Applications and Applied Mathematics: An International Journal (AAM)

Volume 16 | Issue 1 Article 1

6-2021

A Stochastic Knapsack Game: Revenue Management in Competitions

Yingdong Lu
IBM T.J. Watson Research Center

Available at
http://pvamu.edu/aam
Appl. Appl. Math.
ISSN: 1932-9466

Applications and Applied Mathematics:
An International Journal
(AAM)

Vol. 16, Issue 1 (June 2021), pp. 1 – 11

A Stochastic Knapsack Game: Revenue Management in Competitions

Yingdong Lu

Mathematical Sciences
IBM T.J. Watson Research Center
1101 Kitchawan Rd
Yorktown Heights, NY, USA
yingdong@us.ibm.com;

Abstract

We study a mathematical model for revenue management under competitions with multiple sellers. The model combines the stochastic knapsack problem, a classic revenue management model, with a non-coorperative game model that characterizes the sellers' rational behavior. We are able to establish a dynamic recursive procedure that incorporate the value function with the utility function of the games. The formalization of the dynamic recursion allows us to establish some fundamental structural properties.

Keywords: Revenue management; Dynamic recursion; Competition; Optimization; Monotonicity; Concavity; Game theory

MSC 2010 No.: 90B36, 49L20, 91A80

1. Introduction

A key model in revenue (yield) management is the following, a seller needs to sell a fixed amount of certain commodity before a fixed deadline to different buyers with individual price they are willing to pay, and the seller can dynamically adjust the selling price to maximize his/her overall revenue over time. Stochastic knapsack problem, also known as stochastic dynamic knapsack problem, a mathematical problem that captures the essences of this model, quantifies some of the most fundamental trade-offs in revenue management, and serves as an important building block for more complex and sophisticated models for real life applications. Consequently, the stochastic (dynamic) knapsack problem and its variations have been studied extensively, see, e.g., Gallego and van Ryzin (1994), Bitran and Mondschein (1995), Feng and Gallego (1995), Papastavrou et al. (1996), Feng and Gallego (2000), Feng and Xiao (2000), Van Slyke and Young (2000), Zhao and Zheng (2000), and Lin et al. (2008). It is one of the fundamental models surveyed in Den Boer (2015); please refer to that paper for more details, as well as references.

It is natural to ask the question of what would happen if there are multiple sellers competing for the same demand stream from the buyers. In this paper, we generalize the classic stochastic knapsack problem, and formulate a mathematical model to capture the basic relations in this situation. An immediate goal is to formulate a dynamic recursion for calculating the optimal policies for sellers. In the single seller case, this is accomplished through the formulation of a dynamic program that computes the maximum expected revenue starting at any time with any amount of remaining inventory. However, in the case of multiple sellers, at each time period, the sellers' decisions are inter-dependent. It is, therefore, not a trivial task to decide what will be the next best action even if every seller has the same forecast of the future demand arrivals. Another difficulty is that when multiple sellers are willing to sell the product, the buyer can have different ways to choose one of them to fulfill the demand, the difference in these selection rules has significant impact on the evolution of the system. To overcome these difficulties, we model the sellers as rational individual or institutions, and introduce a non-cooperative game at each step of the dynamic recursion characterizing their behavior. Furthermore, we follow a static probabilistic selection rule, which will be described precisely later, that the buyer will use to select sellers. This selection rule, on one hand, reflects market power of the sellers, on the other hand, it allows the uncertainty that is natural in business reality. With this mechanism, the utility functions of the games are properly connected with the value functions of the dynamic recursion, thus help to identify pure strategy Nash equilibriums. Under the selection rule assumed, we are able to demonstrate that there is a unique Nash equilibrium of the game. In turn, assuming that the Nash equilibriums will be the strategy followed by all the sellers at each step, the dynamic recursion is able to proceed. Once establishing the dynamic recursion, we are able to extend the arguments that are effective for the single seller dynamic programming, and demonstrate that, in some cases, the value function exhibits remarkable rich monotonicity properties that provide insights to key trade-offs to the problem and can be helpful to dynamic pricing in practice. A related but different model is considered in Gallego and Hu (2014), it is concluded that, under a differential game setting, the equilibrium structure enjoys simple structural properties. While the model studied here is quite different, but results are similar in spirit.

The rest of the paper will be organized as follow. In Section 2, we will introduce the basic mathematical models, and review preliminaries including some basic concepts in game theory that will be needed for our analysis. In Section 3, we will discuss in details the dynamic recursion in which the game aspect of the problem is incorporated. In Sections 4 and 5, we establish some fundamental structural properties of the value functions of the dynamic recursion. Finally, we conclude the paper in Section 6 with a summary of our findings.

2. Models and Preliminaries

2.1. Model Descriptions

Suppose that there are N sellers, and each seller n, $n = 1, 2, \ldots, N$, has an initial inventory of C_n units of product (could be either goods or services) at the beginning of a common selling horizon. The selling horizon is discrete and of length $T < \infty$. At each time $t = 1, 2, \ldots, T$, demand for one unit of the product will emerge, and the buyer will post a price that he/she is willing to pay. To accommodate the event of no arrival, we can always include a class of demand with exceedingly low price. The sellers who have positive inventory need to decide whether they should accept or reject this demand. The buyer will then select one seller among all the sellers that accept the demand according to certain *selection rule*, and the selected seller will supply the product and collect the revenue. At the end of the selling horizon, all the remaining product will be savaged. The goal for each seller is to maximize his/her expected revenue.

We assume that each seller does not have the information of the exact value of the initial inventory of other sellers, but has a distributional estimation of that quantity. We also assume that the distributional information of the future demand price is given to each seller, and no seller has any extra knowledge. In particular, we assume that the price of the demand realization at each period follow an independent and identically distributed discrete probability distribution P, with $\mathsf{P}[P = p_i] = \theta_i$, $i = 1, 2, \ldots, I$.

Suppose that, at each time t, when the demand is of class i, i,e. the price is p_i, a subset of sellers, denoted by $A_t(i)$ (which can be shortened to A_t when there is no ambiguity), will accept the demand, decided based on the remaining time, demand type, remaining inventory and the selling history up to time t. The buyer will select only one seller among them, which means that there is a possibility that no seller is selected. There could be various selection rule models reflecting different market mechanisms, for example, a static rule (the buyer chooses one product over the other overwhelmingly, which happens often in some local and monopoly market) and weighted rule (buyer assigns weights to the each product, then randomly, with probabilities determined by the weights, select ones that are available). In this paper, we will focus on a random allocation rule with static probabilities: each seller is associated with a probability π_n, $\sum_{n=1}^{N} \pi_n = 1$. At each time, if a seller accepts, the probability of it being selected is always π_n, and with probability $1 - \sum_{n \in A_t(i)} \pi_n$, no one is selected.

At each time t, the phenomenon that the sellers are making independent decisions based on distributional information on the other sellers can be best modeled by a non coorperative strategic game, see, e.g., Osborne and Rubinstein (1994).

3. A Dynamic Recursion Formulation

Our goal is to identify a strategy for a seller to achieve the best outcome, in terms of average revenue, under a reasonable assumption on other sellers' behavior. Recall that each seller n, $n =$

$1, 2, \ldots, N$ with initial inventory C_n is also given the distributional information of the inventory of all other sellers, either though statistical forecast or other business information inquiry, and any two sellers will be given the exact same distribution on the third seller. In addition, all the sellers do observe all the sells outcomes up to each decision time epoch, i.e., they know the amount each seller sold so far. It is our intention to derive a dynamic recursion for calculating the best outcome, hence the optimal strategy for each seller. Equivalently, given $\mathbf{s} = (s_1, s_2, \ldots, s_N)$ representing the amount of inventory has been sold so far by each seller, we seek to calculate $v_n(t, d_n, \mathbf{s})$, $n = 1, 2, \ldots, N$, the maximum expected revenue seller n can collect starting from time t and with remaining inventory d_n, for any time $t = 1, 2, \ldots, T$.

We assume that the behavior of the sellers is modeled as a N-person game, and if sellers follow the Nash Equilibriums at each time period, a dynamic recursion can proceed. This will be argued inductively. At the last time period T, given a price realization, p_i, there are two strategies for each seller, accept or reject. The utility function of the game for seller n will be the expected revenue collected by taking either action. If reject, of course, there is no revenue. It is clear that, if the random selection rule with static probabilities is followed, there exists a unique Nash equilibrium, that is, every seller will accept, as long as they have a positive inventory. In this case, the value function $v_n(T, \mathbf{d})$ has the following form, $v_n(T, \mathbf{d}) = \pi_n \mathsf{E}[P]$, where π_n are the probabilities in the section rule model.

Now, suppose that we can calculate recursively all the value function $v_n(t+1, d_n, \mathbf{s})$ for any feasible $\mathbf{s}$, we demonstrate that there exists a unique pure strategy Nash equilibrium at time period t, and show how it is related to the calculation of the value function for time t, $v_n(t, d_n, \mathbf{s})$. There are two actions for each seller, accept or reject. The payoff function will be the expected revenue to be collected until time T. Therefore, the seller will consider the following *balance inequality*, whose left hand side (LHS) represents the price we get immediately, and right hand side (RHS) represents the future reward,

$$p \geq \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s}) - v_n(t+1, d_n - 1, \mathbf{s} + \mathbf{e}_n)], \tag{3.1}$$

where $\mathsf{E}_{n,t}$ is the expectation with respect to the information available at time t for seller n, p the generic price the class indicator is suppressed when there is no ambiguity). If (3.1) holds, then the order will be accepted. Otherwise, if we have,

$$p < \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s}) - v_n(t+1, d_n - 1, \mathbf{s} + \mathbf{e}_n)], \tag{3.2}$$

the order will be rejected.

Remark 3.1.

The operator $\mathsf{E}_{n,t}$ can be treated in a way as an conditional expectation, the information update each time is basically the confirmation that the random variable of each seller's inventory is larger than the cumulative sales, which is updated at the end of each time period.

Proposition 3.1.

The above defined strategy is a unique Nash Equilibrium.

Proof:

To prove that it is a Nash Equilibrium, let us discuss separately for those sellers depends upon their decisions. Suppose that for a particular seller n, the action is to accept, three events can happen,

- sell n is selected, with probability π_n;
- some other seller j in A_t is selected, with probability π_j;
- no seller is selected, with probability $1 - \sum_{A_t} \pi_i$.

Sum them up, the pay-off function has the following form,

$$\pi_n \mathsf{E}_{n,t}[p + v_n(t+1, d_n - 1, \mathbf{s})] + \sum_{j \neq n, j \in A_t} \pi_j \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s} + \mathbf{e}_j)] + \left(1 - \sum_{A_t} \pi_i\right) \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})].$$

If seller n deviates from the strategy, i.e., rejects the demand, its payoff will be,

$$\sum_{j \neq n, j \in A_t} \pi_j \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s} + \mathbf{e}_j)] + \left(1 - \sum_{A_t - \{n\}} \pi_i\right) \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})].$$

From the equation (3.1), we know that seller n could not be better off.

In the case seller n reject, the pay off is,

$$\sum_{j \neq n, j \in A_t} \pi_j \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s} + \mathbf{e}_j)] + \left(1 - \sum_{A_t - \{n\}} \pi_i\right) \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s} + \mathbf{e}_j)].$$

If the seller deviates from this strategy, the pay-off will become,

$$\pi_n [p + \mathsf{E}_{n,t}[v_n(t+1, d_n - 1, \mathbf{s})]] + \sum_{j \neq n, j \in A_t} \pi_j \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s} + \mathbf{e}_j)] + \left(1 - \sum_{A_t} \pi_i\right) \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})].$$

However, we know that $p + \mathsf{E}_{n,t}[v_n(t+1, d_n - 1, \mathbf{s})] < \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})]$, therefore, the seller will be worse off.

Suppose any other strategy that has a seller n, such that, $p + v_n(t+1, d_n - 1, \mathbf{s} + \mathbf{e}_n) < v_n(t+1, d_n \mathbf{s})$, but seller n accepts the demand. We can see that deviation will lead to better pay-off. Meanwhile if there is a seller n with $p + v_n(t+1, d_n - 1, \mathbf{s} + \mathbf{e}_n) \geq v_n(t+1, d_n, \mathbf{s})$, but seller n rejects, a deviation will lead to higher pay-off. ■

The above arguments allow us to present the following dynamic recursion for the value function,

$$v_n(t, d_n, \mathbf{s}) = \sum_{i=1}^{I} \theta_i w_n(t+1, d_n, \mathbf{s}, p_i), \tag{3.3}$$

$$\begin{aligned} w_n(t+1, d_n, \mathbf{s}, p_i) = &\left(1 - \sum_{m=1}^{N} \pi_m\right) \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})] + p_i \pi_n \mathbf{1}_{A_n} \\ &+ \sum_{m=1}^{N} \pi_m [v_n(t+1, d_n, \mathbf{s} + \mathbf{e}_m) \mathbf{1}_{A_m} \\ &+ \mathsf{E}_{n,t}[v_n(t+1, d_n + \mathbf{s}) \mathbf{1}_{A_n}], \end{aligned} \tag{3.4}$$

$$v_N(T, d_n, \mathbf{s}) = \pi_N \mathsf{E}[p], \tag{3.5}$$

with $A_n := \{\mathsf{E}_{n,t}[v_n(t+1, d_n - 1, \mathbf{s} + \mathbf{e}_n) + p_i] \geq \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})\}$.

Remark 3.2.

The information available at time t is on the distribution on the initial capacity of all the other sellers, as well as the sales records in the past period. At time $t+1$, the sales records will be amended with what happened during time period t, the distribution inform hence is naturally updated. For example, if the original distributional estimation is D, and at time t, the total sales has been s, then that information should be updated to $D; D \geq s$. At time t, if there are sales by that seller, it should be updated to $D; D \geq s+1$, otherwise, it will stay at $D; D \geq s$.

4. Monotonicity of the Value Functions and its Implications in Revenue Management

In this section, we will establish monotonicity properties of the value function $v_n(t, d_n, \mathbf{s})$, based on the dynamic recursion formulated in Section 3. The main result is stated in the following theorem, and its proof is presented in Section 5.

Theorem 4.1.

Under the random selection rule with static probabilities, the value function of the knapsack problem $v_n(t, d_n, \mathbf{s})$ for the n-th seller satisfies the following monotonicity properties.

(i) Monotone in inventory $\mathbf{d}$, i.e., $\mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s}) \geq \mathsf{E}_{n,t-1}[v_n(t, d_n - 1, \mathbf{s})]$;

(ii) Monotone in selling amount of competitors,

$$\mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s})] \leq \mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s} + \mathbf{e}_i)];$$

(iii) Monotone in time t, i.e., $\mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s})] \geq \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})]$;

(iv) Second order relationship with respect to d_n, i.e.,

$$\begin{aligned} &\mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s}) - v_n(t, d_n - 1, \mathbf{s} + \mathbf{e}_n)] \\ \geq &\mathsf{E}_{n,t-1}[v_n(t, d_n + 1, \mathbf{s}) - v_n(t, d_n, \mathbf{s} + \mathbf{e}_n)]; \end{aligned} \tag{4.1}$$

(v) Second order relationship with respect to t, i.e.,

$$\begin{aligned} &\mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s})] - \mathsf{E}_{n,t-1}[v_n(t, d_n - 1, \mathbf{s} + \mathbf{e}_n)] \\ &\geq \mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s})] - \mathsf{E}_{n,t}[v_n(t+1, d_n - 1, \mathbf{s} + \mathbf{e}_n)]; \end{aligned} \tag{4.2}$$

(vi) Second order relationship with respect to s, i.e.,

$$\begin{aligned} &\mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s})] - \mathsf{E}_{n,t-1}[v_n(t, d_n - 1, \mathbf{s} + \mathbf{e}_n)] \\ &\geq \mathsf{E}_{n,t-1}[v_n(t, d_n, \mathbf{s} - \mathbf{e}_m)] + \mathsf{E}_{n,t-1}[v_n(t, d_n - 1, \mathbf{s} + \mathbf{e}_n - \mathbf{e}_m)], \end{aligned} \tag{4.3}$$

with $m \neq n$.

Proof:

See the proof in Section 5. ■

From the above theorem, we can draw the following conclusions on the optimal policy.

Corollary 4.1.

For each individual seller, his/her total average revenue is a monotone decrease function of his/her competitors inventory surplus levels.

Proof:

At any time t, given the the amount of remaining inventory being d_n, and s, the total average revenue for seller n takes the function form $v_n(t, d_n, \mathbf{s})$ by definition. Property (ii) in Theorem 4.1 implies that $v_n(t, d_n, \mathbf{s})$ is a component-wise monotone increasing function of the selling quantity s, therefore, $v_n(t, d_n, \mathbf{s})$ is a component-wise monotone decreasing function of the inventory surplus of his/her competitors. ■

Remark 4.1.

It is apparent that the more the overall supply is, the less is the expected marginal gain for each individual unit.

Corollary 4.2.

If it is optimal for a seller to accept the at certain point, then it is also optimal to accept when the competitors have more inventory.

Proof:

Recall that a seller checks the *blance inequality* (3.1), i.e., whether the price will be above the quantity $\mathsf{E}_{n,t}[v_n(t+1, d_n, \mathbf{s}) - v_n(t+1, d_n - 1, \mathbf{s} + \mathbf{e}_n)]$ to decide whether to accept or reject. Property (iv) in Theorem 4.1 says that this quantity decreases when the competitors have more inventory. Therefore, if it is optimal for a seller to accept the at certain point, then it is also optimal to accept when the competitors have more inventory. ■

Remark 4.2.

The intuition is that when there are more inventory in the hands of the competitors, they will be more aggressive, and it will then lower a seller's expected marginal gain. Thus, the seller will be more likely to accept a lower price.

Corollary 4.3.

The quantity $\mathsf{E}_{n,t-1}[v_n(t,d,\mathbf{s})] - \mathsf{E}_{n,t-1}[v_n(t,d-1,\mathbf{s}+\mathbf{e}_n)]$ is a decreasing function of t, and a increasing function of each coordinate of $\mathbf{s}$.

Proof:

The monotonicity with respect to t holds because of (4.2), which is the main conclusion of property (v) in Theorem 4.1. Meanwhile, inequality (4.3) in property (vi) confirms that the quantity $\mathsf{E}_{n,t-1}[v_n(t,d,\mathbf{s})] - \mathsf{E}_{n,t-1}[v_n(t,d-1,\mathbf{s}+\mathbf{e}_n)]$ decreases when $\mathbf{s}$ is replaced by $\mathbf{s}-\mathbf{e}_m$ for any $m \neq n$. ■

Remark 4.3.

Here, it is understood that a lower selling amount of his/her competitors will make a seller less likely to accept a fixed price; conversely, a higher selling amount will make the same seller more likely to accept the same price. Intuitively, observing more sells from ones competitors will make a seller more aggressive. Similarly, approaching the selling deadline will also make a seller to be more aggressive.

5. Proof of Theorem 4.1

Proof:

It is easy to see that (i) and (iii) are trivial. We will prove the rest by backward induction on time t. First, it is trivial to check all of them at the end of selling season, time T. Next, suppose that at time period $t+1$ and later, the properties (ii) and (iv) through (vi) hold. We want to extend all the result to time period t. Since selection is based on the static probabilities, to facilitate our discussion, denote Π_0 the event that no seller is selected, and $\Pi_i, i = 1,\ldots,N$, the event that seller i is selected. From our assumptions, it is clear that the probabilities of these events are $\pi_i, i = 0,1,\ldots,N$, respectively. Furthermore, since all the demand random variables are i.i.d, it is suffice to focus on the event that the price of the demand is $p_i, i = 1,\ldots,N$. We will use a generic notation p to denote the price, for the ease of exposition.

Validity of (ii)

Recall that, we need to establish $\mathsf{E}_{i,t-1}[v_i(t,d_i,\mathbf{s})] \leq \mathsf{E}_{i,t-1}[v_i(t,d_i,\mathbf{s}+\mathbf{e}_j)]$, for $j \neq i$. Without loss of generality, it suffices to show, $\mathsf{E}_{1,t-1}[v_1(t,d_1,\mathbf{s})] \leq \mathsf{E}_{1,t-1}[v_1(t,d_1,\mathbf{s}+\mathbf{e}_j)]$, for any $j > 1$. We will argue that the inequality holds on each event $\Pi_i, i = 0,1,\ldots,N$. On Π_0, since no seller is selected, it is easy to see that the inequality holds by induction, and the induction arguments also applies to $\Pi_i, i \neq 1$ and $i \neq j$. On Π_1, examine what happens at time t, the only case that is not straightforward is that seller one only accept given that the history is $\mathbf{s}$ but reject when it is $\mathbf{s}+\mathbf{e}_j$. In this case, the left hand side (LHS) of the inequality becomes $\mathsf{E}_{1,t}[v_1(t+1, d_1-1, \mathbf{s}+\mathbf{e}_1)]+p$. By induction, it is less than or equal to $\mathsf{E}_{1,t}[v_1(t+1, d_1-1, \mathbf{s}+\mathbf{e}_1+\mathbf{e}_j)]+p$. Meanwhile, $\mathsf{E}_{1,t}[v_1(t+1, d_1-1, \mathbf{s}+\mathbf{e}_j+\mathbf{e}_1))]+p \leq \mathsf{E}_{1,t}[v_1(t+1, d_1, \mathbf{s}+\mathbf{e}_j)]$ due to the fact that this demand is not accepted when the history is $\mathbf{s}+\mathbf{e}_j$. Hence, the inequality follows. On $\Pi_j, j > 1$, there are two cases need to be considered depending on whether seller j accepts the demand. Case I, seller j only accepts when the history is $\mathbf{s}$ not when it is $\mathbf{s}+\mathbf{e}_j$. In this case, we have both the LHS and the right hand side (RHS) equal to $\mathsf{E}_{1,t}[v_1(t+1, d_1, \mathbf{s}+\mathbf{e}_j)]$. Case II, seller j accepts in both cases. Then, the desired inequality is a consequence of $\mathsf{E}_{1,t}[v_1(t+1, d_1, \mathbf{s}+\mathbf{e}_1)] \leq \mathsf{E}_{1,t}[v_1(t+1, d_1, \mathbf{s}+2\mathbf{e}_1)]$, which is the consequence of induction.

Validity of (iv)

Without loss of generality, we only need to show,

$$\begin{aligned}
&\mathsf{E}_{1,t-1}[v_1(t, d_1, \mathbf{s})] - \mathsf{E}_{1,t-1}[v_1(t, d_1-1, \mathbf{s}+\mathbf{e}_1)] \\
\geq &\mathsf{E}_{1,t-1}[v_1(t, d_1+1, \mathbf{s})] - \mathsf{E}_{1,t-1}[v_1(t, d_1, \mathbf{s}+\mathbf{e}_1)].
\end{aligned}$$

Let us first consider case by case based on whether demand will be accepted by seller one. From the induction assumption for time $t+1$, we know that there are only the following cases,

I. the demand is only accepted when the inventory is at d_1+1 not when it is d_1;
II. the demand is accepted when the inventory levels are at both d_1+1 and d_1;
III. the demand is rejected in either case.

And we will discuss each case for events Π_0, Π_1 and $\Pi_j, j > 1$.

In Case I, on event Π_0, the inequality follows from induction, i.e., the concavity with respect to the inventory, at time $t+1$. On the event Π_1, the LHS becomes $\mathsf{E}_{1,t}[v_1(t+1, d_1, \mathbf{s}) - v_1(t+1, d_1-1, \mathbf{s}+\mathbf{e}_1)]$, and the RHS becomes p, then the inequality follows because the balance inequality is violated, which is exactly the reason the demand is not accepted when the inventory is at $(d_1, \mathbf{s})$. On $\Pi_j, j \geq 2$, since the decision of seller j will not depend on the actual amount of inventory seller one has, but just the distribution, the RHS becomes, $\mathsf{E}_{1,t}[v_1(t+1, d_1+1, \mathbf{s}+\mathbf{e}_j) - v_1(t+1, d_1, \mathbf{s}+\mathbf{e}_j+\mathbf{e}_1)]$. Hence, the inequality will follow from the concavity with respect to inventory from time $t+1$ due to induction assumption.

In Case II, again, we only need to look at event Π_1, where the LHS becomes p and the RHS becomes $\mathsf{E}_{1,t}[v_1(t+1, d_1, \mathbf{s}) - v_1(t+1, d_1-1, \mathbf{s}+\mathbf{e}_1)]$, and the inequality follows from the balance inequality. Finally, in Case III, the inequality follows from induction.

Validity of (v)

Again, we need to show that,

$$\begin{aligned}&\mathsf{E}_{1,t-1}[v_1(t,d_1,\mathbf{s})]-\mathsf{E}_{1,t-1}[v_1(t,d_1-1,\mathbf{s}+\mathbf{e}_1)]\\ \geq&\mathsf{E}_{1,t-1}[v_1(t+1,d_1,\mathbf{s})]-\mathsf{E}_{1,t-1}[v_1(t+1,d_1-1,\mathbf{s}+\mathbf{e}_1)].\end{aligned}$$

We will examine the inequality on each event Π_i, $i=0,1,\dots,N$. On Π_0, the inequality follows directly from the induction assumption. On Π_1, let us consider three subcases. First, it is again a straightforward conclusion from the induction assumption if the demand is not accepted for either inventory level. On the other hand if it is accepted for both inventory levels, then the inequality holds due to the induction assumption on the validity of (4) at time t and $t+1$. If seller one only accepts when the inventory level is at d_1, but not when it is at d_1-1, the LHS will become p, then by the condition of accept, i.e., the balance inequality, it is larger than the RHS. On $\Pi_j, j\geq 2$, the inequality follows from the induction assumption on (6) if the demand is accepted for both inventory levels. By the distributional assumption, that is all that needs to be considered.

Validity of (vi)

It is our task to show that, for $j\geq 2$,

$$\begin{aligned}&\mathsf{E}_{1,t-1}[v_1(t,d_1,\mathbf{s})]-\mathsf{E}_{1,t-1}[v_1(t,d_1-1,\mathbf{s}+\mathbf{e}_1)]\\ &\geq \mathsf{E}_{1,t-1}[v_1(t,d_1,\mathbf{s}-\mathbf{e}_j)]-\mathsf{E}_{1,t-1}[v_1(t,d_1-1,\mathbf{s}-\mathbf{e}_j+\mathbf{e}_1)].\end{aligned}$$

On the event Π_1, we know that, by induction assumption, we only need to consider the case that the seller one accepts the demand when the inventory level is at d_1, but not when it is at d_1-1. In this case, the LHS becomes p. For the RHS, consider the two cases that seller one accepts in both cases and only accepts when the inventory is d_1 but not d_1-1. In the first case, it becomes

$$\mathsf{E}_{1,t}[v_1(t+1,d_1-1,\mathbf{s}-\mathbf{e}_j+\mathbf{e}_1)]-\mathsf{E}_{1,t}[v_1(t+1,d_1-2,\mathbf{s}-\mathbf{e}_j+\mathbf{e}_1)].$$

Then the inequality follows from the condition that the seller accepts when the inventory and history is $(d_1-1,\mathbf{s}-\mathbf{e}_j)$. In the second case, both the LHS and RHS become p. Now for the event Π_j, again, the one non-trivial case is similar. Hence, the LHS becomes,

$$\begin{aligned}&\mathsf{E}_{1,t}[v_1(t+1,d_1,\mathbf{s}+\mathbf{e}_j)]-\mathsf{E}_{1,t}[v_1(t+1,d_1-1,\mathbf{s}+\mathbf{e}_j+\mathbf{e}_1)]\\ &\geq \mathsf{E}_{1,t}[v_1(t+1,d_1,\mathbf{s})]-\mathsf{E}_{1,t}[v_1(t+1,d_1-1,\mathbf{s}+\mathbf{e}_1)],\end{aligned}$$

and the inequality thus follows by induction.

This concludes the proof. ∎

6. Conclusions

In this paper, we extended the classic stochastic knapsack problem to model competitions between several sellers and effects on their dynamic pricing decisions. By utilizing dynamic programming techniques, together with a game theoretical model on the sellers' behavior, we are able to identify a simple strategy, i.e., checking the balance inequality, for each seller, and a dynamic recursion for calculating the value functions required. Furthermore, we showed that the value functions have several important first and second order monotonicity properties that are of important theoretical values and critical practical implications.

REFERENCES

Bitran, G. R. and Mondschein, S. V. (1995). An application of yield management to the hotel industry considering multiple day stays, Operations Research, Vol. 43, No. 3, pp. 427–443.

Den Boer, A. V. (2015). Dynamic pricing and learning: Historical origins, current research, and new directions, Surveys in Operations Research and Management Science, Vol. 20, No. 1, pp. 1–18.

Feng, Y. and Gallego, G. (1995). Optimal starting times for end-of-season sales and optimal stopping times for promotional fares, Manage. Sci., Vol. 41, No. 8, pp. 1371–1391.

Feng, Y. and Gallego, G. (2000). Perishable asset revenue management with Markovian time dependent demand intensities, Management Science, Vol. 46, No. 7, pp. 941–956.

Feng, Y. and Xiao, B. (2000). Optimal policies of yield management with multiple predetermined prices, Operations Research, Vol. 48, No. 2, pp. 332–343.

Gallego, G. and Hu, M. (2014). Dynamic pricing of perishable assets under competition, Management Science, Vol. 60, No. 5, pp. 1241–1259.

Gallego, G. and van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons, Manage. Sci., Vol. 40, No. 8, pp. 999–1020.

Lin, G. Y., Lu, Y. and Yao, D. D. (2008). The stochastic knapsack revisited: Switch-over policies and dynamic pricing, Oper. Res., Vol. 56, No. 4, pp. 945–957.

Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*, The MIT Press.

Papastavrou, J. D. Rajagopalan, S. and Kleywegt, A. J. (1996). The dynamic and stochastic knapsack problem with deadlines, Management Science, Vol. 42, No. 12, pp. 1706–1718.

Van Slyke, R. and Young, Y. (2000). Finite horizon stochastic knapsacks with applications to yield management, Operations Research, Vol. 48, No. 1, pp. 155–172.

Zhao, W. and Zheng, Y.-S. (2000). Optimal dynamic pricing for perishable assets with nonhomogeneous demand, Management Science, Vol. 46, No. 3, pp. 375–388.