# Dynamic Pricing with Buyer Learning

Anton J. Kleywegt*
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332
anton@isye.gatech.edu

October 23, 2006

**Abstract**

Many traditional dynamic pricing models such as the ones widely used in revenue management assumed that the demand at each point in time depends on the price at that point in time only, that is, it is independent of prices at other points in time. Recently some models of so-called strategic customer behavior have been studied, in which buyers' purchasing decisions at a point in time depend on the prices at other points in time, or more generally, on the sellers' pricing policies. Many new questions are associated with such models. One question is how the buyers can be expected to obtain and process all the information necessary to make such complicated decisions. We study several models in which buyers learn quantities that are simpler than the pricing policies of the sellers. We investigate the convergence of the buyers' estimates, and compare the limits with equilibria associated with full information.

## 1 Dynamic Pricing with Learning by the Buyer

The demand models used in most revenue management research assume that the demand at a particular point in time does not depend on the prices or availabilities of products at other points in time. It seems that such models underestimate the sophistication of the buyers' decision processes. Recently, some models have been developed in which buyers' choices depend on the prices or availabilities of products at multiple points in time. Some of this work includes Besanko and Winston (1990), Xu and Hopp (2004), Aviv and Pazgal (2004, 2005), Su (2005), Liu and Van Ryzin (2005), Gallego and Şahin (2005), Gallien (2006), Levin et al. (2006), and Zhang and Cooper (2006). In most of this work, it is assumed that buyers know the seller's dynamic pricing policy (most of the work considers a single seller), and the buyers choose a best buying policy in response to the seller's dynamic pricing policy. In contrast with traditional revenue management models, it seems that such models overestimate the sophistication of buyers' decision processes. Typically, sellers' dynamic pricing policies can be quite complicated, and it is questionable whether buyers have the data, the insight, and the computational power to learn the sellers' dynamic pricing policies and to compute their best responses to these policies. If sellers use heuristics unknown to the buyers, it may be an even harder task for buyers to figure out what the resulting pricing policies are. Thus we are interested in studying what happens if buyers do not know sellers' dynamic pricing policies, but rather learn simpler quantities, and use the resulting simple models to make their decisions.

Next we briefly describe two problems that are motivated by the considerations described above. Section 2 considers a setting in which buyers attempt to learn the probability distribution of the

*Supported by the National Science Foundation under grant ITR/DMI-0427446.

spot price of a product, and Section 3 considers a setting in which buyers learn about product availability.

## 2 Buyer Learning of Spot Price Distribution

Consider a setting with both longer term contracts and a spot market for a product or service. Such a setting is typical in markets for freight transportation services, in which carriers and customers can enter into longer term contracts that specify prices before it is known exactly how much freight will be transported, and in which customers can also purchase transportation services when it is known how much freight will be transported, even if they had not entered into a contract before. There is quite a large literature on the interaction between contracts and spot markets; see for example Cohen and Agrawal (1999), Lee and Whang (2002), Wu et al. (2002), Kleindorfer and Wu (2003), Mendelson and Tunca (2003), Burnetas and Ritchken (2005), Wu and Kleindorfer (2005), and Tunca and Zenios (2006). However, none of this work considers learning by the market participants. In the model that we consider, the buyers do not know how the prices will be determined on the spot market when they make their contracting decisions. Instead, the buyers observe a sequence of spot market prices and attempt to estimate a probability distribution for the spot price. At each step in the sequence, the buyers base their contracting decisions on their current estimate of the spot price distribution. The seller's choice of the spot price is affected by the contracting decisions of the buyers. Thereafter, buyers make their spot market purchasing decisions, which are affected by the spot price and the quantities contracted by the buyers. Thus, the buyers affect the spot prices through their contracting decisions, and thus through their estimates. Thus the sequence of observed spot prices is neither independent nor identically distributed, but depends on the decisions and thus the estimates of the buyers at every step. As mentioned, the buyers do not know exactly how this complicated dependence works, and construct a sequence of probability distributions as though they are estimating a single exogenous distribution, instead of estimating a family of distributions that explicitly depend on their decisions. It seems intuitive that if each buyer's decisions have only an infinitesimal effect on the spot price, then the sequence may converge to a point that is also an equilibrium for the game in which the players know the details of the system's behavior. We also identify sufficient conditions for such convergence even if each buyer's decisions have a significant effect on the spot price. Thus two interesting questions are the following:

1. Sufficient conditions for the sequence of distribution estimates to converge. (Recall that the sequence of random variables is not i.i.d., but is generated by quite a complicated process with feedback.)
2. If the sequence converges, how the limit is related to equilibria in the game with full information.

## 3 Buyer Learning of Product Availability

Often customers buy airline tickets before they are certain that they will take the trip. It is not unusual for 1/3 or more of the bookings on hand for a flight a week before departure time to be cancelled during the last week. There are several incentives for customers to buy airline tickets before they are certain that they will take the trip. One incentive is that the tickets often tends to become more expensive as the departure time comes closer. Another reason is that customers want to make sure that they get a booking on their preferred flight.

We consider a model in which the buyers are not sure whether they will need the product later, but there is an incentive to buy the product early because the price increases over time and because the product may not be available later. The probability that the product is available later depends on the earlier purchasing decisions. Each buyer does not observe the earlier purchasing decisions of the other buyers and/or does not know the exact dependence of the later availability on the earlier decisions. Instead, each buyer observes product availability when they attempt to buy the product later. The process repeats, and each buyer uses such a sequence of observations to estimate a probability that the product is available later if the buyer does not buy it early. At each step, each buyer uses the buyer's current estimate to make a decision, namely whether to buy the product early or to wait until later. A buyer's decisions at one step in the sequence affects the observations of the other buyers, and thereby the estimates and the decisions of the other buyers at later steps, and hence the observations of the buyer at later steps. However, each buyer does not take this complicated dependence into account, and does not model availability as a function of the decisions. Rather, each buyer only estimates availability as though it is an exogenous quantity. Questions to be addressed in this work are similar to those for the previous problem, namely (1) conditions under which the estimates converge, and (2) if the estimates converge, how the limit is related to equilibria with no modeling error.

Next we give a few more details regarding the latter problem. The seller sells a single type of product, and chooses a price for each of two time periods. The seller has a given amount $C$ of the product available for sale over the two time periods. The seller announces the prices $x_1$ and $x_2$ for each of the two time periods, and each buyer decides whether to buy in the first period, or whether to wait until the second period. Other revenue management literature that considers buyer decisions over two periods include Liu and Van Ryzin (2005), Gallego and Şahin (2005), and Zhang and Cooper (2006). In the first time period, buyers are not certain whether they will end up needing the product (for example, travelers may later change their minds and decide not to travel). In the first time period, each buyer knows a probability $p$ that the buyer will end up needing the product, and if so, a utility $u$ for the product. The values of $p$ and $u$ follow a joint distribution $F$ in the population of potential buyers. In the second period, each buyer knows whether he/she needs the product or not. It may turn out that there is not enough product for all the buyers who want the product in the second period. Buyers do not know in advance whether the product will be available in the second period.

There are various versions of the problem depending on the number of buyers and whether the seller changes prices or not. Here we describe a version with $m > 1$ buyers, $m < \infty$. In the first period, each buyer $i$ knows $x_1$, $x_2$, and the buyer's values $p_i$ and $u_i$, and has to decide whether to buy or wait. The seller does not know the values $p_i$ and $u_i$, and each buyer does not know the values of the other buyers. The strategy of each buyer $i$ can be represented by a function $L_i : [0,1] \times [0,\infty) \mapsto \{0,1\}$, with $L_i(p_i, u_i) = 1$ denoting that buyer $i$ decides to buy in the first period. Let $\mathcal{L}$ denote the set of such functions $L$ with $L^{-1}(1)$ a measurable set. The objective function of buyer $i$ is given by $f_i : \mathcal{L}^m \times [0,\infty)^2 \times [0,1] \times [0,\infty) \mapsto \mathbb{R}$. (Objective function $f_i(L_1, \ldots, L_m, x_1, x_2, p_i, u_i)$ gives the expected utility of buyer $i$ given $x_1, x_2, p_i, u_i$, taking into account the dependence of second period availability on the strategy combination $(L_1, \ldots, L_m)$ and on $(x_1, x_2)$. It is easy to give an expression for $f_i$, but it requires additional notation that we want to avoid here.) Given the prices $x_1$, $x_2$, a strategy combination $(L_1^*, \ldots, L_m^*) \in \mathcal{L}^m$ is an equilibrium for the buyers if, for each $i$, $p_i$, and $u_i$, $f_i(L_1^*, \ldots, L_i^*, \ldots, L_m^*, x_1, x_2, p_i, u_i) \geq f_i(L_1^*, \ldots, L_{i-1}^*, L, L_{i+1}^*, \ldots, L_m^*, x_1, x_2, p_i, u_i)$ for all $L \in \mathcal{L}$. The objective function of the seller is given by $g : \mathcal{L}^m \times [0,\infty)^2 \mapsto \mathbb{R}$. Suppose that, for each $(x_1, x_2)$, there is a unique equilibrium $(L_1^*(x_1, x_2), \ldots, L_m^*(x_1, x_2))$ for the buyers. Then the seller's problem is $\max_{x_1, x_2 \geq 0} g(L_1^*(x_1, x_2), \ldots, L_m^*(x_1, x_2), x_1, x_2)$.

The setup in the previous paragraph defines equilibria with full information. As mentioned before, we are interested in the setting where buyers do not know the dependence of second period availability on the strategy combination $(L_1, \ldots, L_m)$. Instead, when buyer $i$ decides to wait until the second period and then decides to buy the product, the buyer either gets a product or not, and the buyer records the observation. The process repeats, producing a sequence of decisions and observations, each involving two periods. At each step $n$ of the sequence, each buyer $i$ uses the buyer's observed data to estimate the "probability" $\hat{H}_{i,n}$ that the product will be available in the second period of that step. Each buyer $i$ makes a decision in the first period of step $n$ to maximize his/her own "expected" utility, choosing $\max\{p_{i,n}u_{i,n} - x_1, p_{i,n}[u_{i,n} - x_2]^+ \hat{H}_{i,n}\}$, using the estimated probability $\hat{H}_{i,n}$ that the product will be available in the second period. The first and second period demands, $D_{1,n}$ and $D_{2,n}$ respectively, are determined by aggregating the customers' decisions. Then the actual fraction $H_n$ of second period demand that is satisfied is given by $H_n = \min\{[C - D_{1,n}]^+, D_{2,n}\}/D_{2,n}$. Note that $H_n$ depends on $(\hat{H}_{1,n}, \ldots, \hat{H}_{m,n})$, although the notation does not show the dependence. In one version of the problem, each buyer does not observe $H_n$, but only whether that buyer gets a product in the second period or not (if the buyer does not want to buy product in the second period, then the buyer makes no observation in that step). In another version of the problem, each buyer observes $H_n$ at step $n$.

Next, a number of questions arise naturally. One question is whether there exist values of $(\hat{H}_1, \ldots, \hat{H}_m)$ such that the resulting expected actual fraction $H = \hat{H}_i$ for all $i$, that is, the expected actual fraction $H$ is consistent with the "probability" estimates $\hat{H}_i$ used by the buyers, or in other words, $\hat{H}_i$ is a fixed point of $H$ regarded as a function of $\hat{H}_i$. If such values exist, then a second question is whether the estimates of the buyers would converge to such a consistent point if the buyers started with a different estimate and used a good forecasting method to update their estimates (again, the sequence of observations of each buyer is neither independent nor identically distributed). A third question is how such a consistent point compares with equilibria with full information.

## References


Y. Aviv and A. Pazgal. A partially observed markov decision process for dynamic pricing. 2004.

Y. Aviv and A. Pazgal. Optimal pricing of seasonal products in the presence of forward-looking customers. 2005.

D. Besanko and W. L. Winston. Optimal price skimming by a monopolist facing rational consumers. *Management Science*, 36(5):555–567, 1990.

A. N. Burnetas and P. Ritchken. Option pricing with downward-sloping demand curves: The case of supply chain options. *Management Science*, 51(4):566–580, 2005.

M. A. Cohen and N. Agrawal. An analytical comparison of long and short term contracts. *IIE Transactions*, 31:783–796, 1999.

G. Gallego and Ö. Şahin. Inter-temporal valuations, product design and revenue management. Working paper, Columbia University, 2005.

J. Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.

P. R. Kleindorfer and D. J. Wu. Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Science*, 49(11):1597–1615, 2003.

H. L. Lee and S. Whang. The impact of the secondary market on the supply chain. *Management Science*, 48(6):719–731, 2002.

Y. Levin, J. I. McGill, and M. Nediak. Optimal dynamic pricing of perishable items by a monopolist facing strategic consumers. 2006.

Q. Liu and G. Van Ryzin. Strategic capacity rationing to induce early purchases. Working paper, Columbia University, Graduate School of Business, 2005.

H. Mendelson and T. I. Tunca. Liquidity in industrial exchanges. 2003.

X. Su. Inter-temporal pricing with strategic customer behavior. 2005.

T. I. Tunca and S. A. Zenios. Supply auctions and relational contracts for procurement. 2006.

D. J. Wu and P. R. Kleindorfer. Competitive options, supply contracting, and electronic markets. *Management Science*, 51(3):452–466, 2005.

D. J. Wu, P. R. Kleindorfer, and J. E. Zhang. Optimal bidding and contracting strategies for capital-intensive goods. *European Journal of Operational Research*, 137:657–676, 2002.

X. Xu and W. J. Hopp. Customer heterogeneity and strategic behavior in revenue management: A martingale approach. 2004.

D. Zhang and W. L. Cooper. Managing clearance sales in the presence of strategic customers. Working paper, University of Minnesota, 2006.