Dynamic Pricing Strategies in an Uncertain World

Paul Johnson, Andreas Papayiannis and Kévin Poquet

January 23, 2012

Abstract

In this paper we investigate the problem of the seller who has fixed time in which to sell stock but is uncertain about the model governing price and demand. We assume there exists a finite set of possible models from which the true one is randomly chosen and remains fixed through time. The seller has different optimal strategies for each model, so we develop a new algorithm that chooses between strategies taking into account model misspecification. We test using simulations and the results show that by using Bayesian learning we can see an increase of up to 10% in sales revenues.

1 Introduction

Dynamic pricing policies are now common place in many areas of retail and hospitality sector, as the internet allows for sellers to use a centralised strategic pricing model, and an integral part of developing a strategic pricing model is the estimation of effect of price on the demand. When companies start to introduce the classic yield management type pricing policies into their business they are often face the problem that they have not done enough exploration of pricing in the past to give adequate estimation of the underlying effect price has on demand for their products. This uncertainty about the effect price will have means that they must choose their strategy and their models wisely, especially if they are selling perishable products where charging too high a price may result in the seller being left with unsold stock. In this paper we assume that the seller must identify a finite set of possible models for the effect of price and demand along with some probability that each of them are true. The different models may be identified by the functional form, or parameters within the model. Next the seller must identify how best to choose which model they should use to derive an optimal strategy, given the uncertainty about which which model might be the correct one. This choice should be informed not just by the probability that a model might be correct, but by what happens if you make the wrong choice. For instance, assume that there are two possible models for demand where price elasticity is either high or low. If the seller predicts that price elasticity is low then we can assume that an optimal strategy will opt to charge a higher price resulting in higher expected revenues *if* this model is correct. However *if* the model is wrong (and price elasticity is high) then the consequences could be disastrous as units are left unsold. Thus the benefit to the seller of choosing the strategy for low price elasticity is the sum of the gain in revenue from getting it right multiplied by the probability they are right plus the loss if they are wrong multiplied by the probability they are wrong. Therefore it is not certain that the seller would choose the strategy for low price elasticity even if it were the most likely scenario. Then in order to find the optimal pricing policy which maximizes expected revenues, we must determine the probability that the current set of observations have originated from each of the different models. This creates an incentive for Bayesian learning and updating.

Bayesian dynamic programming usually requires two distinct phases; an exploration (learning) phase where the demand is recorded and an exploitation phase where the estimates on the optimal prices are applied. Often, the learning phase takes place over a small time interval, to limit losses in revenues over the entire horizon. Nonetheless, by having it too small, it makes it hard to derive accurate estimates for the true demand, a fact that has a direct impact on the revenues during the remaining interval. Therefore, the managements problem is not only to determine good estimates to the true parameters but also to decide the optimal interval size for the learning phase [Araman and Caldentey, 2009]. Several researchers have studied Bayesian analysis techniques for a single- product RM problem with the objective of optimally adjusting prices to maximize expected revenues within a finite horizon T. Related literature involves the work of Aviv and Pazgal [2005] and Farias and Van Roy [2010], for which a single parameter of the demand function is not known.

Besbes and Zeevi [2009] on the other hand, extend the idea by examining two different settings:

1. a parametric setting, where the demand function is known to be of a particular form, but the exact parameter(s) values are not determined;
2. a non-parametric setting, where nothing is known about the demand curve other than it belongs to a broad class of functions satisfying mild-regularity conditions.

They use a full information benchmark (the demand is deterministic and it is known throughout the horizon) to evaluate the performance of the algorithms (minimax regret method is used). They begin by giving in a small time interval $[0, \tau]$ of the total $[0, T]$ for the learning phase of the algorithm; in this interval they test different prices and record the demand λ. Then, based on the results, a pricing phase takes place in the interval $[\tau, T]$, to identify and apply an accurate estimate of the price under the full information benchmark. Their results generate suboptimal policies of good performance. In particular, they show that the performance of the algorithm increases as the level of uncertainty decreases, implying that the loss in revenues from the non-parametric setting is greater than the one in the parametric setting, similar work has been pursued by Lim and Shanthikumar [2007].

Our model is differentiates itself from other Bayesian type models as we assume that the time in which the stock must be sold is too small to justify a full exploration of prices. Instead we assume that the seller chooses a single strategy to implement until the information arriving leads them to think they would be better off to change and implement a new strategy. Essentially we view this decision as a Real option that the company holds, namely the option to switch pricing strategy given the current state of the system

2 Brief Description of the Model

2.1 Assumptions

We follow the basic model and assumptions found in Gallego and Van Ryzin [1994] briefly recited here. A company is to sell some perishable product using a variable price p in some finite time horizon T. The company only has a finite stock of the variable, say Q units, and the value of the stock will be zero at time $t = T$, so you would imagine they would like to sell all of the stock. We assume that the firm operates in conditions in which they may alter demand for the product by varying the price p. Now we express the demand as a rate λ (units/time) that must depend only on the current price chosen by the company so that $\lambda = \lambda(p)$.

To make things stochastic we shall allow the demand intensity to be modelled by a Poisson process so that the probability that the company sells a unit over the time interval δt is $\lambda \delta t + o(\delta t)$ and the probability that they sell no units is $1 - \lambda \delta t + o(\delta t)$. We assume that the demand is a one-to-one function of price so that we may write the price p as a function of demand $p(\lambda)$. Then the revenue received over an interval δt is given by:-

$$\lambda p(\lambda) \delta t$$

the expected units sold

$$dQ = \lambda \delta t$$

and we may write the revenue rate r as

$$r = \lambda p(\lambda) \tag{1}$$

Now let $J_u(Q, \tau)$ denote the net present value of all future sales revenue received by the company, given the time left to execute the sales $\tau = T - t$ and units left unsold Q. Then given some pricing policy $u \in \mathcal{U}$

the value may be expressed as:-

$$J_u(Q,\tau) = E_u\left[\int_{\tau}^{0} p(\lambda^*)dQ\right] \tag{2}$$

where λ^* is the demand intensity chosen at each point in time under the current pricing policy.

Let us choose the optimal demand intensity $\lambda^*(Q,\tau)$ in order to maximise the revenue generated by the company, and let J^* denote the optimal value of revenue. Then using the Hamilton-Jacobi-Bellman principle we can write

$$\begin{aligned} J^*(Q,\tau) = \sup_{\lambda} \big[&\lambda p(\lambda)\delta\tau + \lambda J^*(Q-1,\tau-\delta\tau)\delta\tau \\ &+ (1-\lambda\delta\tau)J^*(Q,\tau-\delta\tau) + o(\delta\tau)\big] \end{aligned}$$

and after rearranging and taking limits reduces to

$$\frac{\partial J^*(Q,\tau)}{\partial\tau} = \sup_{\lambda}\left[r(\lambda) - \lambda\big(J^*(Q,t-\delta t) - J^*(Q-1,t-\delta t)\big)\right] \tag{3}$$

2.2 Uncertain Parameters

Let us make the assumption that the demand rate λ and the price rate p are known at time $t = 0$. This input to the model is obtained via prior knowledge, such as the rate of sales previously, marketing data or data from similar products. This will give us a single point on the demand function – but not the shape. For simplicity we will assume that the functional form of the equation is known, and choose the exponential form in honour of Gallego and van Ryzin (1994). Given that this equation has just two parameters, and we have been given one fixed point on the curve, we are left with just one degree of freedom, resulting in a known expected demand level and an unknown elasticity in the market.

Now let us define a model denoted M for the demand intensity as functional relationship between price and demand

$$\lambda(p) = ae^{-\alpha p} \tag{4}$$

where a and α are fixed parameters. Now assume that there exists a finite set of n possible models $\mathcal{M} \in \{M_1, M_2, \dots, n\}$ where each model M has its own particular choice for the parameters a and α. Then let the optimal pricing strategy $\pi_i \in \mathcal{U}$ be the optimal strategy associated with the model M_i, which maximises the expected revenue subject to (3) with the parameters identified within the model M_i. The expected value if parameters from the model come from M_i is optimal value or

$$J^*(Q,\tau) = J_{\pi_i}(Q,\tau;M_i) = J_{i,i}(Q,\tau) \tag{5}$$

If the model is wrong and the real parameters come from model M_j then we must solve (3) without the optimisation instead supplying the pricing strategy π_i. Then we may write the expected value from this system as

$$J_{\pi_i}(Q,\tau;M_j) = J_{i,j}(Q,\tau). \tag{6}$$

Since choosing the wrong strategy is suboptimal we know for certain that

$$J_{i,j}(Q,\tau) \le J_{j,j}(Q,\tau) \qquad \forall i \tag{7}$$

2.3 Belief Vector

Now we must introduce a belief vector $\boldsymbol{\theta}(t) \in \{\theta_1, \theta_2, \dots, \theta_n\}$, in which each element $\theta_i(t)$ is our belief at time t that model M_i is the true underlying model. If we can initialise the state of the belief vector with some prior information, then we may write the current expected value V of choosing strategy π_i at time t as

$$V_i(Q,T-t) = \sum_j \theta_j(t)J_{i,j}(Q,T-t) \tag{8}$$

Then we propose that the strategy of the seller should be to choose the pricing strategy π_i such that V_i is maximised. Then at any time the seller has the option to switch strategies if a different strategy yields a higher expected value.

Now we must define how we can update our 'belief' about the system given all the current observations to date. We propose to achieve this by using the following Bayesian transformation given the observation over the time interval dt. So long as the resulting vector gives a true reflection of the probabilities that the observations came from a certain model we should be able to maximise the value of the option to switch.

References

V. Araman and R. Caldentey. Dynamic pricing for non-perishable products with demand learning. *Operations research*, 57(5):1169–1188, 2009.

Y. Aviv and A. Pazgal. A partially observed markov decision process for dynamic pricing. *Management Science*, pages 1400–1416, 2005.

O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations research*, 57(6):1407–1420, 2009.

V.F. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Operations Research*, 58(1): 16–29, 2010.

G. Gallego and G. Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science*, pages 999–1020, 1994.

A.E.B. Lim and J.G. Shanthikumar. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research*, 55(2):198, 2007.