future customers. The objective of the seller is to dynamically adjust the price in order to maximize the expected total revenue when the sale horizon ends. We formulate the dynamic pricing problem as a dynamic programming model and derive the structural properties of the optimal policy and the optimal value function. For cases in which the customer's reservation price is exponentially distributed, we derive the optimal policy in a closed form. For general reservation price distribution, we consider an approximation of the original model by discretizing both time and the allowable price set. We then present an algorithm for numerically computing the optimal policy in this discrete time model. Numerical examples show that if the discrete price set is carefully chosen, the expected total revenue is nearly the same as that when the allowable price set is continuous.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Dynamic pricing is the practice of changing the price of a product in real time in order to maximize profits. In particular, it is vitally important for those industries such as airlines, hotels, cruise lines, car rental companies, and theaters. In those industries, sellers have various restrictions on their selling activity; they need to sell a fixed stock of items to customers within a given finite time horizon and, during the horizon, the stock cannot be replenished and the unsold items have little salvage value after the sale horizon ends. The problem is that the items need to be sold over a given finite time horizon, which is typically very short. For example, an airline has only a few months to sell tickets for a certain flight and a fashion retailer has to sell summer clothes targeted only for a single summer season. To make the matter worse, it is difficult to forecast the demand accurately. Therefore, sellers have an incentive to use the price to control the demand in order to maximize their total revenue. In the airline example, they post a higher price if the sale goes well initially since they can reliably expect that late arriving customers are substantially willing to pay the premium for tickets, while they post a lower price if the sale does not go well in order to induce customer demand. This is the principal reason why the sellers in the aforementioned industries price items differently for different customers.

The dynamic pricing problem obliges the seller to experience the trade-off between the *yield loss* and *spoilage loss* (Bialogorsky et al. (1999)); the former is the revenue loss associated with selling an item at a low price and losing the opportunity to sell it at a high price and the latter is the revenue loss resulting from waiting in vain for higher price customers and losing the opportunity to sell the item to lower price customers. A dynamic pricing approach mitigates these opportunity losses by aiding the sellers in determining posted prices in a timely manner. It is literally the tactical approach for pricing items dynamically differently during the sale horizon to capture customer demand (customer willingness to pay) and maximize profits. It has been recognized as an effective tool for solving the revenue management problem in the aforementioned industries.

1.1 Literature Review

Dynamic pricing has been studied and developed extensively in the literature of the revenue management (previously called the yield management). Since the revenue management itself originates and have developed from the activity of airlines selling discounted price tickets, which started in the 1970's, many of the studies in the literature, especially the earlier studies, concentrate on what is called the airline seat inventory control problem; this is the problem of knowing when to reject the request for discounted seats from lower fare class customers to reserve some seats for higher fare class customers in order to maximize profits. Brumelle and McGill (1993) consider the problem and find the optimal booking limit policies for multiple fare classes for a single leg flight, with a commonly used assumption that lower fare class customers arrive for tickets earlier than higher fare class customers. Robinson (1995) also studies the single leg flight problem, but does not assume the previous assumption. In particular, the models proposed in the earlier studies are static when customers arrive strictly according to the order of the fare classes (from lower to higher ones) in that the demand of customers can be considered as random variables independent of classes.

In contrast to the above static model, recent works are categorized as a dynamic model in which customer demand is viewed as a stochastic process using primarily a dynamic programming approach. Lee and Hersh (1993) study the airline seat inventory control problem using the dynamic programming approach without the aforementioned assumption and also study group booking problems. Gallego and Ryzin (1994) study the continuous time model where the customer arrival rate is defined as a function of a posted price. They proposed the two structural monotonic properties of the optimal policy: the inventory monotonic property and the time monotonic property. Among the more recent works, Feng and Xiao (2000) consider the seat inventory control problem in a continuous time framework with multiple fare classes in which prices are assumed to change monotonically only at pre-specified periods, and they derive the optimal pricing policy directly from the continuous model. Feng and

Xiao (2001) also study the seat inventory control problem in which customers are assumed to travel from multiple origins through one hub to their final destinations. Zhao and Zheng (2000) study a more general case where customers arrive according to a non-homogeneous Poisson process, and the distribution of the customer reservation price changes over time.

There are those papers proposing a heuristic that discretizes a continuous time horizon and solves a continuous time model efficiently. Bitran and Mondscheine (1997) propose the dynamic pricing model with the periodic pricing review where posted prices are changed at predetermined review times. They use the optimal policy from a continuous time model as an upper bound and show that the expected revenue depends on the number of the review times. Lin (2001) proposes a sequential dynamic pricing model where customers arrive sequentially and a seller knows the distribution of the total number of future customers. He derives the optimal policy for each of the cases when the total number of future customers is distributed such as deterministic, geometric, and bounded. He applies those results to the continuous time model where customers arrive according to a non-homogeneous Poisson process, and shows that the numerical results are nearly optimal.

There are also many papers exclusively dealing with the problem of overbooking, cancellation, no-shows and so forth. Chatwin (1998) studies the airline overbooking problem with a single fare class for a single-leg flight and applies the results to the discount seat allocation problem. He subsequently studies, in Chatwin (1999), the continuous time overbooking problem and proposes that the piecewise constant booking limit policy is optimal provided that fares and refunds are a piecewise constant function of the remaining time to flight.

Subramanian, Stidam, and Lautenbacher (1999) examine the airline seat allocation problem by considering the events of cancellations, no-shows, and overbooking, and they formulate it as the optimal control problem of customer arrivals into a queueing system. Biyalogorsky et al. (1999) study the effect of *overselling with opportunistic cancellation* on the expected revenue and propose that when overselling is permitted, a seller should oversell items on purpose if high fare customers arrive, in order to maximize profits.

Beyond the studies concerning the airline industry revenue management, Bitran and Gilbert (1996) study a dynamic pricing approach for the hotel reservation process and develop heuristic algorithms to accept reservations. They emphasize that it is important to take into account the events of cancellations and no-shows in order to realize potential revenue.

1.2 Overview of the Research

In this research, we consider a dynamic pricing model with a focus on the event of cancellation. Upon cancellation, a seller receives a returned item, issues a refund, and sells the item later to another customer. Cancellation does not require a seller to make any decisions, but provides the seller with potential resources during the time horizon. Hence, the concern of the seller is how to dynamically and optimally price items considering the possibility of cancellation and its resultant resale in order to maximize his/her total revenue.

In this paper, we consider a dynamic pricing model where a seller has a single item to sell over a given time horizon. The item is perishable and cannot be replenished during the horizon, and there is no salvage cost if it is unsold at the end of the sale. Customers arrive in accordance with a Poisson process with a known rate. An arriving customer buys the item if his/her reservation price is greater than a price posted by the seller. The customer reservation price is independent and identically distributed (i.i.d.) over time, and it is known to the seller. The customer who has bought the item holds it for an exponential amount of time with a known rate; that is, the time interval of cancellation is exponentially distributed with a known rate. Upon cancellation, the seller receives a returned item and issues a full refund. The task of the seller is to choose the optimal posted price in order to maximize total revenue at the end of the sale horizon.

Chapter 2

Continuous Time Model

In this section, we formulate the continuous time dynamic pricing model. Since only one item is available, a seller needs to post a price only if the item is in inventory when a customer arrives. The seller's concern, thus, is with the time available to sell the item. We say that the seller is in state t if he/she holds the item and has t time remaining. The time horizon is described in a reversed order starting at time T and ending at time 0. Customers arrive in accordance with a Poisson process with rate λ and the reservation prices of customers are i.i.d. with the cumulative distribution function F. We assume that F strictly increases in a posted price p. A customer buys an item if his/her reservation price exceeds a posted price, and the seller receives the instantaneous revenue p from a customer with the probability $\overline{F}(p)$. The customer who has bought the item makes cancellation according to the exponential distribution with rate μ. Upon cancellation, the seller issues a full refund. Overbooking and no-shows are not permitted. The task of the seller is to search for the optimal posted price p^* that maximizes the expected total revenue at the end of the sale horizon.

2.1 Formulation for Optimal Policy

Now we define the optimal value function $V(t)$ as the expected revenue from time t and onward when a customer arrives at time t if a seller has the item and uses the optimal policy. Now let $U(t)$ denote the expected revenue from time t and onward if the seller has the item at time t and uses the optimal policy. (Note that the difference between $U(t)$ and $V(t)$ is only whether there is a customer arrival at time t or not.) By conditioning on the time of the next customer arrival, $U(t)$ is determined as,

$$U(t) = \int_0^t V(t-x)\lambda e^{-\lambda x} dx \tag{2.1}$$

Now suppose that a customer arrives at time t. $V(t)$ is determined by conditioning on whether the arriving customer buys the item or not. If the customer does not buy the item (or rejects a posted price p), with probability $F(p)$, there is no revenue from the customer at time t, which is equivalent to no customer arriving at time t. Hence, the expected revenue from t and onward is now given by $U(t)$. In contrast, if the customer buys the item (or accepts a posted price p), the expected revenue from t and onward is determined by the summation of the expected revenue from the arriving customer (from time t and onward) and the expected revenue after cancellation occurs up until the end of the sale horizon. Since the interval of cancellations is now exponentially distributed with rate μ, the probability that no cancellation occurs within the remaining time is given as $e^{-\mu t}$. Therefore, the expected revenue from the customer from time t and onward is given as $pe^{-\mu t}$. In addition, conditioning on the possible time at which cancellation occurs after time t, we have,

$$\int_0^t U(t-y)\mu e^{-\mu y} dy$$

as the expected revenue after cancellation occurs up until the end of the sale horizon in the future. As a result, $V(t)$ is formulated as,

$$V(t) = \max_p \Big[F(p)U(t) + \overline{F}(p) \Big\{ pe^{-\mu t} + \int_0^t U(t-y)\mu e^{-\mu y} dy \Big\} \Big] \tag{2.2}$$

Therefore, the optimal pricing policy at time t is derived as the maximizer to the function $V(t)$. The optimal policy is obviously a continuous function of time t, namely, $p^*(t)$.

Remark 1 The boundary condition of the optimal policy $p^*(0)$ is derived as the maximizer to $p \cdot \overline{F}(p)$ since there is no customer after time 0 and therefore a seller only needs to maximize the instantaneous revenue from a customer at time 0. It is also important to be aware that it is never optimal to post any price less than $p^*(0)$ at any given time. Since posting such a price results in lower instantaneous expected revenue from a customer and a higher probability that the item is sold out, the expected revenue with $\forall p < p^*(0)$ cannot be higher than that with $p^*(0)$. Therefore, it is useless to post any price less than $p^*(0)$. (Refer to Lemma 2.1 in Lin (2001))

2.2 Structural Property

Now we solve the optimal value function $V(t)$ along with $U(t)$ and derive the structural properties of the optimal value function and the optimal policy.

First of all, Eq (2.1) is transformed as follows by replacing $t-x$ with a,

$$\begin{aligned} U(t) &= \int_0^t V(a)\lambda e^{-\lambda(t-a)}da \\ &= \lambda e^{-\lambda t}\int_0^t V(a)e^{\lambda a}da \end{aligned}$$

Differentiating both sides of the equation,

$$\begin{aligned} U'(t) &= -\lambda^2 e^{-\lambda t}\int_0^t V(a)e^{\lambda a}da + \lambda e^{-\lambda t}\frac{d}{dt}\int_0^t V(a)e^{\lambda a}da \\ &= -\lambda^2 e^{-\lambda t}\int_0^t V(a)e^{\lambda a}da + \lambda e^{-\lambda t}V(t)e^{\lambda t} \\ &= \lambda(V(t) - \lambda e^{-\lambda t}\int_0^t V(a)e^{\lambda a}da) \\ &= \lambda(V(t) - U(t)) \end{aligned} \tag{2.3}$$

Arranging Eq (2.2), we have

$$V(t) = U(t) + \max_p\Big[\overline{F}(p)\Big\{pe^{-\mu t} + \int_0^t U(t-y)\mu e^{-\mu y}dy - U(t)\Big\}\Big] \tag{2.4}$$

Now define $g(p)$ as,

$$g(p) = \overline{F}(p)\Big\{pe^{-\mu t} + \int_0^t U(t-y)\mu e^{-\mu y}dy - U(t)\Big\}$$

and also define the hazard rate function of F as $r(p) = f(p)/\overline{F}(p)$ where f is the probability density function of F.

The first derivative of $g(p)$ is,

$$g'(p) = -f(p)\Big(pe^{-\mu t} + \int_0^t U(t-y)\mu e^{-\mu y} dy - U(t)\Big) + \overline{F}(p)e^{-\mu t}$$

Denote the maximizer for $g(p)$ by $p^*(t)$. Since $g'(p) = 0$ when $p = p^*(t)$,

$$-f(p^*(t))\Big(p^*(t)e^{-\mu t} + \int_0^t U(t-y)\mu e^{-\mu y} dy - U(t)\Big) + \overline{F}(p^*(t))e^{-\mu t} \quad = \quad 0$$

Hence,

$$\int_0^t U(t-y)\mu e^{-\mu y} dy = U(t) + \left(\frac{1}{r(p^*(t))} - p^*(t)\right)e^{-\mu t} \tag{2.5}$$

The following proposition is for the uniqueness of the optimal policy associated with Eq (2.5).

Proposition 1 *If the customer reservation price has an increasing failure rate (IFR), then the optimal price $p^*(t)$ is uniquely determined at any time t.*

Proof: Since $1/r(p^*(t))$ decreases in $p^*(t)$ if the reservation price has an IFR, $\Big(1/r(p^*(t)) - p^*(t)\Big)$ monotonically decreases in $p^*(t)$. Therefore, Eq (2.5) has a unique solution with respect to $p^*(t)$ if the reservation price has an IFR. $\Box$

Given that the reservation price has an IFR,

$$\frac{dr(p(t))}{dp(t)} = \frac{d}{dp(t)}\left(\frac{f(p(t))}{\overline{F}(p(t))}\right) = \frac{\overline{F}(p(t))\frac{df(p(t))}{dp(t)} + \{f(p(t))\}^2}{\{\overline{F}(p(t))\}^2} \geq 0$$

and thus,

$$\frac{df(p(t))}{dp(t)} \geq 0$$

Hence, for the second derivative of $g(p(t))$ at $p^*(t))$,

$$g''(p^*(t)) = \Big(- \frac{1}{r(p^*(t))}\frac{df(p^*(t))}{dp^*(t)} - 2f(p^*(t))\Big)e^{-\mu t} \leq 0$$

Therefore, it is proved that $g(p(t))$ is concave in the neighborhood of $p^*(t)$.

In addition, from Eq (2.5),

$$\begin{aligned} g(p^*(t)) &= \overline{F}(p^*(t))\Big\{p^*(t)e^{-\mu t} + \int_0^t U(t-y)\mu e^{-\mu y}dy - U(t)\Big\} \\ &= \overline{F}(p^*(t))\Big\{p^*(t)e^{-\mu t} + U(t) + \Big(\frac{1}{r(p^*(t))} - p^*(t)\Big)e^{-\mu t} - U(t)\Big\} \\ &= \frac{\overline{F}(p^*(t))}{r(p^*(t))} \cdot e^{-\mu t} \end{aligned}$$

From Eq (2.4),

$$\begin{aligned} V(t) &= U(t) + \max_p g(p) \\ &= U(t) + g(p^*(t)) \\ &= U(t) + \frac{\overline{F}(p^*(t))}{r(p^*(t))} \cdot e^{-\mu t} \end{aligned}$$

Hence, from Eq (2.3),

$$U'(t) = \lambda e^{-\mu t}\Big(\frac{\overline{F}(p^*(t))}{r(p^*(t))}\Big) \geq 0 \qquad (2.6)$$

Remark 2 Supposing that Seller A has more available time to sell the item than Seller B, Seller A will be able to yield at least as much revenue as that of Seller B since Seller A can exactly imitate the optimal policy of Seller B until his/her business is closed, whether or not cancellation is taken into account in the process. Therefore, it is no wonder that the expected revenue that a seller yields with the item increases in the available time to sell it even when cancellation is taken into account in the process.

Furthermore, the left-hand side of Eq (2.5) is transformed by replacing $t-y$ with z as follows.

$$\int_0^t U(t-y)\mu e^{-\mu y}dy = e^{-\mu t}\int_0^t U(z)\mu e^{\mu z}dz$$

and hence,

$$\int_0^t U(z)\mu e^{\mu z}dz = U(t)e^{\mu t} + \left(\frac{1}{r(p^*(t))} - p^*(t)\right)$$

Differentiating both sides of the above equation with respect to t,

$$U(t)\mu e^{\mu t} \quad = \quad U'(t)e^{\mu t} + \mu e^{\mu t}U(t) + \frac{d}{dt}\left(\frac{1}{r(p^*(t))}\right) - \frac{dp^*(t)}{dt}$$

Hence,

$$\frac{dp^*(t)}{dt} = U'(t)e^{\mu t} + \frac{d}{dt}\left(\frac{1}{r(p^*(t))}\right) \tag{2.7}$$

From Eqs (2.6) and (2.7), we finally have the following.

$$\frac{dp^*(t)}{dt} = \lambda \cdot \frac{\overline{F}(p^*(t))}{r(p^*(t))} + \frac{d}{dt}\left(\frac{1}{r(p^*(t))}\right) \tag{2.8}$$

Therefore, the optimal policy $p(t)$ is derived once the distribution of the reservation price is known. In addition, it is clear that for any given time t, $p(t)$ increases in the arrival rate λ, regardless of the cancellation rate μ. From our perspective, no matter how frequently a customer cancels an item, it would be reasonable to post a higher price for each arriving customer if there is the higher demand (i.e., the higher arrival) for an item. In short, since it never hurts a seller no matter how many customers arrive for the item, for any given t,

$p(t)$ is higher as λ is higher.

Furthermore, from Eq (2.8), we present the proposition for the time-monotonic property of the optimal policy.

Proposition 2 *If the customer reservation price has an IFR, then the optimal price increases in the remaining time t.*

Proof: From Eq (2.8),

$$\begin{aligned}\frac{dp^*(t)}{dt} &= \lambda \cdot \frac{\overline{F}(p^*(t))}{r(p^*(t))} + \frac{d}{dt}\left(\frac{1}{r(p^*(t))}\right) \\ &= \lambda \cdot \frac{\overline{F}(p^*(t))}{r(p^*(t))} + \left(- \frac{1}{\{r(p^*(t))\}^2} \cdot \frac{dr(p^*(t))}{dt} \cdot \frac{dp^*(t)}{dt}\right)\end{aligned}$$

Hence,

$$\left(1 + \frac{1}{\{r(p^*(t))\}^2} \frac{dr(p^*(t))}{dt}\right)\frac{dp^*(t)}{dt} = \lambda \cdot \frac{\overline{F}(p^*(t))}{r(p^*(t))}$$

Clearly, if the reservation price has an IFR, we have $dr(p^*(t))/dt \geq 0$ and then $dp^*(t)/dt \geq 0$.□

Remark 3 In Proposition 2, the IFR of the reservation price is a restricted condition since $dp^*(t)/dt \geq 0$ still holds for some $dr(p^*(t))/dt \leq 0$ (DFR). However, $dr(p^*(t))/dt \geq 0$ would be more reasonable from a practical point of view. In practice, for a fixed number of items, a seller posts higher prices earlier in the time horizon and lower prices later. It is then typical that the higher the posted price, the more critical customers feel about the increase in price, since they are more reluctant to pay the premium for a higher valued item; the IFR of the reservation price would be explicitly appropriate for this situation. We remain with the IFR for the monotonic property of the optimal price for this reason.

2.3 Exponential Reservation Price Distribution

Now we implement the continuous time model for the case in which the customer reservation price is exponentially distributed. First of all, we derive the optimal policy $p(t)$ and the optimal value function $U(t)$ and $V(t)$. We then implement the numerical experiment to heuristically analyze the optimal expected total revenue with the optimal policy.

Suppose that the distribution of the reservation price is exponential with rate α. From Eq (2.8),

$$\begin{aligned} \frac{dp^*(t)}{dt} &= \lambda \cdot \frac{e^{-\alpha p^*(t)}}{\alpha} + \frac{d}{dt}\left(\frac{1}{\alpha}\right) \\ &= \frac{\lambda}{\alpha} \cdot e^{-\alpha p^*(t)} \end{aligned}$$

Hence,

$$e^{\alpha p^*(t)} \cdot \frac{dp^*(t)}{dt} = \frac{\lambda}{\alpha}$$

Solving the above differential equation with respect to t, we have,

$$\begin{aligned} \int e^{\alpha p^*(t)} \cdot \frac{dp^*(t)}{dt} dt &= \int \frac{\lambda}{\alpha} dt \\ \Rightarrow e^{\alpha p^*(t)} &= \lambda t + C \end{aligned}$$

Since the optimal pricing policy at time 0 is $1/\alpha$, given by $\max_p (p\overline{F}(p)) = \max_p (p \cdot e^{-\alpha p})$, we have $C = e$ and so,

$$\begin{aligned} e^{\alpha p^*(t)} &= \lambda t + e \\ \Rightarrow \alpha p^*(t) &= \log(\lambda t + e) \end{aligned}$$

Therefore, the optimal policy $p^*(t)$ is,

$$p^*(t) = \frac{1}{\alpha} \cdot \log(\lambda t + e)$$

Now the optimal policy monotonically increases in the remaining time t (i.e., the time-monotonic property) and goes to infinity as t goes to infinity, which is interpreted as the longer the available time to sell the item, the higher the prices posted by a seller. This is intuitively understandable since it never hurts the seller no matter how many customers arrive for the item. Specifically, the more customers arrive with interest in the item, the more likely it is for the seller to yield higher revenue since customer arrival does not cost them anything at all.

We then have the following by plugging $p^*(t)$ into Eq (2.6),

$$\begin{aligned} U'(t) &= \lambda e^{-\mu t} \cdot \frac{\overline{F}(p^*(t))}{r(p^*(t))} \\ &= \frac{\lambda}{\alpha} \cdot e^{-\alpha p^*(t) + \mu t} \\ &= \frac{\lambda}{\alpha} \cdot e^{-\log(\lambda t + e)} \cdot e^{-\mu t} \\ &= \frac{\lambda}{\alpha(\lambda t + e)} \cdot e^{-\mu t} \end{aligned}$$

Then,

$$U(t) = \frac{\lambda}{\alpha} \cdot \int_0^t \frac{e^{-\mu u}}{\lambda u + e} du$$

and so,

$$\begin{aligned} V(t) &= U(t) + \frac{\overline{F}(p^*(t))}{r(p^*(t))} \cdot e^{-\mu t} \\ &= \frac{\lambda}{\alpha} \cdot \int_0^t \frac{e^{-\mu u}}{\lambda u + e} du + \frac{e^{-\alpha p^*(t)}}{\alpha} \cdot e^{-\mu t} \\ &= \frac{1}{\alpha} \cdot \left(\lambda \cdot \int_0^t \frac{e^{-\mu u}}{\lambda u + e} du + \frac{e^{-\mu t}}{\lambda t + e} \right) \end{aligned}$$

Remark 4 We can compute the function $U(t)$ and $V(t)$ by using the identity:

$$\int_b^t \frac{e^{ax}}{x} dx = \log|x| + \frac{ax}{1 \cdot 1!} + \frac{a^2 x^2}{2 \cdot 2!} + \cdots + \frac{a^n x^n}{n \cdot n!} + \cdots$$

Figure 2.1 shows the plots of $p^*(t)$, $U(t)$, and $V(t)$ respectively when the length of the time horizon $T = 1$ (e.g., interpreted as 1 year); the mean of the reservation price is \$500, that is, $\alpha = 0.002$; the rate of the cancellation $\mu = 5$ (e.g, in a unit of month, the average length of the interval of the cancellation is 0.2 months = about 6 days); and the customer arrival rate $\lambda = 50$. It should be noted that $U(t)$ and $V(t)$ are converged as $\lim_{t\to\infty} V(t) = \lim_{t\to\infty} U(t)$ as the available time to sell the item t goes to infinity, which implies that the effect of having an arriving customer is negligible when the available time is sufficiently large. This is reasonable when seen in the following way: since it is already known that $p^*(t)$ increases in the remaining time t and $\overline{F}(p)$ decreases in a posted price p, $\overline{F}(p)$ decreases in the remaining time t. Now it would be reasonable to assume that the rate of the change in $\overline{F}(p)$ with respect to p is higher than that of p itself since, in a practical sense, the higher the prices posted for an item, the lower significantly the probability that customers buy it. In particular, when the posted prices are sufficiently high, the probability would be almost equal to 0. For this reason, the instantaneous revenue from a customer would be converged to 0 as the posted price goes to infinity, that is, the function $p \cdot \overline{F}(p)$ decreases in p and converges to 0 as p goes to infinity. Therefore, $p^*(t) \cdot \overline{F}(p^*(t))$ could be assumed to decrease in t and converge to 0 as t goes to infinity. This is true as long as the instantaneous revenue $p^*(t) \cdot \overline{F}(p^*(t))$ goes to 0 when $p^*(t)$ goes to infinity as t goes to infinity, regardless of the distribution of the reservation price. (Another interpretation for $\lim_{t\to\infty} V(t) = \lim_{t\to\infty} U(t)$ could be that customer arrivals do not mean anything when a seller has a sufficiently large amount of time for sale since the more time remaining, the higher the probability that a customer cancels the item.)

Now, as previously algebraically derived, $U(t)$ monotonically increases in the remaining time t in general as seen in Figure 2.1. In the current example, $V(t)$ also monotonically increases in t. However, this is not necessarily true. It actually changes its behavior depending on the ratio of the arrival rate to the cancellation rate, namely λ/μ.

Figures 2.2 and 2.3 show the expected total revenue $V(1)$ regarding the arrival rate λ and the cancellation rate μ respectively. It is clear that the expected revenue is higher when the ratio λ/μ is in the higher range. In this case, a seller has more chances to sell the item and fewer chances to experience cancellation since there are more customers arriving with interest in the item, which is a more frequent event than cancellation. In other words, the customer demand for the item is higher and hence the seller is more likely to sell out of the

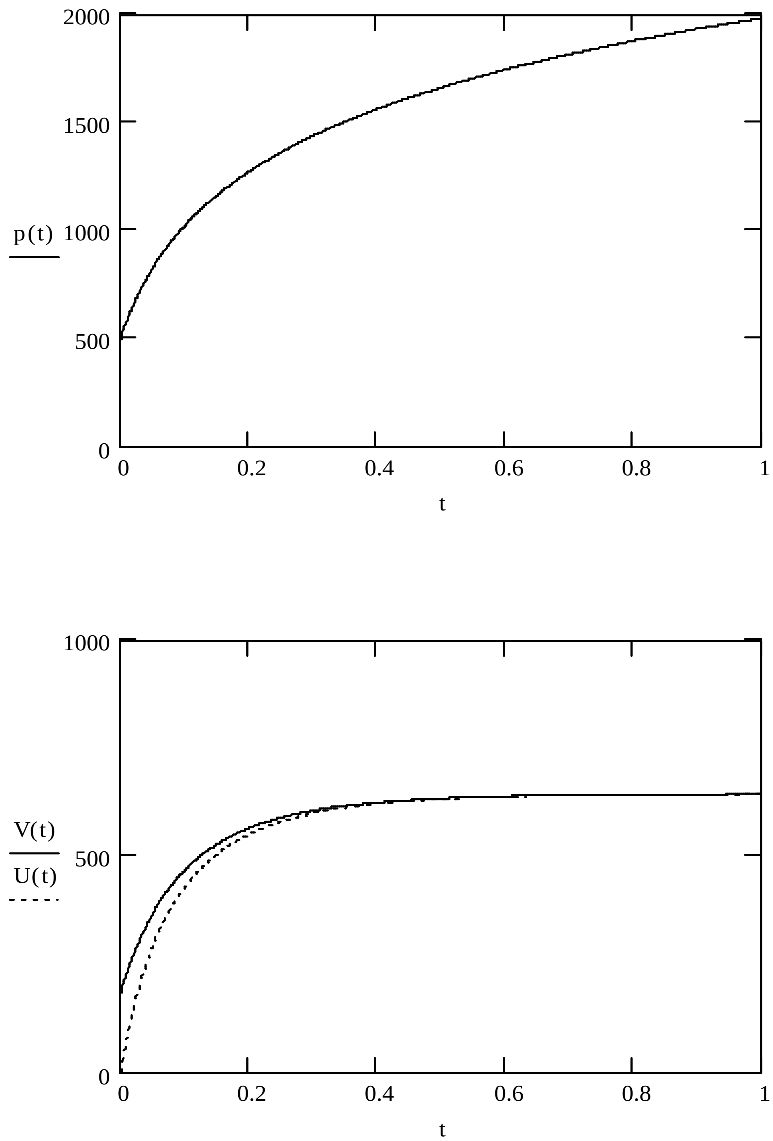

Figure 2.1: Plots of $p^*(t)$, $U(t)$, and $V(t)$ ($\lambda = 50$, $\mu = 5$, and $T = 1$)

item at the end of the sale horizon. Therefore, the item is more valuable, the seller posts higher prices, and then the expected revenue is higher. Figures 2.4, 2.5 and 2.6 illustrate this argument by showing the plots of $p^*(t)$, $U(t)$ and $V(t)$ respectively for some different values of the arrival rate λ. It can clearly be seen that, for any given time, the (optimal) posted price ($p^*(t)$) and the expected revenue, with or without a customer arrival ($U(t)$ or $V(t)$), increases as the ratio λ/μ increases. In particular, $V(t)$ monotonically increases in t when the ratio is in the higher range as seen in Figure 2.6. In this case, the more time available for sale, the more advantageous for the seller to sell the item since the seller has more customer arrivals and thus more opportunity to sell the item with less possibility of cancellation. In addition, the more customers arrive, the more distributed reservation prices, and so the more time remaining for sale, the higher reservation prices customers are more likely to have. Therefore, it would be reasonable to expect higher revenue with higher posted prices earlier in the time horizon (targeting for higher reservation price customers), while lower revenue with lower posted prices later in the horizon (targeting for lower reservation price customers) would be expected.

On the other hand, when the ratio λ/μ is in the lower range, the expected revenue is much lower as seen in Figure 2.3. In this case, a customer is more likely to cancel the item immediately after purchasing it, that is, the revenue that a seller yields from a customer is more likely to be refunded to the customer. Hence, even though the seller posts higher prices in response to the higher arrival rate and then yields some higher revenue, this revenue is more likely to be fully cancelled out immediately after that. As a result, the item is not

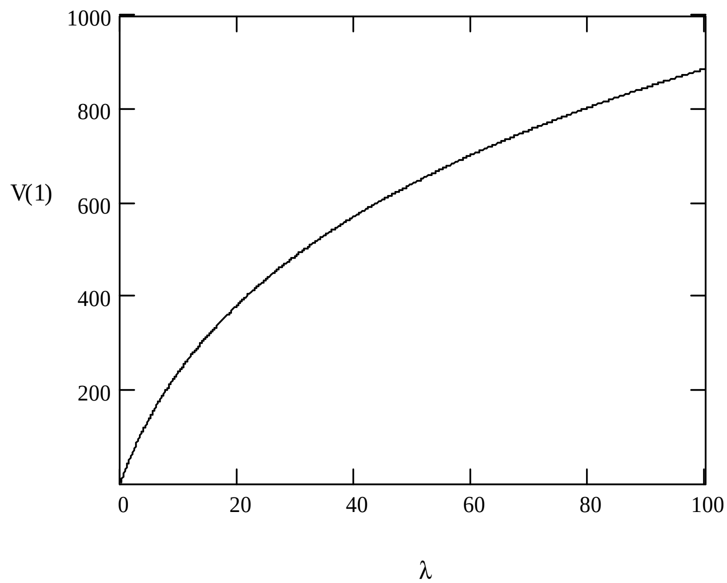

Figure 2.2: The expected total revenue $V(1)$ with respect to the arrival rate λ ($T = 1$, $\mu = 5$)

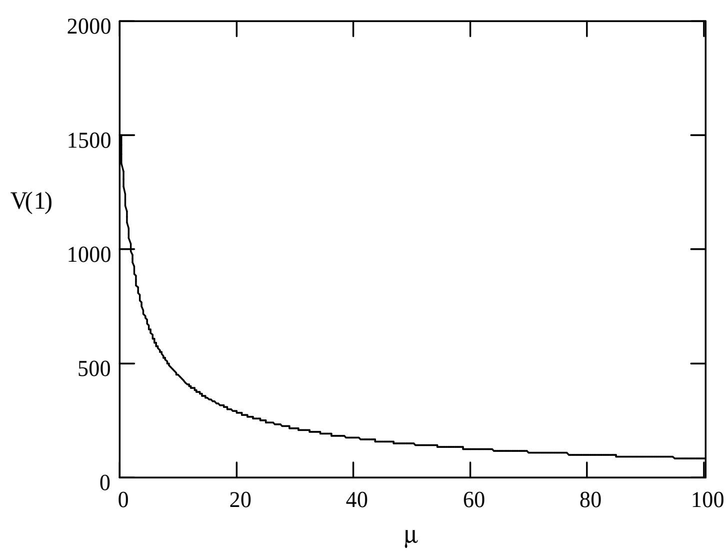

Figure 2.3: The expected total revenue $V(1)$ with respect to the cancellation rate μ ($T = 1$, $\lambda = 50$)

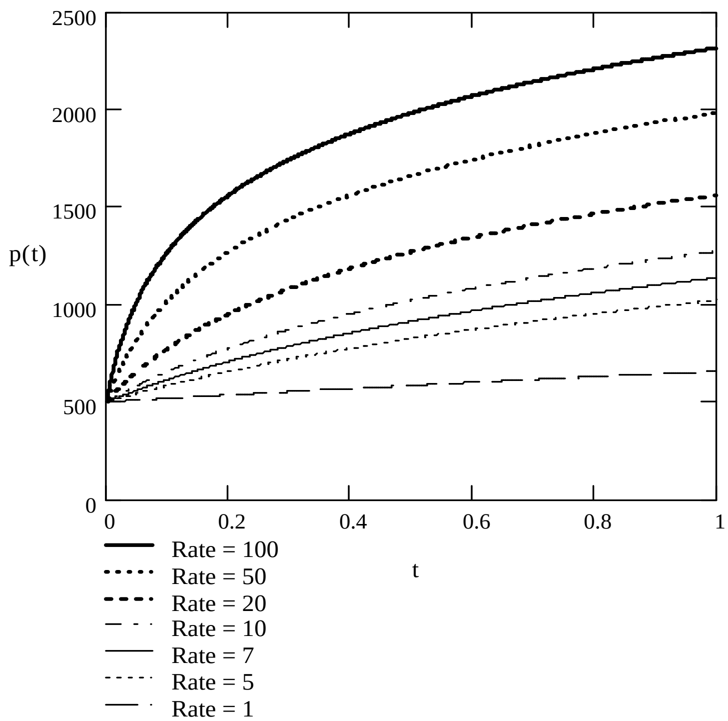

Figure 2.4: $p^*(t)$ for different values of the arrival rate $\lambda(T = 1,\, \mu = 5)$

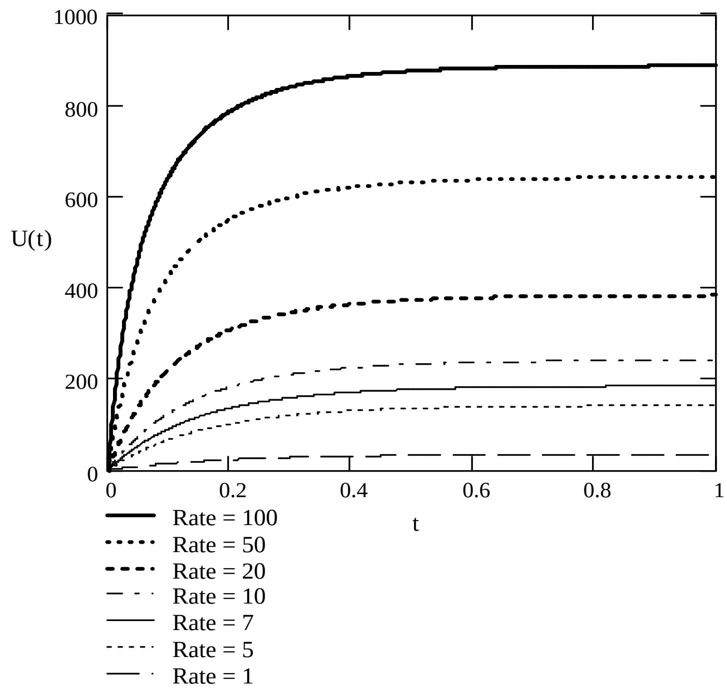

Figure 2.5: $U(t)$ for different values of the arrival rate λ ($T = 1$, $\mu = 5$)

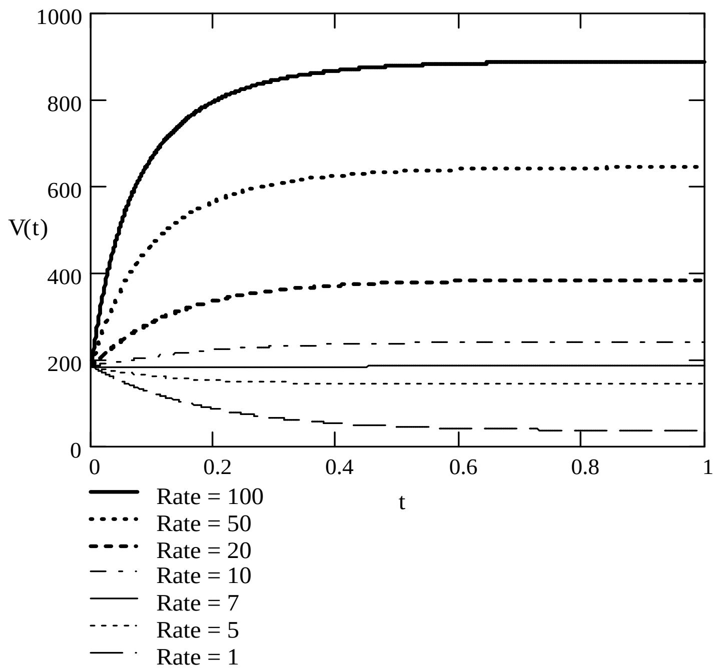

Figure 2.6: $V(t)$ for different values of the arrival rate $\lambda(T = 1,\, \mu = 5)$

expected to be sold out at the end of the sale horizon in this case. Therefore, the item is expected to be less valuable and the revenue is expected to be much lower, no matter how much in price the seller would post in response to the magnitude of the arrival rate (i.e., posted prices still depend on the magnitude of the arrival rate only). As in the previous case, Figures 2.7 and 2.8 illustrate this argument respectively by showing the plots of $U(t)$ and $V(t)$ for some different values of the cancellation rate μ. It can clearly be seen that, for any given time, the expected revenue, with or without a customer arrival ($U(t)$ or $V(t)$), decreases as the ratio λ/μ decreases. In particular, $V(t)$ monotonically decreases in t when the ratio is in the lower range, as seen in Figure 2.8. In this case, the less time available time for sale, the more advantageous it is for the seller, since there is less probability that a customer cancels the item and thus the seller has a slightly better chance to sell out of the item by the end of the sale later in the time horizon rather than earlier. Therefore, the expected revenue would be slightly higher later in the time horizon.

In summary, how much revenue a seller can expect depends on the ratio of the arrival rate to the cancellation rate. If the ratio is in the higher range, the revenue is expected to be higher in response to the arrival rate. Otherwise, the effect of cancellations would be so significant that it would adversely affect the revenue of the seller. Intuitively, in a full refund system, cancellation decreases the revenue of the seller every time it occurs since upon cancellation, the seller loses all revenue from the cancelling customer and simultaneously experiences the opportunity loss of time to sell the item; this causes the seller to post the prices in the rest of the time horizon that are lower than the first posted price. This means

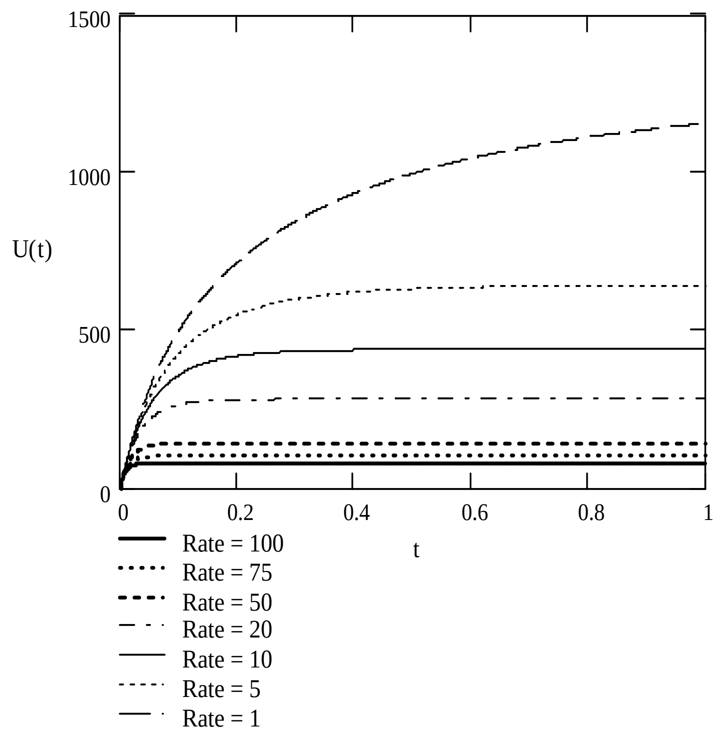

Figure 2.7: $U(t)$ for different values of the cancellation rate μ ($T = 1$, $\lambda = 50$)

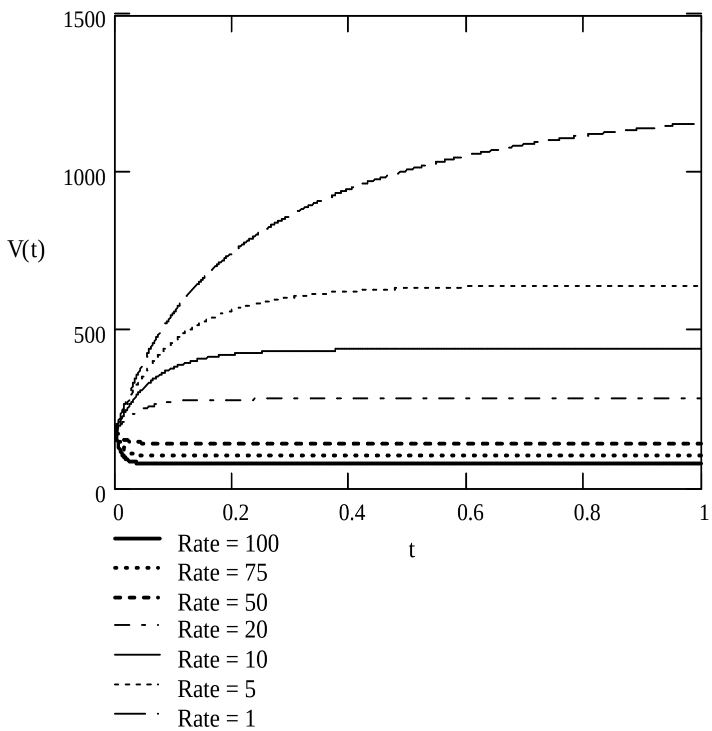

Figure 2.8: $V(t)$ for different values of the cancellation rate μ ($T = 1$, $\lambda = 50$)

that, once cancellation occurs, the total revenue afterward cannot be higher than the first yielded revenue. As a result, cancellation prevents the seller from yielding high revenue as long as a full refund policy is adopted.

Chapter 3

Approximation with the Discrete Time Model

In this section, we consider a discrete time dynamic pricing model to approximate the continuous time model and solve for the optimal pricing policy efficiently. We have seen that the optimal pricing policy is derived efficiently in a closed form when the customer reservation price is exponentially distributed. However, if the reservation price is not exponentially distributed, it is generally difficult to derive the optimal policy from a continuous time model. In addition, it is not efficient to implement a continuous time model for practical applications since the model involves such an intractable amount of computation. Most of the research in the literature, thus, focuses on developing heuristics to solve for the optimal pricing policy efficiently. The heuristic approaches are mainly categorized into two types: one that discretizes the time horizon in a continuous time model and one that uses a predetermined discrete set of prices. For example, it can be seen in Zhao and Zheng (2000) that the optimal value function in the continuous time model is approximately solved by discretizing the time horizon into a number of negligible time periods. Bitran and Mondscheine (1997) study the dynamic pricing model in which prices are changed only at the periodic review times that are specified in advance. In this paper, we consider a discrete time model that discretizes the continuous time horizon and uses a predetermined set of prices. We first formulate the model in a similar manner to the continuous time model and then analyze the approximation with numerical examples.

3.1 Formulation for Optimal Policy

We formulate the discrete time model in the same framework as the continuous one except for the point that a seller looks for the optimal price among the predetermined set of prices. There is no change in the interaction between the seller and the customers. We first divide the continuous time horizon T into K subintervals with the length of each interval Δt (i.e., $\Delta t = T/K$). We then enumerate the created $K+1$ discrete time points in a reversed order as Period K, Period $K-1$, Period $K-2$, and so forth. We assume that in any given period, say k, a customer arrives with a certain probability denoted by h and a customer with the item cancels it with a certain probability denoted by q. In addition, we assume that at most one of these two events (i.e., cancellation and customer arrival) occurs in any given period and so the seller is not allowed to sell the returned item in the same period during which a cancellation occurs. As in the continuous time model, the seller needs to make a decision on posting price only if the item is in inventory when a customer arrives. Hence, the state of the system is described as the remaining time periods, namely k.

Let $J(k)$ denote the optimal expected revenue from period k and onward. $J(k)$ is determined by conditioning on whether a customer arrives in period k or not. If a customer does not show up in period k, with probability $1-h$, there is no yielding revenue for the seller and the expected revenue from k and onward depends on the revenue from the next period and onward, that is, $J(k-1)$. On the other hand, if a customer arrives in period k, the expected revenue is determined by additionally conditioning on whether the arriving customer buys the item or not. If the arriving customer does not buy the item (or rejects a posted price p), with probability $F(p)$, there is no obtainable revenue in period k and the expected revenue from k and onward depends on the revenue from the next period and onward, that is, $J(k-1)$. In contrast, if the arriving customer buys the item (or accepts a posted price p), then $J(k-1)$ is determined by the summation of the expected revenue from the arriving customer (from period k and onward) and the expected revenue from the returned item after cancellation occurs in the future. Since now the number of periods between cancellations has a geometric distribution with parameter q, the probability that no cancellation occurs within the remaining periods is given as $(1-q)^{k-1}$. Therefore, the expected revenue from the customer in period k is given as $p(1-q)^{k-1}$. In addition, conditioning on the possible period when cancellation occurs after period k, we have,

$$\sum_{i=1}^{k-1} J(k-i-1)q(1-q)^{i-1}$$

as the expected revenue from the returned item after cancellation occurs up until the end of the sale. As a result, $J(k)$ is formulated as,

$$\begin{aligned} J(k) &= h \max_p \Big[F(p)J(k-1) + \overline{F}(p) \Big\{ p(1-q)^{k-1} + \\ &\quad + \sum_{i=1}^{k-1} J(k-i-1)q(1-q)^{i-1} \Big\} \Big] + (1-h)J(k-1) \end{aligned}$$

with the boundary condition $J(0) = 0$.

For the approximation, the continuous time horizon T is discretized into a number of the subintervals K so that the length of each interval Δt is sufficiently small, which makes it possible to assume that at most one customer arrives or at most one cancellation occurs in Δt. The probability of a customer arrival and that of cancellation in period k can then be approximated by the rates of the arrival and cancellation in Δt respectively. Therefore, h and q are determined as $h = \lambda \cdot \Delta t$ and $q = \mu \cdot \Delta t$.

When it comes to choosing a set of prices, selected prices should be variable and flexible enough to have a comprehensive scope of the reservation prices of customers. To reiterate, the concept of the dynamic pricing is to post a different price for each arriving customer since customers have different reservation prices. Hence, if selected prices are biased or too restricted, a seller will capture only a particular portion of the reservation prices and then experience the significant opportunity loss in revenue from the resultant poor approximation; this would be more significant when the arrival rate is higher since higher customer arrivals require more variable reservation prices since the problem of the seller lacking the capability to post appropriate prices for customers would be more critical.

To chose variable prices for placement into a set, we henceforth set prices in the following procedure using the probability density function (p.d.f.) of the reservation price: First of all, consider the area under the curve of the p.d.f. corresponding to the probability that the reservation price is greater than $p^*(0)$ (recall that any price less than $p^*(0)$ is never optimal and so it is futile to include such a price in a set.). Then divide the area by the number of prices n so that we have n identical sub-areas, and finally take those points dividing the original area as our choice of prices. The example when $n = 5$ is presented in Figure 3.1.

This procedure makes it possible for the customer reservation price to be equally likely to fall into each of the intervals of the selected prices, which ensures that the prices are variable enough to capture the comprehensive scope of the reservation prices of customers. Apparently, the higher the number of prices, the higher the number of sub-areas and the seller can then capture the reservation prices more precisely; the approximation for the optimal continuous policy is more refined as well.

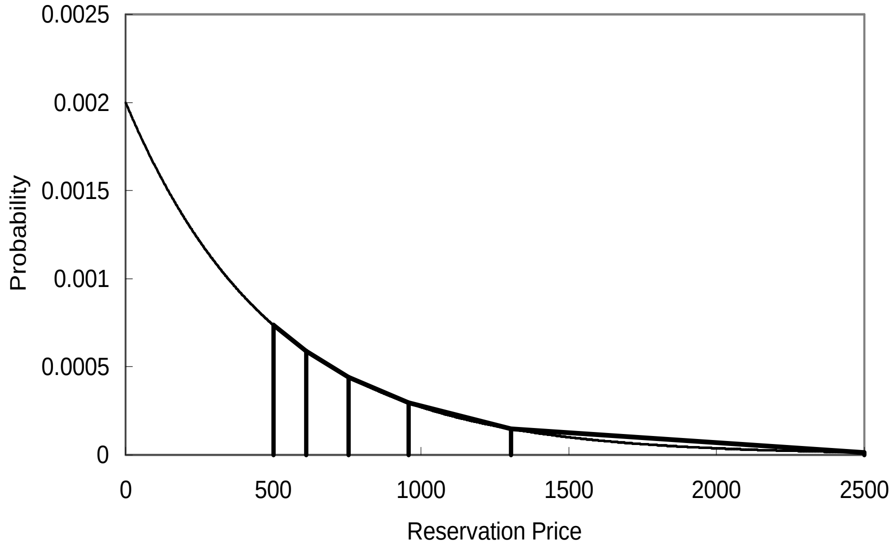

Figure 3.1: The choice of 5 prices using p.d.f. of the exponential reservation price

3.2 Numerical Example

In this section, we present numerical examples to approximate the continuous time model where the reservation price is exponentially distributed, and analyze the effectiveness of the approximation. Now we set the continuous time horizon $T = 1$, say 1 month, and divide it into the subintervals K, say to hourly units, that is, $K = 720$ $(= 24 \cdot 30)$ subintervals. We assume that the customer arrival rate is $\lambda = 50$ and the cancellation rate is $\mu = 5$ in the time horizon. Correspondingly, the probabilities of customer arrival and cancellation are determined respectively as $h = 50/720$ and $q = 5/720$. As previously, the mean of the exponential reservation price is set at \$500 (i.e., the parameter of the reservation price $\alpha = 0.002$).

Now we choose the number of prices $n = 8$ and, according to the procedure explained in Section 3.1, the set of prices is determined as shown in Table 3.1 with the corresponding probabilities that a customer buys an item with these prices (i.e., $\overline{F}(p)$) according to the exponential distribution. Figure 3.2 shows the plots of the optimal pricing policy and the expected revenue with the counterpart theoretical policy and expected revenue respectively. The approximated expected total revenue (i.e., $J(720)$) is actually computed as 639.57 while the theoretical expected total revenue (i.e., $V(1)$) is computed as 642.43.

Figure 3.3 shows the approximated expected total revenue with respect to the number of subintervals K. It is clear that the approximated expected total revenue is better with larger Ks. To reiterate, we need to divide the continuous time horizon T into a sufficiently large K to make the probability of a customer arrival and that of a cancellation in any time interval Δt sufficiently small; this is because we now approximate the continuous time model, that is, we imitate the process in the model using K, which is a Poisson process in this example. Thus, if K were chosen inappropriately, for example, if K were values of between 100 and 300 as in Figure 10, the approximation would be done quite poorly since we create an undesirable discrete time model with an insufficient arrival probability h in any period k. Therefore, to be specific, we should choose a K that is sufficiently larger than the expected total number of customer arrivals and cancellations.

When the discrete time model is implemented, however, a gap between the approximated expected total revenue and the theoretical revenue is inevitable, even with a sufficiently large K; this is because the discrete set of prices is more restrictive than the continuous set of

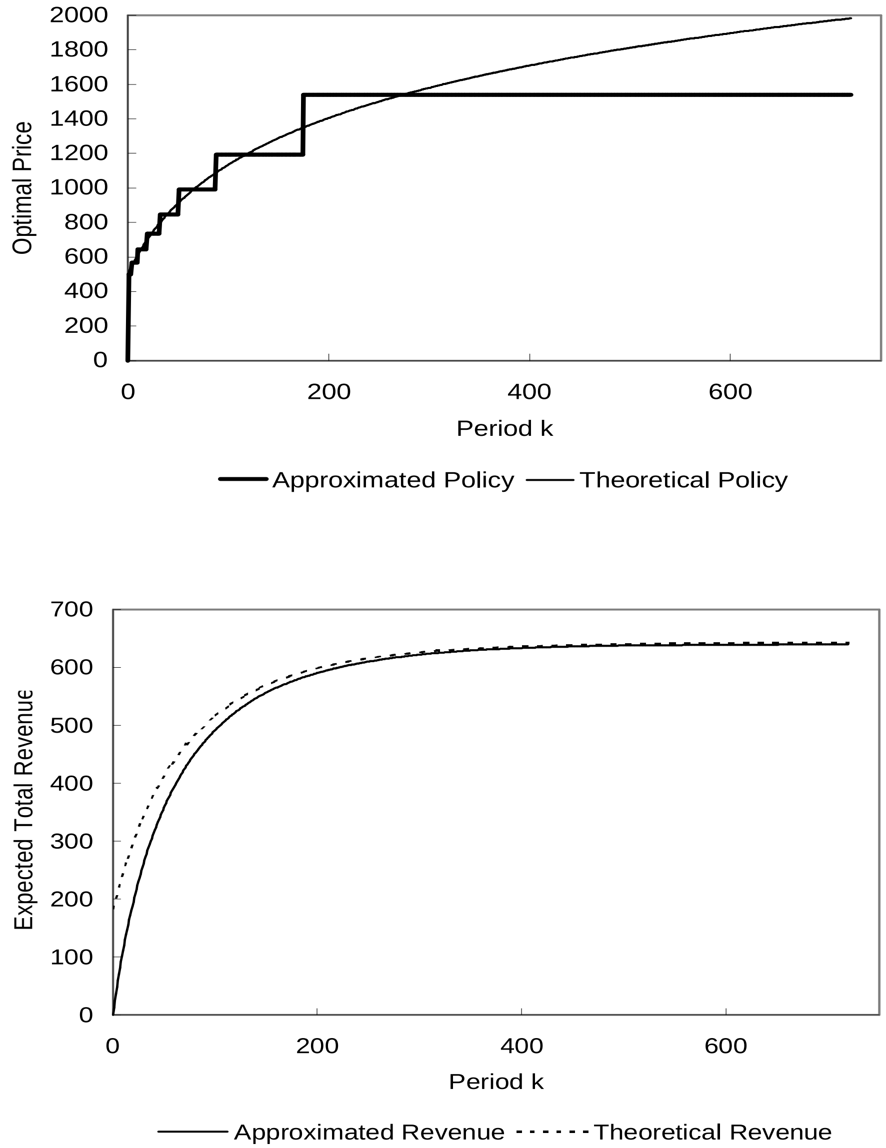

Figure 3.2: The approximation for the optimal policy and expected revenue($\lambda = 50$, $\mu = 5$, $n = 8$, $K = 720$)

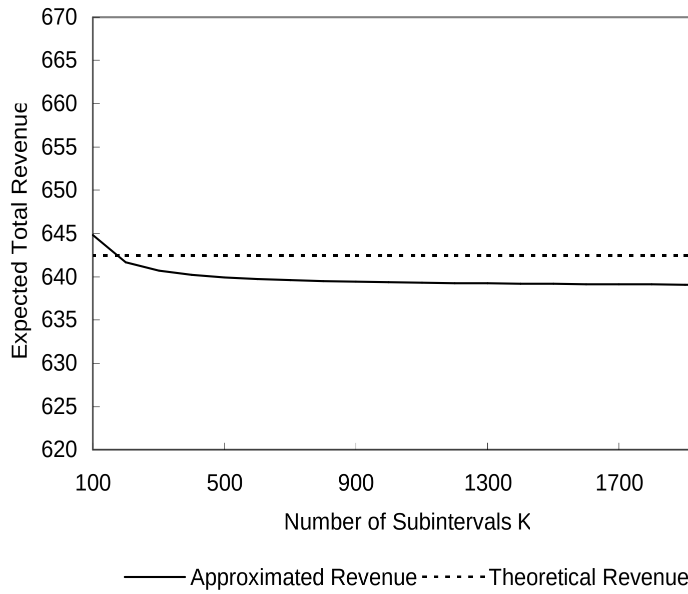

Figure 3.3: The approximated total revenue with respect to the number of subintervals K ($\lambda = 50$, $\mu = 5$, $n = 8$)

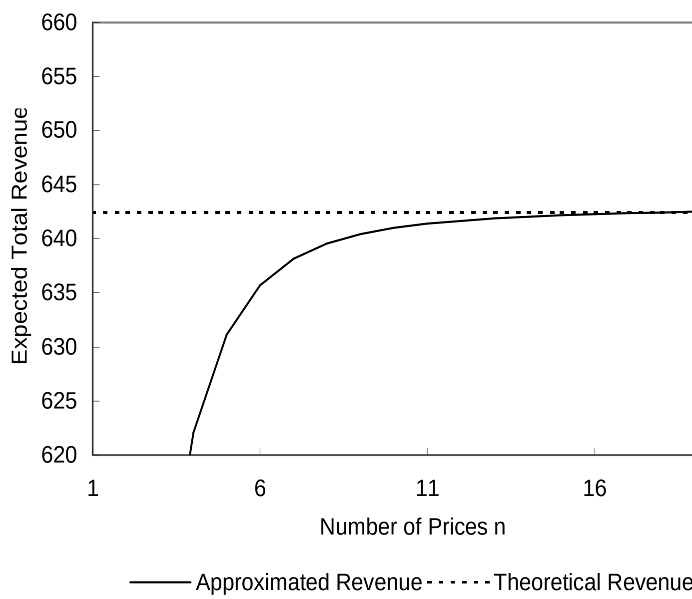

Figure 3.4: The approximated total revenue with respect to the number of prices n ($\lambda = 50$, $\mu = 5$, $K = 720$)

prices, which causes a seller to be unable to post appropriate prices for customers. Hence, every time a customer arrives, the seller experiences the opportunity loss in revenue, which is the main reason for the revenue gap. In order to reduce the opportunity loss, the seller should increase the number of prices and prepare more variable prices so that they can post more appropriate prices for customers. Figure 3.4 shows the approximated expected total revenue with respect to the number of prices n. It is clear that the approximated revenue is better with more prices in use. In the ongoing example, the approximated revenue is actually improved from 639.57 to 642.26 by enlarging the number of prices from 8 to 16.

However, now the seller has the tradeoff between the accuracy of the approximation and the computational effort. Of course, the more prices the seller has, the better approximation they can achieve, but at the same time, the more computational effort they need to make. We should keep in mind that we implement a discrete time model to mitigate the practically intractable computational effort involved in a continuous time model. Therefore, it makes no sense to have too many prices in the pursuit of a perfect approximation. The focus is now on how the seller can choose the number of prices to use.

The number of prices is actually influenced by the ratio of the arrival rate to the cancellation rate. Figure 3.5 shows the accuracy of the approximation with respect to the arrival rate for different numbers of prices focusing on ratios in the higher range. (Refer to Table 3.1 for the selected prices in each of the sets.) It is obvious that the higher the ratio is, the more prices the seller needs for better approximation. In this case, due to the higher arrival rate, there would be higher customer demand for the item and hence there would be more variation in the reservation prices of customers. Therefore, it is obvious that the more (variable) prices the seller has, the more precisely they can respond to these different reservation prices. Then, the opportunity revenue loss can be less and better approximation can be achieved. Hence, the seller should increase the number of prices in response to increasing magnitudes of the arrival rate.

On the other hand, Figure 3.6 shows the accuracy of the approximation with respect to the arrival rate for different numbers of prices focusing on ratios in the lower range. In this case, the approximation does not show much improvement even though the number of prices is increased. Now a customer is more likely to cancel the item and receive a full refund for the returned item. This means that from a seller's point of view, the revenue from a customer is more likely to be nullified by the customer's cancellation. Hence, even though a seller posts an inappropriate price for a customer and experiences some opportunity revenue loss, this opportunity loss appears more likely to be nullified by the customer's cancellation. As a result, no matter how many variable prices the seller has and no matter how tremendously he/she experiences the opportunity revenue loss, it would be expected to be much lower anyway. Therefore, the seller does not need to bother with so many prices. As a matter of fact, the seller should put more focus on the lower reservation prices of customers to reduce the opportunity revenue loss in this case: the seller is now more likely to sell the item later in the time horizon than earlier since the less time remaining, the lower the probability that a customer cancels the item. The seller, thus, is more likely to experience the opportunity revenue loss later in the time horizon. In addition, the seller has fewer chances to have higher

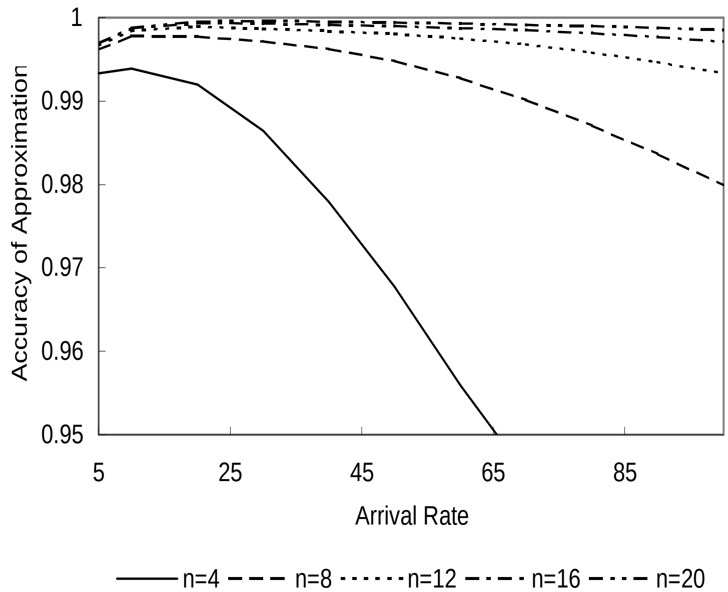

Figure 3.5: The accuracy of the approximation with respect to the arrival rate λ for different numbers of prices n ($\mu = 5$, $K = 2000$)

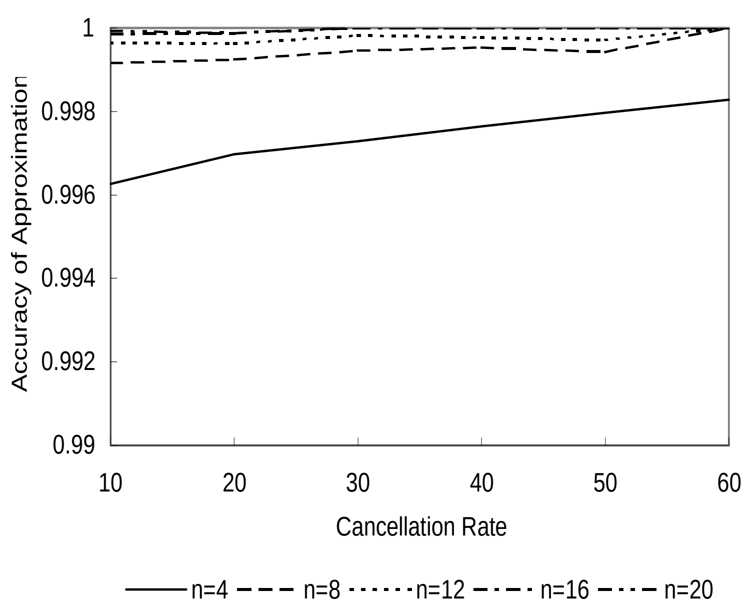

Figure 3.6: The accuracy of the approximation with respect to the cancellation rate μ ($\lambda = 10$, $K = 2000$)

Table 3.1: Set of prices when $n = 4, 8, 12, 16$, and 20 respectively

$n = 4$		$n = 8$		$n = 12$		$n = 16$		$n = 20$	
p	$\overline{F}(p)$	p	$\overline{F}(p)$	p	$\overline{F}(p)$	p	$\overline{F}(p)$	p	$\overline{F}(p)$
500.00	0.368	500.00	0.368	500.00	0.368	500.00	0.368	500.00	0.368
643.84	0.276	566.77	0.322	543.51	0.337	532.27	0.345	525.65	0.349
846.57	0.184	643.84	0.276	591.16	0.307	566.77	0.322	552.68	0.331
1193.15	0.092	735.00	0.230	643.84	0.276	603.82	0.299	581.26	0.313
-	-	846.57	0.184	702.73	0.245	643.84	0.276	611.57	0.294
-	-	990.41	0.138	769.50	0.215	687.35	0.253	643.84	0.276
-	-	1193.15	0.092	846.57	0.184	735.00	0.230	678.34	0.258
-	-	1539.72	0.046	937.73	0.153	787.68	0.207	715.39	0.239
-	-	-	-	1049.31	0.123	846.57	0.184	755.41	0.221
-	-	-	-	1193.15	0.092	913.34	0.161	798.92	0.202
-	-	-	-	1395.88	0.061	990.41	0.138	846.57	0.184
-	-	-	-	1742.75	0.031	1081.58	0.115	899.25	0.166
-	-	-	-	-	-	1193.15	0.092	958.15	0.147
-	-	-	-	-	-	1336.99	0.069	1024.91	0.129
-	-	-	-	-	-	1539.72	0.046	1101.99	0.092
-	-	-	-	-	-	1886.29	0.023	1193.15	0.092
-	-	-	-	-	-	-	-	1304.72	0.074
-	-	-	-	-	-	-	-	1448.56	0.055
-	-	-	-	-	-	-	-	1651.29	0.037
-	-	-	-	-	-	-	-	1997.87	0.018

reservation price customers later in the time horizon since there are fewer customer arrivals, which causes the reservation prices of customers to be less distributed. Therefore, the seller should have more coverage over lower reservation prices to reduce the more opportunity loss when the ratio of the arrival rate to the cancellation rate is lower.

Here is the summary for the algorithm to numerically approximate the optimal policy with the discrete time model.

1. Divide the continuous time horizon T into a sufficiently larger number of subintervals K. Use Δt to approximate h and q as $h = \lambda \cdot \Delta t$ and $q = \mu \Delta t$.

2. Choose a set of prices in accordance with the procedure (explained in Section 3.1).

3. Evaluate the value function $J(k)$ with each of the selected prices and pick up the price with the highest value of $J(k)$ as the (approximated) optimal policy in period k.

We should remember that the number of subintervals K should be sufficiently larger than both the arrival rates λ and μ and also that the number of prices n should be determined based on the ratio of the arrival rate λ to the cancellation rate μ; if the ratio is in the higher range, then n should be larger in response to the magnitude of λ, but it does not matter how large n is if the ratio is in the lower range. Rather, a seller should have more coverage over the lower reservation prices in that case. As a result, if appropriate values are chosen for K and n respectively, then the approximation will be successful. However, we should really keep in mind the tradeoff between the computational tractability and the approximation efficiency.

Chapter 4

Conclusion

In this paper, we have studied the opportunity of resale following cancellation in the continuous time dynamic pricing model where a single item is sold over a finite time horizon. We have formulated the dynamic pricing problem as a dynamic programming model and have derived the structural properties of the optimal pricing policy and the optimal value function. In particular, we have derived the optimal policy in a closed form when the customer reservation price is exponentially distributed and have gained some insight into the optimal policy. Furthermore, we have considered the discrete time dynamic pricing model to approximate the continuous time model by discretizing the continuous time horizon and the continuous optimal policy. We have presented numerical examples to show that the optimal policy is efficiently solved when we choose appropriate numbers of subintervals and of prices.

Consequently, we have learned, through the use of heuristics, that the ratio of the arrival rate to the cancellation rate influences the expected revenue of a seller. Actually, it turns out that cancellation is so significant that its effect cannot be ignored. If the customer arrival rate is not much higher than the cancellation rate, the seller cannot yield higher revenue. Also, intuitively, the expected total revenue in a full refund system is worsened every time a cancellation occurs. As a result, the seller cannot take advantage of the cancellation with resale as an opportunity to improve the total revenue in a full refund system. Therefore, our research describes the worst case scenario for the dynamic pricing models with some type of refund policy and could serve as a benchmark for those models with partial refund policy.

Regard to the direction of future research, some sort of partial refund policy (e.g., time-dependant refund policy) should be considered to see if a seller could take advantage of the cancellation for increasing revenue. Also, it would be interesting to consider multiple (identical) items for sale in a full refund or a partial refund system. In addition, in order to describe the process more dynamically, we should consider the time-dependent reservation price distribution and the time-dependent arrival and cancellation rate. It would also be of interest to consider the batch demand for items to make the model more realistic.

Bibliography

Bitran, G. R. and S. M. Gilbert, "Managing Hotel Reservations With Uncertain Arrivals," *Operations Research*, **44**(1), 35-49, 1996.

Bitran, G. R. and S. V. Mondschein, "Periodic Pricing of Seasonal Products in Retailing," *Management Science*, **43**(1), 427-443, 1997.

Biyalogorsky, E., Z. Carmon, G. E. Fruchter, and E. Gerstner, "Research Note: Overselling with Opportunistic Cancellations," *Marketing Science*, **18**(4), 605-610, 1999.

Brumelle, S. L. and J. I. McGill, "Airline Seat Allocation With Multiple Nested Fare Classes," *Operations Research*, **41**(1), 127-137, 1995.

Chatwin, R. E., "Multiperiod Airline Overbooking with A Single Fare Class," *Operations Research*, **46**(6), 805-819, 1998.

Chatwin, R. E., "Continuous-Time Airline Overbooking with Time-Dependent Fares and Refunds," *Transportation Science*, **33**(2), 182-191, 1999.

Feng, Y. and G. Gallego, "Perishable Asset Revenue Management with Markovian Time Dependent Demand Intensities," *Management Science*, **46**(7), 941-956, 2000.

Feng, Y. and B. Xiao, "Optimal Policies of Yield Management with Multiple Predetermined Prices," *Operations Research*, **48**(2), 332-343, 2000.

Feng, Y. and B. Xiao, "A Dynamic Airline Seat Inventory Control Model and Its Optimal Policy," *Operations Research*, **49**(6), 938-949, 2001.

Gallego, G. and G. van Ryzin, "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Science*, **40**(8), 999-1020. 1994.

Lee, T. C. and M. Hersh, "A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings," *Transportation Science*, **27**(3), 252-265, 1993.

Liang, Y., "Solution to the Continuous Time Dynamic Yield Management Model," *Transportation Science*, **33**(1), 117-123, 1999.

Lin, K. Y., "A Generic Dynamic Pricing Model and Its Application," *Manuscript*, 2002.

Lippman, S. and S. M. Ross, "Search with Posted Prices and Price Discrimination," *Manuscript*, 2002.

McGill, J. I. and G. J. van Ryzin, "Revenue Management: Research Overview and Prospects," *Transportation Science*, **33**(2), 233-256, 1999.

Robinson, L. W., "Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes," *Operations Research*, **43**(2), 252-263, 1995.

Subramanian, J., S. Stidham Jr, and C. J. Lautenbacher, "Airline Yield Management with Overbooking, Cancellations, and No-shows," *Transportation Science*, **33**(2), 147-167, 1999.

Zhao, W. and Y. Zheng, "Optimal Dynamic Pricing for Perishable Assets with Nonhomogeneous Demand," *Management Science*, **46**(3), 375-388, 2000.

Vita

Kwan-Ang An is Korean born in Japan on December 23, 1976 . He received a B.S. degree in Industrial and Systems Engineering from Aoyama Gakuin University (Tokyo, Japan) in March, 1999. He has been pursuing a M.S. degree within the area of Operations Research in Grado Department of the Industrial and Systems Engineering at Virginia Polytechnic Institute and State University since January, 2001.