

Dynamic Capacity Management with General Upgrading

Yueshan Yu

Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130, yuyu@wustl.edu

Xin Chen

Industrial Enterprise and Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, xinchen@illinois.edu

Fuqiang Zhang

Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130, fzhang22@wustl.edu

This paper studies a capacity management problem with upgrading. A firm needs to procure multiple classes of capacities and then allocate the capacities to satisfy multiple classes of customers that arrive over time. A general upgrading rule is considered, i.e., unmet demand can be satisfied using multistep upgrade. No replenishment is allowed and the firm has to make the allocation decisions without observing future demand. We first characterize the structure of the optimal allocation policy, which consists of parallel allocation and then sequential rationing. Specifically, the firm first uses capacity to satisfy the same-class demand as much as possible, then considers possible upgrading decisions in a sequential manner. We also propose a heuristic based on certainty equivalence control to solve the problem. Numerical analysis shows that the heuristic is fast and delivers close-to-optimal profit for the firm. Finally, we conduct extensive numerical studies to derive insights into the problem. It is found that under the proposed heuristic, the value of using sophisticated multistep upgrading can be quite significant; however, using simple approximations for the initial capacity leads to negligible profit loss, which suggests that the firm's profit is not sensitive to the initial capacity decision if the optimal upgrading policy is used.

Keywords: capacity management; inventory; upgrading; dynamic programming; revenue management.

Subject classifications: inventory/production: policies, uncertainty: stochastic.

Area of review: Operations and Supply Chains.

History: Received August 2013; revisions received March 2014, March 2015, August 2015; accepted September 2015.

Published online in *Articles in Advance* November 25, 2015.

1. Introduction

Driven by intensified market competition and rapidly changing consumer trends, many firms have expanded their product lines to cater to different customer segments. On one hand, by offering products with a wide range of quality, design, and characteristics, firms can reach more consumers, generate additional sales, and extract higher profit margins. On the other hand, this has caused significant difficulties in matching supply with demand because the demand is less predictable at the individual segment level than at the aggregate level. Various operational strategies (e.g., postponement, component commonality, modular design) have been proposed for firms to enjoy the benefit of product differentiation while mitigating the risk of mismatches between supply and demand. This paper studies the influential practice of upgrading, where higher-quality products can be used to satisfy demand for a lower-quality product that is sold out. Such a practice takes advantage of risk pooling (product substitution essentially allows product/demand pooling), which results in several immediate benefits: First, it increases revenue by serving more demand; second, it enhances customer service by reducing lost sales;

third, it may lead to lower inventory investment by hedging against demand uncertainty.

The practice of upgrading or substitution has been widely adopted in the business world. In the automobile industry, firms may shift demand for a dedicated capacity to a flexible capacity when the dedicated capacity is constrained (Wall 2003). In the semiconductor industry, faster memory chips can substitute for slower chips when the latter are no longer available (Leachman 1987). More examples in production/inventory control settings can be found in Bassok et al. (1999) and Shumsky and Zhang (2009). Similar practice is ubiquitous in the service industries as well. For instance, airlines may assign business-class seats to economy-class passengers, car rental companies may upgrade customers to more expensive cars, and hotels may use luxury rooms to satisfy demand for standard rooms.

Both practitioners and academics surely understand the importance of the upgrading practice. Substantial research has been conducted on how to manage upgrading in a variety of problem settings. Here we contribute to this large body of literature by studying a dynamic capacity management problem under general upgrading structure. For convenience, we use the terms “product” and “capacity” interchangeably

throughout the paper, and similarly for “upgrading” and “substitution” (strictly speaking, upgrading is one-way substitution). A brief description of our problem is as follows. Consider a firm selling N products with differentiated qualities in a fixed horizon consisting of T periods. There are N classes of customers who arrive randomly in each period. Each customer requests one unit of the product; in the case of a stock-out, the customer can be satisfied with a higher-quality product at no extra charge. Unsatisfied demand is backlogged and the firm incurs a goodwill cost. The firm needs to first determine the procurement quantity of each product at the beginning of the horizon, and then decide how to distribute the products among incoming customers. Because of long ordering lead time, the firm cannot replenish inventory before the end of the horizon; as a result, the firm must dynamically allocate the products over time, before observing future demand.

This paper represents an extension of the recent work by Shumsky and Zhang (2009, referred to as SZ hereafter). As one of the first studies that incorporate dynamic allocation into substitution models, SZ make a simplifying assumption to maintain tractability. Specifically, they consider single-step upgrading, i.e., a demand can only be upgraded by the adjacent product. Clearly, this is a restrictive assumption because in many practical situations firms may have incentives to use multistep upgrading to satisfy demand. Thus there is a need for a theoretical model that captures the realistic upgrading structure. The purpose of this paper is to fill this gap in the literature. While relaxing the single-step upgrading assumption, we attempt to address the following questions as in SZ: What is the optimal initial capacity? How should the products be allocated among customers over time? Are there any effective and efficient heuristics for solving the capacity management problem? The main findings from this paper are summarized as follows.

We start with the dynamic capacity allocation problem. In each period, the firm needs to use the available products to satisfy the realized demand. When a product is depleted while there is still demand for that product, the firm may use upgrading to satisfy customers. How to make such upgrading decisions is a key in substitution models. With the general upgrading structure, the optimal allocation policy is complicated by the fact that the upgrading decisions within a period are interdependent. Under the backlog assumption, we show that a Parallel and Sequential Rationing (PSR) policy is optimal among all possible policies. The PSR policy consists of two stages: In stage 1, the firm uses parallel allocation to satisfy demand as much as possible (i.e., demand is satisfied by the same-class capacity). Then in stage 2, the firm sequentially upgrades left-over demand, starting from the highest demand class; when upgrading a given demand class, the firm starts with the lowest capacity class. The optimality of such a sequential rationing scheme depends on an important property. That is, when using a particular class of capacity to upgrade

demand, the upgrading decision does not depend on the status of the portion of the system below that class. The PSR can greatly reduce the computational complexity because the upgrading decisions do not have to be solved jointly. As an extension, we also consider the multihorizon model with capacity replenishment and show that the PSR policy remains optimal in each horizon. Our theoretical results, though intuitive, turn out to be very challenging to prove. Indeed, our proofs rely on intricate arguments and fully exploit the special structure of the upgrading problem.

Despite the simplified policy structure given by the PSR, solving the problem is still quite involved because of the curse of dimensionality. So there is a need to search for fast heuristics that perform well for the firm. We present a heuristic that adapts certainty equivalence control (CEC) to exploit the sequential rationing property in the PSR policy. Such a heuristic is more appealing than the commonly used CEC heuristic, and we call it refined certainty equivalence control (RCEC) heuristic. Through extensive numerical experiments, we find that the RCEC heuristic delivers close-to-optimal profit for the firm.

The RCEC heuristic enables us to solve large problems effectively. Thus we can use numerical studies based on such a heuristic to derive more insights into the dynamic capacity management problem. First, compared to single-step upgrading, general upgrading (multistep upgrading) can be highly valuable, especially when the capacities are severely unbalanced. Second, our numerical studies indicate that the firm’s profit is not sensitive to the initial capacity decision, given that the optimal upgrading policy is used. For instance, either the newsvendor capacities (calculated assuming no upgrading) or the static capacities (calculated assuming complete demand information) provide nearly optimal profit for the firm. However, the negative impact of using suboptimal allocation policies could be quite significant. These findings suggest that from the practical perspective, deriving the optimal allocation policy should receive a higher priority than calculating the optimal initial capacity.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the model setting. The optimal allocation policy is characterized by §§4 and 5. Section 6 extends the base model to multiple horizons with capacity replenishment. Section 7 proposes the RCEC heuristic and §8 presents the findings from numerical studies. The paper concludes with §9. All proofs are given in the appendix (available as supplemental material at <http://www.doi.org/10.1287/opre.2015.1446>).

2. Literature Review

This paper falls in the vast literature on how to match supply with demand when there are multiple classes of uncertain demand. To facilitate the review, we may divide this literature into four major categories using the following criteria: (1) whether there are multiple capacity types or a

single capacity type; and (2) whether the nature of capacity allocation is static or dynamic. A problem is called static if capacity allocation can be made after observing full demand information. The category that involves the single capacity and static allocation essentially reduces to the newsvendor model that is less relevant. Thus, our review below focuses on the representative studies from the other three categories.

The first category of studies involves multiple capacity types and static capacity allocation. In these studies, firms invest in capacities before demand is realized and then allocate capacities to customers after observing all demand. Because of the existence of multiple capacity types, the issue of substitution naturally arises. Van Mieghem (2003) and Yao and Zheng (2003) provide comprehensive surveys of this category of studies, which can be further divided into two groups. One group of papers study the optimal capacity investment and/or allocation decisions under substitution. Parlar and Goyal (1984) and Pasternack and Drezner (1991) are among the first to consider the simplest substitution structure with two products. Bassok et al. (1999) extend the problem to the general multiproduct case. Hsu and Bassok (1999) introduce random yield into the substitution problem. By assuming single-level substitution, Netessine et al. (2002) study the impact of demand correlation on the optimal capacity levels. Van Mieghem and Rudi (2002) propose the notion of newsvendor networks that consist of multiple newsvendors and multiple periods of demand. Similar settings can be found in the studies on multiperiod inventory models with transshipment, including Robinson (1990), Archibald et al. (1997), and Axsäter (2003). Although these studies involve multiple periods, replenishment is allowed and capacity allocation in each period is made with full demand information. The other group of studies focuses on the value of capacity flexibility. Fine and Freund (1990) and Van Mieghem (1998) consider two types of capacities (dedicated and flexible) and study the optimal investment in flexibility. Bish and Wang (2004) and Chod and Rudi (2005) incorporate pricing decisions when studying the value of resource flexibility. Jordan and Graves (1995) investigate a manufacturing flexibility design problem and discover the well-known chaining rule: Limited capacity flexibility, configured to connect all production facilities and products in a complete chain, can almost deliver the benefit of full flexibility. Their classic work on the design of flexibility has inspired numerous follow-up studies. For example, recently, Chou et al. (2010, 2011) have provided analytical evaluations of the chaining structure for both symmetric and asymmetric problem settings with large scales.

The second category of related literature studies the allocation of a single capacity to multiclass demand in a dynamic environment. This category dates back to the early work by Topkis (1968), who characterizes the optimal rationing policy that assigns capacity to different customer classes over time. Since then similar rationing policies have

been applied to various industry settings. For instance, many revenue management studies focus on how to maximize firms' revenue through capacity rationing when there are multiple fare classes for a single seat type; see Talluri and van Ryzin (2004b) for a review of this literature. A stream of studies on production and inventory control has also derived threshold policies when serving multiple customer classes; see Ha (1997, 2000), de Véricourt et al. (2001, 2002), Deshpande et al. (2003), Savin et al. (2005), Ding et al. (2006) and the references therein.

The third category of studies involves multiple capacity types and dynamic capacity allocation. It differs from the first category mainly in that firms need to allocate capacities to customers without full demand information. There are relatively few papers in this category. Shumsky and Zhang (2009) consider a dynamic capacity management problem with single-step upgrading. They characterize the optimal upgrading policy and provide easy-to-compute bounds for the optimal protection limits that can help solve large problems. Xu et al. (2011) consider a two-product dynamic substitution problem where customers may or may not accept the substitution choice offered by the seller. Our paper extends Shumsky and Zhang (2009) to allow general upgrading. We show that a sequential upgrading policy is optimal for such a problem and provide a fast heuristic that can effectively solve the optimal capacity investment and allocation decisions. Our problem can be framed as a network revenue management model with full upgrading, where the fares are fixed and each demand requests one unit of the corresponding resource (see Gallego and van Ryzin 1997). Gallego and Stefanescu (2009) introduce two continuous optimal control formulations for capacity allocation but concentrate on the analysis of deterministic cases. Steinhardt and Gönsch (2012) study a similar network revenue management problem but allow at most one customer in each period, which means there is at most one upgrading decision in each period. In contrast, we consider multiple interdependent upgrading decisions in each period that involve a much more challenging structure. Our work is also related to the studies on airline revenue management that involve multiple fare products. Talluri and van Ryzin (2004a) study revenue management under a general customer choice model. Zhang and Cooper (2005) consider the selling of parallel flights with dynamic customer choice among the flights. More recent developments include Liu and van Ryzin (2008) and Zhang (2011). In these studies, firms need to decide the subset of products from which a customer can choose; AQ“While” is used only as a temporal term, never causal. Confirm change to “although” here and similar edits throughout.although in our paper, firms decide how to allocate capacities to realized demand. Therefore, both the model settings and results are quite different between these studies and our paper.

3. Model Setting

Consider a firm managing N types of products to satisfy customer demand. The products are indexed in decreasing

quality so that product 1 has the highest quality whereas product N has the lowest quality. There are N corresponding classes of customer demand, i.e., a customer is called class j if she requests product j ($1 \leq j \leq N$). The sales horizon consists of T discrete periods. The initial capacities of the products must be determined prior to the first period and no capacity replenishment is allowed during the sales horizon. (In §6, we extend the model to consider multiple horizons and allow for replenishment.) Customers arrive over time and the demand in each period is random. Let $\mathbf{D}^t = (d_1^t, d_2^t, \dots, d_N^t)^\top \in \mathbb{R}_+^N$ denote the demand vector for period t ($1 \leq t \leq T$), where superscript \top stands for the transpose operation. Throughout the paper we use bold letters for vectors and matrices, and use $(\mathbf{Z})_i$ for the i th component of vector \mathbf{Z} (or $(\mathbf{Z})_{ij}$ for the corresponding element in matrix \mathbf{Z}). For instance, $(\mathbf{D}^t)_i = d_i^t$ is the demand for product i in period t . We assume demand is independent across periods; however, demands for different products within a period can be correlated.

Let r_j be the revenue the firm collects from satisfying a class j customer. If product j is out of stock, then a class j customer could be upgraded at no extra charge by any product i as long as $i < j$. If a class j demand cannot be satisfied in period t , then it will be backlogged to the next period and the firm has to incur a goodwill cost g_j .¹ Define $\mathbf{G} = (g_1, \dots, g_N) \in \mathbb{R}_+^N$. To incorporate service settings like the car rental industry, we include a usage cost denoted by u_i for product i . We make the following assumptions:

ASSUMPTION 1 (A1). $r_1 > r_2 > \dots > r_N$.

ASSUMPTION 2 (A2). $g_1 > g_2 > \dots > g_N$.

ASSUMPTION 3 (A3). $u_1 > u_2 > \dots > u_N$.

We may define $\alpha_{ij} = r_j + g_j - u_i$ ($i \leq j$) as the profit margin for satisfying a class j customer using product i . Based on the above assumptions, we know $\alpha_{ij} > \alpha_{ik}$ and $\alpha_{jk} > \alpha_{ik}$ ($i < j < k$). In other words, for a given capacity, it is more profitable to satisfy a higher class of demand; for a given demand, it is more profitable to use a lower class of capacity. These assumptions are similar to, but more general than, those made in SZ: We have relaxed the single step upgrading assumption in SZ ($\alpha_{ij} > 0$ only if $j = i + 1$) and added Assumption (A2) about the backorder costs. Note that the above assumptions do not require all α_{ij} to be positive. Specifically, if $\alpha_{ij} < 0$ for some i and j , then the assumptions imply that $\alpha_{1j} < \dots < \alpha_{ij} < 0$ and $\alpha_{iN} < \dots < \alpha_{ij} < 0$, which are reasonable in practice.

The firm's objective is to maximize the expected profit over the sales horizon. There are two major decisions for the firm. First, the firm needs to determine the initial capacity before the start of the selling season; second, the firm needs to allocate the available capacities to satisfy demands in each period. Let $\mathbf{C} = (c_1, \dots, c_N) \in \mathbb{R}_+^N$ denote the capacity cost vector, $\mathbf{X}^1 = (x_1^1, x_2^1, \dots, x_N^1)^\top \in \mathbb{R}_+^N$ the starting capacities in period 1, and $\tilde{\mathbf{D}}^t = (\tilde{d}_1^t, \tilde{d}_2^t, \dots, \tilde{d}_N^t)^\top \in \mathbb{R}_+^N$ the backordered demand at the beginning of period t . We

use \mathbf{Y}^t for the capacity allocation matrix in period t , i.e., $(\mathbf{Y}^t)_{ij} = y_{ij}^t$ is the amount of product i offered to satisfy class j demand ($y_{ij} = 0$ if $i > j$). Define $\Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ as the optimal revenue-to-go function in period t given the state variable $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$. Then the buyer's problem can be formulated as follows:

$$\max_{\mathbf{X}^1 \in \mathbb{R}_+^N} \Pi(\mathbf{X}^1) = \Theta^1(\mathbf{X}^1, \mathbf{0}) - \mathbf{C} \cdot \mathbf{X}^1, \quad (1)$$

and for each period t ($1 \leq t \leq T$):

$$\begin{aligned} \Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t) &= \mathbb{E}_{\mathbf{D}^t} \{ \Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t | \mathbf{D}^t) \} \\ &= \mathbb{E}_{\mathbf{D}^t} \left\{ \max_{\mathbf{Y}^t} [H(\mathbf{Y}^t | \tilde{\mathbf{D}}^t; \mathbf{D}^t) + \Theta^{t+1}(\mathbf{X}^{t+1}, \tilde{\mathbf{D}}^{t+1})] \right\}, \end{aligned} \quad (2)$$

where

$$H(\mathbf{Y}^t | \tilde{\mathbf{D}}^t; \mathbf{D}^t) = \sum_{1 \leq i \leq j \leq N} \alpha_{ij} y_{ij}^t - \mathbf{G} \cdot (\tilde{\mathbf{D}}^t + \mathbf{D}^t), \quad (3)$$

$$\mathbf{X}^{t+1} = \mathbf{X}^t - \mathbf{Y}^t \cdot \mathbf{1} \geq \mathbf{0}, \quad (4)$$

$$\tilde{\mathbf{D}}^{t+1} = \tilde{\mathbf{D}}^t + \mathbf{D}^t - (\mathbf{Y}^t)^\top \cdot \mathbf{1} \geq \mathbf{0}, \quad (5)$$

$$\mathbf{Y}^t \geq \mathbf{0}, \quad \mathbf{1} = (1, 1, \dots, 1)^\top.$$

We assume leftover products have zero value and unmet demand is lost at the end of the selling season, so $\Theta^{T+1} \equiv 0$. Note that the optimal revenue-to-go function $\Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ is recursively defined in (2). Given the allocation decision \mathbf{Y}^t , $H(\mathbf{Y}^t | \tilde{\mathbf{D}}^t; \mathbf{D}^t)$ in (3) denotes the single period revenue, which is the difference between the upgrading revenue and the goodwill cost. The state transition between two consecutive periods is governed by (4) and (5), which represent two constraints for the allocation decision \mathbf{Y}^t in period t .

4. Parallel and Sequential Rationing (PSR)

This section starts analyzing the upgrading problem given in (1). Here we introduce several useful definitions and qualitatively characterize the optimal allocation policy. The formal optimality proof will be presented in the next section. As the first step, since

$$\Pi(\mathbf{0}) = -\mathbf{G} \cdot \sum_{t=1}^T (T+1-t) \mathbb{E}[\mathbf{D}^t] > -\infty, \quad (6)$$

$$\lim_{\mathbf{X}^1 \rightarrow \infty} \Pi(\mathbf{X}^1) = \sum_{t=1}^T \sum_{i=1}^N (r_i - u_i) \mathbb{E}[d_i^t] - \lim_{\mathbf{X}^1 \rightarrow \infty} \mathbf{C} \cdot \mathbf{X}^1 = -\infty,$$

and the fact that $\Pi(\mathbf{X}^1)$ is continuous in $\mathbf{X}^1 \in \mathbb{R}_+^N$, we know there exists a finite $\mathbf{X}^* \in \mathbb{R}_+^N$ that solves the optimization problem in (1).

From Murty (1983) and Rockafellar (1996), for any demand realization \mathbf{D}^T in period T , it is straightforward

to see $\Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t | \mathbf{D}^t)$ is concave in the state variable $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$, which are the right-hand side variables in the linear program defined by (2). Since concavity is preserved under the expectation operation on \mathbf{D}^t ($1 \leq t \leq T$) and the maximization operation with respect to the allocation decision \mathbf{Y}^t (see, e.g., Simchi-Levi et al. 2014, Propositions 2.1.3 and 2.1.15), Θ^t is again concave in $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ in each period t . Clearly, the revenue function

$$\hat{\Theta}^t(\mathbf{Y}^t | \mathbf{X}^t, \tilde{\mathbf{D}}^t; \mathbf{D}^t) = H(\mathbf{Y}^t | \tilde{\mathbf{D}}^t; \mathbf{D}^t) + \Theta^{t+1}(\mathbf{X}^{t+1}, \tilde{\mathbf{D}}^{t+1}), \quad (7)$$

in period t given state $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ and demand realization \mathbf{D}^t , is also concave in the allocation decision \mathbf{Y}^t . The concavity property is summarized in the following proposition whose formal proof is omitted.

PROPOSITION 1. *In period t , $\Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ is concave in $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$, and $\hat{\Theta}^t(\mathbf{Y}^t | \mathbf{X}^t, \tilde{\mathbf{D}}^t; \mathbf{D}^t)$ is concave in \mathbf{Y}^t .*

Notice that the allocation decision \mathbf{Y}^t is constrained by a bounded polyhedron defined by (4)–(5) and Θ^t in (7) is continuous in \mathbf{Y}^t . Thus, there always exists an optimal allocation to the general upgrading problem in each period t . For a given state $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ and demand realization \mathbf{D}^t , there are two types of decisions: parallel allocations y_{ii}^t for all i ($1 \leq i \leq N$) and upgrading decisions y_{ij}^t for classes i and j ($1 \leq i < j \leq N$). These are dynamic decisions because they will not only determine revenue H in the current period but also affect future revenue $\Theta^{t+1}(\mathbf{X}^{t+1}, \tilde{\mathbf{D}}^{t+1})$.

It is straightforward to solve the parallel allocation problem. In our model, the maximum revenue we can get from a unit of capacity i is α_{ii} through the parallel allocation, i.e., capacity i is used to fulfill demand class i . It is suboptimal to satisfy demand from lower classes using capacity i when there is still unmet demand i . Furthermore, the expected value of carrying over capacity i to the next period will not exceed α_{ii} , either. Hence the optimal strategy is to use the parallel allocation as much as possible, which implies $y_{ii}^t = \min(d_i^t + \tilde{d}_i^t, x_i^t)$. Another implication is that in the state variable $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$, class i ($1 \leq i \leq N$) cannot be positive in both \mathbf{X}^t and $\tilde{\mathbf{D}}^t$. Thus, we can use a single variable $\mathbf{M}^t = (\mathbf{X}^t - \tilde{\mathbf{D}}^t) = (m_1^t, m_2^t, \dots, m_N^t)^T$ to represent the state at the beginning of period t (before the parallel allocation): $m_i^t > 0$ means there is positive capacity for i whereas $m_i^t < 0$ means there is backordered demand for i . In the rest of the paper we will use \mathbf{M}^t and $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ exchangeably.

The more challenging question is how to make the upgrading decisions after the parallel allocation. The state after the parallel allocation in period t is $(m_1^t - d_1^t, m_2^t - d_2^t, \dots, m_N^t - d_N^t)^T$. Note that $m_i^t - d_i^t > 0$ means that there is leftover capacity i , whereas $m_i^t - d_i^t < 0$ implies that there is unsatisfied demand i and capacity i must have been depleted. The firm needs to decide how much demand should be upgraded using higher class capacities. This is equivalent to a rationing problem, i.e., how much capacity

should be protected to satisfy future demand. The upgrading problem in our model is different from the one studied in SZ. Particularly, with the single-step assumption in SZ, when capacity i is depleted, classes above i and those below i become independent of each other in future periods, and thus the upgrading problem is greatly simplified because all the upgrading decisions can be solved separately. However, with the general upgrading structure in our model, the upgrading decisions after parallel allocation are no longer isolated. In this case, we may have to solve all decisions jointly, which could be computationally intensive. Fortunately, close scrutiny shows that the following two observations can greatly reduce the complexity of the upgrading problem. The intuition and formal proofs of these observations will be presented in §5.1.

First, the upgrading decision y_{ij}^t of using capacity i to upgrade demand j is independent of the demands and the capacities below class j . Second, for demand class j with $m_j^t - d_j^t < 0$, the upgrading decisions y_{ij}^t , $i = 1, \dots, j-1$ can be solved sequentially in i starting from the lowest class i ($i < j$) with positive capacity. Based on these observations, the upgrading problem can be sequentially solved as follows:

Step 1. Identify the smallest j ($1 \leq j \leq N$) with $m_j^t - d_j^t < 0$ (the highest class with unmet demand);

Step 2. For the largest i (the lowest capacity class) less than j with $m_i^t - d_i^t > 0$, determine the upgrading quantity y_{ij}^t in period t (or equivalently, the quantity of capacity i to be protected for the next period). When solving y_{ij}^t , we can ignore the classes lower than j ;

Step 3. Repeat Step 2 until all capacity classes available for upgrading demand j have been considered;

Step 4. Repeat Step 1 until all unmet demand classes have been considered.

To summarize, the firm may allocate capacity using the so-called Parallel and Sequential Rationing (PSR) policy. Under such a policy, the firm first performs the parallel allocation on each class to satisfy new demands, and then sequentially decides upgrading quantities for classes with unmet demand. For easy reference, hereafter we refer to the aforementioned allocation policy as the PSR policy.

The most crucial decision in the sequential upgrading procedure is to determine y_{ij}^t in Step 2. Consider the decision about how much capacity i should be used to upgrade demand j . It is clear that as long as the current upgrade revenue α_{ij} is greater than the expected marginal value in the future, capacity i should be used to upgrade demand j . Such an upgrading or rationing decision essentially specifies the protection levels for the capacities. Let p_{ij} be the optimal protection level of capacity i with respect to demand j , i.e., the firm should stop upgrading demand j by capacity i when the capacity level of i drops to p_{ij} . Since $\Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ is concave in $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ by Proposition 1, the expected marginal value of capacity i is monotonically increasing as capacity i decreases. Hence, the protection level p_{ij} in period t is the unique lower bound above which

using capacity i to upgrade demand j is profitable. Define $(\partial/\partial p)\Theta^t = [(\partial/\partial p^+)\Theta^t, (\partial/\partial p^-)\Theta^t]$ as the subdifferential of Θ^t with respect to some variable p , where $(\partial/\partial p^-)\Theta^t$ and $(\partial/\partial p^+)\Theta^t$ are the left and right derivatives, respectively. Let

$$\mathbf{N}^t = (n_1^t, n_2^t, \dots, n_N^t)^\top$$

denote the state of the system immediately before the epoch of determining y_{ij}^t . Specifically, \mathbf{N}^t is a generic system state and can be recursively defined, i.e., \mathbf{N}^t could be the resulting state after the parallel allocation and possibly some upgrading allocations as well (recall the upgrading decisions are made sequentially starting from the highest demand class). The optimal protection levels can be defined as follows.

DEFINITION 1. The optimal protection level $p_{ij} \geq 0$ under state $\mathbf{N}^t = (n_1^t, n_2^t, \dots, n_N^t)^\top$ is defined as

$$p_{ij} = \begin{cases} p & \text{if } \alpha_{ij} \in \frac{\partial}{\partial p} \Theta^{t+1}(n_1^t, \dots, n_{i-1}^t, p, 0, \dots, \\ & 0, -p, n_{j+1}^t, \dots, n_N^t), \\ 0 & \text{if } \alpha_{ij} > \frac{\partial}{\partial p^+} \Theta^{t+1}(n_1^t, \dots, n_{i-1}^t, p, 0, \dots, 0, \\ & -p, n_{j+1}^t, \dots, n_N^t) \Big|_{p=0}. \end{cases} \quad (8)$$

With the protection levels p_{ij} and \mathbf{N}^t , the optimal upgrading decision y_{ij}^t is simply given by $\min((n_i^t - p_{ij})^+, (-n_j^t)^+)$ where $(x)^+ = \max(x, 0)$. Notice that there are 0's between classes i and j since the PSR policy does not consider y_{ij}^t if there exists a class s ($i < s < j$) with positive capacity or unmet demand. When class s has a positive capacity, it is more profitable to upgrade demand j with capacity s instead of capacity i , and it is unnecessary for us to consider y_{ij}^t if there is capacity s remaining after solving y_{sj}^t . When there is unmet demand for class s , capacity i should upgrade demand s first, and it would be suboptimal to upgrade demand j if class s still has unmet demand after upgrading y_{is}^t .

Before presenting the main results, we wish to further reduce the computation in the general upgrading problem by exploring its structure. With the single-step upgrading rule, SZ shows that whenever a capacity (say, i) is depleted, the entire problem decouples into two independent subproblems, where the first subproblem consists of products above i and the second consists of products below i (see Lemma 4 in SZ). Under the general upgrading rule, such a property in SZ clearly does not hold. However, it can be shown that under a similar but stronger condition, our problem can also be separated into independent subproblems, as stated in the next lemma.

LEMMA 1. Consider an N -class general upgrading problem with state $\mathbf{N}^t = (n_1^t, n_2^t, \dots, n_N^t)^\top$ in period t . If $\sum_{s=k}^i n_s^t \leq 0$ for all class $k \leq i$, then the problem can be separated into two independent subproblems: an upper part consisting of classes $(1, \dots, i)$, and a lower part consisting of classes $(i + 1, \dots, N)$.

Lemma 1 represents a generalized version of the separation property in SZ. For convenience, we say class i is separable if it satisfies the condition stated in Lemma 1. Notice that unlike the property in SZ, $n_i^t \leq 0$ is not enough to split the N -class general upgrading problem since there may be class k ($k < i$), which can upgrade demands in classes $(i + 1, \dots, N)$. However, the condition in Lemma 1 ensures that none of classes $(1, \dots, i)$ has enough capacity to upgrade the demand in $(i + 1, \dots, N)$ when optimal upgrading is performed. Specifically, there may exist class $k < i$ with positive capacity that can upgrade the demand in $(i + 1, \dots, N)$, but it is more profitable for capacity k to satisfy the demand in classes $(k + 1, \dots, i)$ first, which will consume all of class k 's capacity. Therefore, Lemma 1 asserts that the entire upgrading problem can be simplified by decomposition under certain conditions. That is, the profit of the N -class problem can be written as the sum of the profits from independent subproblems $(1, \dots, i)$ and $(i + 1, \dots, N)$ whenever class i is separable. This observation will allow us to present the optimality proof in §5.1 on each of the nonseparable subproblems, significantly simplifying the exposition.

5. Optimality and Properties of PSR

5.1. Optimality

We now present the optimality proof of the PSR policy. The proof essentially shows that two results hold for each period by induction: First, the marginal value of Θ^t in (8) has certain properties; second, the PSR policy is optimal because of these properties. As a preparation for the optimality proof, we introduce the concepts of greedy upgrading and effective state. By greedy upgrading, we refer to a PSR step with zero protection levels. That is, after the parallel allocation, the unmet demand will be sequentially upgraded as much as possible. In addition, we define the effective state as follows.

DEFINITION 2. Consider a state vector $\mathbf{N}^t = (n_1^t, n_2^t, \dots, n_N^t)$ in period t ($1 \leq t \leq T$). For class r ($1 \leq r \leq N$), the effective state $\hat{\mathbf{N}}_r^t = (\hat{n}_1^t, \dots, \hat{n}_r^t, n_{r+1}^t, \dots, n_N^t)$ is defined as the resulting state after applying the greedy upgrading for classes $(1, \dots, r)$.

In fact, given any state \mathbf{N}^t and its effective state $\hat{\mathbf{N}}_r^t$ for some class r , let h ($1 \leq h \leq r$) denote the highest class with $\hat{n}_h^t > 0$, then class $h - 1$ is separable in \mathbf{N}^t . To see this, note that given $\hat{n}_h^t > 0$, there is no upgrade between classes $(1, \dots, h - 1)$ and (h, \dots, r) when applying the greedy upgrading. Thus, for all class $k < h$, we have $\sum_{s=k}^{h-1} n_s^t \leq \sum_{s=k}^{h-1} \hat{n}_s^t \leq 0$, where the first inequality holds because there may be upgrade between classes $(1, \dots, k - 1)$ and $(k, \dots, h - 1)$ when performing the greedy upgrading, and the second inequality follows from the definition of class h . Hence, according to Lemma 1, $h - 1$ is separable, and classes $(1, \dots, h - 1)$ can be ignored in the subsequent allocation decisions.

Consider a state vector $\mathbf{N}^t = (n_1^t, \dots, n_N^t)$ in period t . For $1 \leq i < j \leq N$, define

$$\Delta_{ij}^{+-} \Theta^t(\mathbf{N}^t) = \frac{\partial}{\partial n_i^+} \Theta^t(\mathbf{N}^t) - \frac{\partial}{\partial n_j^-} \Theta^t(\mathbf{N}^t),$$

$$\Delta_{ij}^{-+} \Theta^t(\mathbf{N}^t) = \frac{\partial}{\partial n_i^-} \Theta^t(\mathbf{N}^t) - \frac{\partial}{\partial n_j^+} \Theta^t(\mathbf{N}^t).$$

Notice that the protection level p_{ij} in (8) can be equivalently defined as

$$\Delta_{ij}^{+-} \Theta^{t+1}(\mathbf{N}^t)|_{p=p_{ij}} \leq \alpha_{ij} \leq \Delta_{ij}^{-+} \Theta^{t+1}(\mathbf{N}^t)|_{p=p_{ij}},$$

where $\mathbf{N}^t = (n_1^t, \dots, n_{i-1}^t, p, 0, \dots, 0, -p, n_{j+1}^t, \dots, n_N^t)$. Thus, the properties of $\Delta_{ij}^{+-} \Theta^{t+1}(\mathbf{N}^t)$ and $\Delta_{ij}^{-+} \Theta^{t+1}(\mathbf{N}^t)$ will have useful implications on the optimal protection levels.

Now we are in the position to prove the optimality of the PSR by induction.

PROPOSITION 2. *Consider an N -class general upgrading problem in period t ($1 \leq t \leq T$) with state vector \mathbf{N}^t , where $(n_{i+1}^t, \dots, n_{j-1}^t) \leq 0$ and $n_j^t < 0$.*

1. *We have*

$$\begin{aligned} \Delta_{ij}^{+-} \Theta^{t+1}(\mathbf{N}^t) &= \Delta_{ij}^{+-} \Theta^{t+1}(\hat{\mathbf{N}}_{i-1}^t), \\ \Delta_{ij}^{-+} \Theta^{t+1}(\mathbf{N}^t) &= \Delta_{ij}^{-+} \Theta^{t+1}(\hat{\mathbf{N}}_{i-1}^t), \end{aligned} \tag{9}$$

both of which are independent of the values of (n_j^t, \dots, n_N^t) .

2. *The PSR policy solves the general upgrading problem in period t .*

Proposition 2 deserves some discussion. The first part of Proposition 2 states that the optimal protection level p_{ij} in period t is independent of the values of (n_j^t, \dots, n_N^t) , although it is affected by the classes above i through the effective state $(\hat{n}_1^t, \dots, \hat{n}_{i-1}^t)$. These results provide the rationale behind the sequential rationing in the PSR policy. We offer the following intuitive explanation of these results. First, we explain why $\Delta_{ij}^{+-} \Theta^{t+1}(\mathbf{N}^t)$ and $\Delta_{ij}^{-+} \Theta^{t+1}(\mathbf{N}^t)$ are independent of (n_j^t, \dots, n_N^t) . Before deciding p_{ij} or y_{ij}^t , without loss of generality, suppose there are p units of capacity i lined up to upgrade unmet demand j . Meanwhile, the unmet demand in class j can be treated as a waiting line, which will be satisfied in the first-come first-served sequence. To determine p_{ij} we need to compare α_{ij} with the expected value of each unit of capacity i . Consider the capacity unit 1, i.e., the first unit in front of the capacity line. Given the backorder assumption, capacity unit 1 can only satisfy either a future demand in classes $(i, \dots, j-1)$ or the first unit in the waiting line of demand j . Hence, the expected value of capacity unit 1 in class i is independent of states $(n_{j+1}^t, \dots, n_N^t)$. Furthermore, the above argument only relies on the fact that there exists unmet demand j . Thus, the expected value of capacity unit 1 is also independent of n_j^t , the length of the waiting line of demand j . Such an argument applies to any other capacity units of class i .

Next, we explain the equalities in (9). Note that for any class k ($1 < k < i$) with positive capacity, it would not upgrade demand from class i or below in any optimal policy if there exists class r ($k < r < i$) with backordered demand or positive capacity. In fact, it would be more valuable to use capacity k to upgrade demand r in the former case and to use capacity r to upgrade demand from class i or below in the latter case. The remaining capacity of class k after upgrading all backordered demands in classes $(k+1, \dots, i-1)$ equals \hat{n}_k^t as defined in the effective state. Therefore, the expected future value of capacity i in period t should equivalently depend on the effective state $(\hat{n}_1^t, \dots, \hat{n}_{i-1}^t)$.

Finally, for any given period t under the PSR policy, the effective states of all intermediate states for classes $(1, 2, \dots, i-1)$ are the same before we exhaust the capacity of class i . Thus, Proposition 2 implies that when solving p_{ij} , it is sufficient to use the first $i-1$ components of $\mathbf{M}^t - \mathbf{D}^t$ (i.e., the state of the system in period t after the parallel allocation) instead of \mathbf{N}^t (i.e., the state of the system prior to deciding y_{ij}^t) in the PSR policy. This is a unique and interesting property of the general upgrading problem, allowing us to simultaneously and independently solve all protection levels based on $\mathbf{M}^t - \mathbf{D}^t$. Specifically, for any classes i and j ($1 \leq i < j \leq N$) with $n_i^t > 0$ and $n_j^t < 0$, the protection level p_{ij} can be immediately determined by $(\partial/\partial p) \Theta^{t+1}(m_1^t - d_1^t, \dots, m_{i-1}^t - d_{i-1}^t, p, 0, \dots, 0, -p, 0, \dots, 0)$.

5.2. Properties of Protection Levels

After establishing the optimality of the PSR policy, we explore some important properties related to the optimal protection levels from the PSR policy.

First, if both the initial capacity \mathbf{X}^1 and all demands are integer valued, similar to SZ, we can prove that there exists an integer valued optimal policy generated by the PSR policy.

PROPOSITION 3. *If initial capacity \mathbf{X}^1 and demand $\mathbf{D}^1, \dots, \mathbf{D}^T$ are integer valued, there exists an integer valued optimal policy $\mathbf{Y}^1, \dots, \mathbf{Y}^T$ derived by the PSR policy.*

To further characterize the protection level p_{ij} defined in (8), we need to deal with the marginal value of Θ^t with respect to each capacity level and unmet demand level. Intuitively, one may think that the profit will be higher if there is an additional unit of capacity $i-1$ ($1 < i \leq N$) rather than capacity i . But this is not necessarily true. When making upgrading decisions, one more unit of capacity from the higher class $i-1$ always provides more flexibility, but such a flexibility does not necessarily mean a higher profit since $\alpha_{ij} > \alpha_{i-1,j}$ ($i < j$) by our assumption. Similarly, one more unit of demand in a lower class, which can be upgraded by more classes of capacities, has similar advantage but can not guarantee greater profit because $\alpha_{ij} > \alpha_{i,j+1}$ ($i \leq j$). However, we can show two different monotone properties of the protection levels. First, since

lower demand has less value for any capacity, the protection level should increase in the class index of demand.

PROPOSITION 4. *For the same $(n_1^t, \dots, n_{i-1}^t)$ in period t ($1 \leq t \leq T$), $p_{ij} \leq p_{i,j+1}$ when $i < j$.*

Because the general upgrading problem in period T is a transportation problem, $\Theta^T(\mathbf{X}^T, \tilde{\mathbf{D}}^T)$ is submodular in $(\mathbf{X}^T, -\tilde{\mathbf{D}}^T)$ (see Topkis 1998). This implies the protection level p_{ij} in period $T - 1$ under state \mathbf{N}^{T-1} is decreasing in $(n_1^{T-1}, \dots, n_{i-1}^{T-1})$. In fact, the same monotonicity holds in earlier periods.

PROPOSITION 5. *The optimal protection level p_{ij} ($1 \leq i < j \leq N$) in period t ($1 \leq t \leq T$) are component-wise decreasing in $(n_1^t, \dots, n_{i-1}^t)$.*

For any class i ($1 \leq i \leq N$), this result assures that the more capacities (or less back-ordered demands) in classes higher than i , the more upgrades can be offered by class i . Note that larger $(n_1^t, \dots, n_{i-1}^t)$ means higher probability of demand i being upgraded in remaining periods, which decreases the expected marginal value of capacity i and gives class i a greater incentive to upgrade lower demands in the current period.

It is noteworthy that although the result for the last period can be proved using lattice programming in Topkis (1998), the commonly used preservation property of supermodularity under maximization operations, Theorem 2.7.6 in Topkis (1998), does not apply. In fact, the revenue-to-go function $\Theta^T(\mathbf{X}^T, \tilde{\mathbf{D}}^T)$ in period T , used in the dynamic program recursion for period $T - 1$, is the optimal objective value of a transportation problem that is submodular in $(\mathbf{X}^T, -\tilde{\mathbf{D}}^T)$. Therefore, our proof relies heavily on the structure of the general upgrading problem and fully utilizes the optimality of the PSR policy.

One may ask whether the optimal protection levels are decreasing over time, i.e., the protection level would be lower if there are fewer periods to go. Interestingly, though this is true in SZ, it does not hold in our upgrading problem. This is mainly due to the existence of the backorder cost. Note that the purpose of the protection levels is to balance the goodwill loss of carrying backorders and the revenue loss of losing future demand from the same class. For early periods that are still far away from the end of the horizon, because a backorder causes the goodwill loss in each period until it is upgraded, the protection levels may be lower to avoid high backorder costs; in contrast, when it is close to the end of the horizon, the protection levels may come back up because carrying backorders will be less costly.

We may use a two-product three-period example to explain this counter-intuitive result. Let $(2, -2)$ be the state after the parallel allocation, $\mathbf{D}^2 = (0, 0)$ and $\mathbf{D}^3 = (1, 0)$ with probability 1. Working backward to solve the p_{12} in period 2, since

$$\Theta^3(2, -2) - \Theta^3(1, -1) = \alpha_{12} - g_2 < \alpha_{12},$$

$$\Theta^3(1, -1) - \Theta^3(0, 0) = \alpha_{11} - g_2,$$

we have $p_{12} = 1$ in period 2 if $\alpha_{11} - g_2 > \alpha_{12}$. Since $\mathbf{D}^2 = (0, 0)$, there is

$$\Theta^2(2, -2) - \Theta^2(1, -1) = \alpha_{12} - g_2 < \alpha_{12},$$

$$\Theta^2(1, -1) - \Theta^2(0, 0) = \alpha_{11} - 2g_2.$$

Therefore, if $\alpha_{11} - g_2 > \alpha_{12} > \alpha_{11} - 2g_2$, the optimal protection level p_{12} increases from 0 in period 1 to 1 in period 2. That is, the protection level does not necessarily decrease over time in our general upgrading problem.²

6. Multiple Horizons with Capacity Replenishment

Now we extend our model to multiple horizons with capacity replenishment. Specifically, there are K ($K \geq 1$) horizons, each consisting of T periods. Demands across horizons are independent and identically distributed. At the beginning of each horizon k ($1 \leq k \leq K$), the firm observes the leftover capacity \mathbf{X} and unmet demand $\tilde{\mathbf{D}}$ carried over from the previous horizon. There are two decisions for the firm in each horizon: First, the firm decides how much capacity to replenish; second, it allocates capacity to satisfy demand as formulated in (2). For completeness, we assume unmet demand after the K th horizon can also be satisfied by purchasing additional capacity. There is a unit cost vector $\mathbf{C} = (c_1, \dots, c_N) \in \mathfrak{R}_+^N$ for capacity replenishment. The remaining capacity at the end of each horizon incurs a holding cost $\mathbf{h} = (h_1, \dots, h_N) \in \mathfrak{R}_+^N$. The leftover capacity after the K th horizon can be sold at the initial capacity cost, i.e., it has salvage value \mathbf{C} . Revenues and costs are discounted at a rate γ ($0 < \gamma \leq 1$) for each horizon. We introduce a discount factor here because a horizon might be long enough so that the time value of money should be taken into account. The rest of the model setting remains the same as in §3.

We show the optimality of a myopic policy in which the firm only needs to solve a single-horizon problem. For this purpose, define

$$\begin{aligned} &\Theta^{T+1}(\mathbf{X}^{T+1}, \tilde{\mathbf{D}}^{T+1}) \\ &= (\gamma\mathbf{C} - \mathbf{h}) \cdot \mathbf{X}^{T+1} + \gamma(\boldsymbol{\alpha} - \mathbf{C}) \cdot \tilde{\mathbf{D}}^{T+1}, \end{aligned} \quad (10)$$

where $\boldsymbol{\alpha} = (\alpha_{11}, \dots, \alpha_{NN})$ is the revenue from parallel allocation. Since we assume the leftover capacity after the K th horizon can be sold at the initial capacity cost and unmet demand after the K th horizon can also be satisfied by purchasing additional capacity, at the end of the K th horizon, the unit salvage value of the remaining capacity \mathbf{X}^{T+1} is $\gamma\mathbf{C} - \mathbf{h}$, and the value of the back-ordered demand \mathbf{D}^{T+1} is $\gamma(\boldsymbol{\alpha} - \mathbf{C})$. That is, the terminal value at the end of the K th horizon is exactly given by $\Theta^{T+1}(\mathbf{X}^{T+1}, \tilde{\mathbf{D}}^{T+1})$ if the remaining capacity is \mathbf{X}^{T+1} and the back-ordered demand is $\tilde{\mathbf{D}}^{T+1}$, which is essential for the optimality of the myopic policy (similar terminal value assumptions have been commonly used in the literature). Let $\Pi(\mathbf{X}; \gamma\mathbf{C} -$

$\mathbf{h}; \gamma(\boldsymbol{\alpha} - \mathbf{C})$) denote the optimal profit of a single-horizon model with initial capacity \mathbf{X} and a terminal value specified by $\Theta^{T+1}(\mathbf{X}^{T+1}, \tilde{\mathbf{D}}^{T+1})$. Note that although $\gamma\mathbf{C} - \mathbf{h}$ and $\gamma(\boldsymbol{\alpha} - \mathbf{C})$ are constants, they are used as arguments in Π to emphasize the terminal values of the remaining capacity and back-ordered demand. This is different from the single-horizon model with zero end values in §3.

From the proof of Proposition 1, $\Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$, which is similarly defined as (2) with $\Theta^{T+1} \equiv \mathbf{0}$ being replaced by Θ^{T+1} in (10), is still concave in $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$. In particular, $\Pi(\mathbf{X}; \gamma\mathbf{C} - \mathbf{h}; \gamma(\boldsymbol{\alpha} - \mathbf{C}))$ is concave in \mathbf{X} from the concavity of $\Theta^1(\mathbf{X}, \mathbf{0})$. Furthermore, similarly as (6), we can show that there exists an optimizer \mathbf{X}^* for the concave function $\Pi(\mathbf{X}; \gamma\mathbf{C} - \mathbf{h}; \gamma(\boldsymbol{\alpha} - \mathbf{C}))$:

$$\mathbf{X}^* \in \arg \max_{\mathbf{X} \in \mathbb{R}_+^N} \Pi(\mathbf{X}; \gamma\mathbf{C} - \mathbf{h}; \gamma(\boldsymbol{\alpha} - \mathbf{C})). \quad (11)$$

Note that \mathbf{X}^* is the optimal capacity level for the replenishment model with $K = 1$.

The next proposition characterizes the optimal capacity replenishment and allocation policies in the multihorizon model, given that the firm starts with an initial capacity $\mathbf{X} \leq \mathbf{X}^*$. It shows that the structural results from the base model in §3 remain valid in the multihorizon model; thus we will focus on the base model in the rest of the paper.

PROPOSITION 6. *Suppose the firm starts with an initial capacity $\mathbf{X} \leq \mathbf{X}^*$ in (11). The firm's optimal replenishment policy in the multihorizon model is a base stock policy with the optimal base stock level \mathbf{X}^* . Furthermore, the PSR algorithm solves the optimal allocation decisions within each horizon.*

7. Heuristics and Benchmark Models

So far we have characterized the structure of the optimal allocation policy for our dynamic capacity management problem. In this section, we propose an effective heuristic for solving the optimal allocation policy. For future comparison, we also present two benchmark models that are simplified versions of the general upgrading problem.

7.1. Heuristics

We have shown that the PSR policy yields the optimal allocation decisions \mathbf{Y}^t for the firm in period t , which essentially consist of the optimal protection levels for each capacity class. These optimal protection levels are defined by (8) and can be solved by backward induction. For instance, the optimal protection levels in period t depend on the revenue-to-go function Θ^{t+1} , which is determined by the protection levels used in period $t + 1$. To evaluate Θ^{t+1} , one needs to derive the optimal protection levels for all possible states in period $t + 1$ (note that these protection levels, though possessing the appealing properties established earlier, are still state-dependent). Because of the curse of dimensionality, solving the exact optimal

upgrading decisions is quite difficult for large problems.³ Therefore, we need to search for heuristics that can solve the problem effectively.

Since solving the allocation decision is equivalent to solving the Bellman equation (2) in period t , in order to develop efficient heuristics, we focus on the one-step lookahead policy that hinges upon reasonable approximations to Θ^{t+1} . The basic idea is as follows. Suppose $\tilde{\Theta}_{\text{approx}}^{t+1}$ is an easy-to-compute and acceptable approximation to Θ^{t+1} . Given the initial state $(\mathbf{X}^t, \tilde{\mathbf{D}}^t)$ and the realized demand \mathbf{D}^t in period t , we solve the following optimization program

$$\max_{\mathbf{Y}^t} [H(\mathbf{Y}^t | \tilde{\mathbf{D}}^t; \mathbf{D}^t) + \tilde{\Theta}_{\text{approx}}^{t+1}(\mathbf{X}^{t+1}, \tilde{\mathbf{D}}^{t+1})], \quad (12)$$

and obtain the corresponding allocation decision $\mathbf{Y}_{\text{approx}}^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t | \mathbf{D}^t)$ in period t . Let Θ_{approx}^t be the revenue collected by applying the policy $(\mathbf{Y}_{\text{approx}}^t, \dots, \mathbf{Y}_{\text{approx}}^T)$ from period t to T . For simplicity, we do not distinguish between the policy and the decision (e.g., $\mathbf{Y}_{\text{approx}}^t$ and $\mathbf{Y}_{\text{approx}}^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t | \mathbf{D}^t)$), since the proper interpretation is usually clear from the context. Note that $\mathbf{Y}_{\text{approx}}^t$ is a suboptimal policy in the general upgrading problem and $\Theta_{\text{approx}}^t \neq \tilde{\Theta}_{\text{approx}}^t$ in general. Moreover, $\Theta_{\text{approx}}^t(\mathbf{N}^t) \leq \Theta^t(\mathbf{N}^t)$ for any state \mathbf{N}^t in period t since $\Theta^t(\mathbf{N}^t)$ adopts the optimal policy from period t to T .

As pointed out by Bertsekas (2005b), even with readily available revenue-to-go approximations, computing Θ_{approx}^t may still involve substantial computational effort. A number of simplifications of the optimization in (12), including different $\tilde{\Theta}_{\text{approx}}^{t+1}$ functions, have been considered. Here we present two of them that stand out both in terms of computational time and revenue performance. Because of the linearity in the upgrading problem, the first natural candidate is the traditional certainty equivalence control (CEC) heuristic in the literature (see, e.g., Bertsekas 2005a). The CEC is a suboptimal control that treats the uncertain quantities as fixed typical values in the stochastic dynamic program. In our case, we use demand means as typical values in evaluating the function $\tilde{\Theta}_{\text{approx}}^{t+1}$. Thus, under the CEC, expectation calculations are no longer relevant, which can alleviate the computational burden in our problem. Specifically, the optimal allocation policy in period t is solved together with all future periods where the mean demand is used as approximation. That is, the optimal allocation decision $\mathbf{Y}_{\text{CEC}}^t$ in the CEC heuristic will be obtained by solving the following linear program:

$$\begin{aligned} \max_{(\mathbf{Y}_{\text{CEC}}^t, \bar{\mathbf{Y}}^{t+1}, \dots, \bar{\mathbf{Y}}^T) \geq 0} & \left\{ H(\mathbf{Y}_{\text{CEC}}^t | \tilde{\mathbf{D}}^t; \mathbf{D}^t) + \sum_{l=t+1}^T H(\bar{\mathbf{Y}}^l | \tilde{\mathbf{D}}^l; \boldsymbol{\mu}^l) \right\} \\ \text{s.t.} & \tilde{\mathbf{D}}^{t+1} = \tilde{\mathbf{D}}^t + \mathbf{D}^t - (\mathbf{Y}_{\text{CEC}}^t)^T \cdot \mathbf{1}, \\ & \tilde{\mathbf{D}}^{l+1} = \tilde{\mathbf{D}}^l + \boldsymbol{\mu}^l - (\bar{\mathbf{Y}}^l)^T \cdot \mathbf{1}, \quad l = t+1, \dots, T, \\ & \left(\mathbf{Y}_{\text{CEC}}^t + \sum_{l=t+1}^T \bar{\mathbf{Y}}^l \right) \mathbf{1} \leq \mathbf{X}^t, \\ & (\mathbf{Y}_{\text{CEC}}^t)^T \cdot \mathbf{1} \leq \tilde{\mathbf{D}}^t + \mathbf{D}^t, \end{aligned}$$

$$\left(\mathbf{Y}'_{\text{CEC}} + \sum_{l=t+1}^k \bar{\mathbf{Y}}^l \right)^\top \mathbf{1} \leq \tilde{\mathbf{D}}^t + \mathbf{D}^t + \sum_{l=t+1}^k \boldsymbol{\mu}^l, \quad k=t+1, \dots, T, \quad (13)$$

where \mathbf{X}^t , $\tilde{\mathbf{D}}^t$, and \mathbf{D}^t are the capacities, backorders, and realized demand in period t , respectively, and $(\boldsymbol{\mu}^1, \boldsymbol{\mu}^2, \dots, \boldsymbol{\mu}^T)$ denote the mean demand vectors.

The solution to (13) yields the allocation decisions $(\mathbf{Y}'_{\text{CEC}}, \bar{\mathbf{Y}}^{t+1}, \dots, \bar{\mathbf{Y}}^T)$ for periods from t to T , where $(\bar{\mathbf{Y}}^{t+1}, \dots, \bar{\mathbf{Y}}^T)$ are discarded in the subsequent periods. We implement \mathbf{Y}'_{CEC} as the allocation decision for period t and then move on to solve problem (13) in period $t+1$. Let Θ^t_{CEC} be the revenue collected by applying the policy $(\mathbf{Y}'_{\text{CEC}}, \dots, \mathbf{Y}'_{\text{CEC}})$ in periods from t to T . Define $\Pi_{\text{CEC}}(\mathbf{X}) = \Theta^1_{\text{CEC}}(\mathbf{X}, \mathbf{0})$ as the firm's total revenue given initial capacity \mathbf{X} under the CEC heuristic.

Although the above CEC heuristic can simplify our problem, its computational time is still quite long. Consider an N -product general upgrading problem with t periods remaining, the CEC heuristic solves the allocation decisions in the current period as a transportation problem with N classes of capacities and tN classes of demands, whose running time is $O(tN^3(\log(tN) + N \log N))$ (see Brenner 2008). In addition, the optimal allocation is derived from the linear program in (13), which does not use the PSR procedure and the marginal analysis in (8). This means that the CEC might be further improved by exploiting the special properties inherited in our upgrading problem.

To this end, we further simplify the revenue-to-go function by applying greedy upgrading. So the approximation to Θ^{t+1} consists of two components: certainty equivalence control (CEC) and greedy upgrading. Under the CEC, again the mean demand is used as an approximation in all future periods. At the same time, $\bar{\Theta}^{t+1}_{\text{approx}}$ is simplified by adopting greedy upgrading from periods $t+1$ to T rather than solving the linear program as in the CEC heuristic. Such simplification, though suboptimal, is much easier to compute than the linear program.⁴ Given these characteristics of the approximation, we call it refined certainty equivalence control (RCEC) and write $\bar{\Theta}^{t+1}_{\text{approx}}$ as $\bar{\Theta}^{t+1}_{\text{RCEC}}$. In addition to the above approximation, the RCEC heuristic then calculates the protection levels in (8) by replacing Θ^{t+1} with $\bar{\Theta}^{t+1}_{\text{RCEC}}$, and determines the allocation decision $\mathbf{Y}'_{\text{RCEC}}$ in period t by applying the PSR policy to solve the following program

$$\max_{\mathbf{Y}^t} [H(\mathbf{Y}^t | \tilde{\mathbf{D}}^t; \mathbf{D}^t) + \bar{\Theta}^{t+1}_{\text{RCEC}}(\mathbf{X}^{t+1}, \tilde{\mathbf{D}}^{t+1})].$$

Note that $\bar{\Theta}^s_{\text{RCEC}}$ ($s \geq t+1$) can be defined recursively as follows:

$$\bar{\Theta}^s_{\text{RCEC}}(\mathbf{X}^s, \tilde{\mathbf{D}}^s) = H(\mathbf{Y}^s_{\boldsymbol{\mu}} | \tilde{\mathbf{D}}^s; \boldsymbol{\mu}^s) + \bar{\Theta}^{s+1}_{\text{RCEC}}(\mathbf{X}^{s+1}, \tilde{\mathbf{D}}^{s+1}), \quad (14)$$

where $\mathbf{X}^{s+1} = \mathbf{X}^s - \mathbf{Y}^s_{\boldsymbol{\mu}} \cdot \mathbf{1}$, $\tilde{\mathbf{D}}^{s+1} = \tilde{\mathbf{D}}^s + \boldsymbol{\mu}^s - (\mathbf{Y}^s_{\boldsymbol{\mu}})^\top \cdot \mathbf{1}$, $\bar{\Theta}^{T+1}_{\text{RCEC}} \equiv 0$, and $\mathbf{Y}^s_{\boldsymbol{\mu}} = (y^s_{ij}(\boldsymbol{\mu}))_{N \times N}$ is the solution to the following linear program:

$$\max_{\mathbf{Y}^s_{\boldsymbol{\mu}} \geq 0} \left\{ \sum_{1 \leq i \leq j \leq N} \alpha_{ij} y^s_{ij}(\boldsymbol{\mu}) \mid (\mathbf{Y}^s_{\boldsymbol{\mu}})^\top \cdot \mathbf{1} \leq \boldsymbol{\mu}^s + \tilde{\mathbf{D}}^s, \mathbf{Y}^s_{\boldsymbol{\mu}} \cdot \mathbf{1} \leq \mathbf{X}^s \right\}.$$

Given the protection levels derived from $\bar{\Theta}^{t+1}_{\text{RCEC}}$, $\mathbf{Y}'_{\text{RCEC}}$ is the allocation policy in period t solved by the PSR policy, and Θ^t_{RCEC} is the revenue collected by applying policy $(\mathbf{Y}'_{\text{RCEC}}, \dots, \mathbf{Y}'_{\text{RCEC}})$ in period t to T . Define $\Pi_{\text{RCEC}}(\mathbf{X}) = \Theta^1_{\text{RCEC}}(\mathbf{X}, \mathbf{0})$ as the firm's total revenue given initial capacity \mathbf{X} under the RCEC heuristic, and \mathbf{X}_{RCEC} as the optimal capacity that maximizes $\Pi_{\text{RCEC}}(\mathbf{X})$.

We now illustrate how the PSR policy can be used to significantly simplify the computation of the allocation policy $\mathbf{Y}'_{\text{RCEC}}$ (through the protection levels) without the explicit function form of $\bar{\Theta}^{t+1}_{\text{RCEC}}$. Although greedy upgrading (rather the optimal allocation) is used in $\bar{\Theta}^{t+1}_{\text{RCEC}}$, it can be shown that for any state $\mathbf{N}^t = (n^t_1, \dots, n^t_N)$,

$$\frac{\partial}{\partial p} \bar{\Theta}^{t+1}_{\text{RCEC}}(n^t_1, \dots, n^t_{i-1}, p, 0, \dots, 0, -p, n^t_{j+1}, \dots, n^t_N) \quad (15)$$

is decreasing in p .⁵ Thus, the protection levels in the PSR policy can be solved by a binary search, and it suffices to examine whether the protection level p_{ij} is between $\max(n^t_i + n^t_j, 0)$ and n^t_i . The binary search starts with the two initial points $\max(n^t_i + n^t_j, 0)$ and n^t_i . If it proceeds to evaluate more points in the middle (i.e., p_{ij} is strictly between $\max(n^t_i + n^t_j, 0)$ and n^t_i), then there must remain both surplus capacity i and unmet demand j after performing the y^t_{ij} allocation. In this case, there will be no upgrade between classes $(1, \dots, i-1)$ and (j, \dots, N) , and it is unnecessary to compute the protection levels between these two sets. In addition, it indicates that when determining the upgrading between classes r and s for $i \leq s < r \leq j$, $(s, r) \neq (i, j)$, either surplus capacity s or unmet demand r must be zero and the binary search only needs to check the two boundary points. Consequently, the N classes can be partitioned into a number of blocks (say K blocks), and upgrading takes place only within each block. Moreover, in each block there is at most one pair of i and j such that the binary search evaluates more than two points to determine p_{ij} . For block k ($1 \leq k \leq K$) with size n_k ($2 \leq n_k \leq N$), the total number of calls to compute the derivative in (15) is no more than $O(n^2_k + \log |X|)$, where $|X|$ is the upper bound of the initial capacity in each class. Since there is no upgrade between blocks, to solve the allocation decision in each period, the total number of calls to evaluate $\bar{\Theta}^{t+1}_{\text{RCEC}}$ would be bounded by $O(N^2 + N \log |X|)$. It is noteworthy that the advantage of using protection levels is that they can fully characterize the upgrading policy; these protection levels can be computed via binary search that makes a small number of calls to $\bar{\Theta}^{t+1}_{\text{RCEC}}$, so there is no need to undertake the more demanding task of deriving the explicit expression of $\bar{\Theta}^{t+1}_{\text{RCEC}}$.

Consider an N -product general upgrading problem with t periods remaining. Since greedy upgrading can be solved in the running time of $O(tN^2)$, from the above analysis, the RCEC has a running time of $O(tN^3(N + \log |X|))$ in the worst scenario, which is significantly shorter than the CEC when $|X|$ is moderate. More appealingly, the PSR

policy can further reduce the computational complexity in practice. Recall the discussion after Proposition 2, the protection level p_{ij} ($1 \leq i < j \leq N$) in period t only depends on the effective state above i , which is decided by $\mathbf{M}^t - \mathbf{D}^t$. Thus, we can use parallel computing technique and solve all protection levels independently based on $\mathbf{M}^t - \mathbf{D}^t$.

A common feature of the RCEC and CEC heuristics is that both use mean demand in future periods as an approximation. There is a critical difference between these two heuristics. In the RCEC, the PSR procedure is used; in particular, the optimal protection level is determined by (8) (i.e., by comparing the upgrading value to the future marginal value). By contrast, in the CEC, the optimal allocation is derived from the linear program in (13), which utilizes neither the PSR procedure nor (8). It seems the adoption of the PSR procedure in the RCEC plays an important role in both reducing the computational complexity and improving the approximation performance, which will be further discussed in §8.1.

7.2. Benchmark Models

For future comparison, we introduce two benchmark models in this subsection. The first one is called the crystal ball (CB) model. In this model, the firm has perfect demand forecast when allocating the capacities in each period. Such a benchmark has been widely adopted in the literature because it offers the “perfect hindsight” upper bound of the firm’s optimal profit. For instance, it has been used in SZ but is called static model because the firm essentially faces a static capacity allocation problem given complete demand information. Let ω represent a sample path of demand $(\mathbf{D}^1, \dots, \mathbf{D}^T)$ over the sales horizon, and $\mathbf{D}^t(\omega)$ the demand in period t on sample path ω . Then, the firm’s expected profit from period t to T is defined as $\mathbb{E}_\omega[\Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t; \omega)]$, where

$$\begin{aligned} \Theta^t(\mathbf{X}^t, \tilde{\mathbf{D}}^t; \omega) &= \max_{\mathbf{Y}^t, \dots, \mathbf{Y}^T} \sum_{l=t}^T H(\mathbf{Y}^l | \tilde{\mathbf{D}}^l; \mathbf{D}^l(\omega)) \\ \text{s.t. } \tilde{\mathbf{D}}^{l+1} &= \tilde{\mathbf{D}}^l + \mathbf{D}^l(\omega) - (\mathbf{Y}^l)^\top \cdot \mathbf{1}, \quad l = t, \dots, T, \\ \sum_{l=t}^T \mathbf{Y}^l \cdot \mathbf{1} &\leq \mathbf{X}^t, \\ \sum_{l=t}^k (\mathbf{Y}^l)^\top \cdot \mathbf{1} &\leq \tilde{\mathbf{D}}^t + \sum_{l=t}^k \mathbf{D}^l(\omega), \quad k = t, \dots, T, \\ \mathbf{Y}^l &\geq 0, \quad l = t, \dots, T. \end{aligned}$$

The firm’s optimal profit in the crystal ball model is given by

$$\max_{\mathbf{X}^1 \in \mathfrak{N}_+^N} \Pi_{\text{CB}}(\mathbf{X}^1) = \max_{\mathbf{X}^1 \in \mathfrak{N}_+^N} \left\{ \mathbb{E}_\omega[\Theta^1(\mathbf{X}^1, \mathbf{0}; \omega)] - \mathbf{C} \cdot \mathbf{X}^1 \right\}, \quad (16)$$

which can be used to benchmark the performance of our heuristic in the dynamic upgrading problem.

The second benchmark is the model without product upgrading. In this case, the firm’s problem reduces to N independent newsvendors (NV) with backorders. The firm’s expected profit can be written as

$$\begin{aligned} &\max_{\mathbf{X}^1 \in \mathfrak{N}_+^N} \Pi_{\text{NV}}(\mathbf{X}^1) \\ &= \max_{\mathbf{X}^1 \in \mathfrak{N}_+^N} \left\{ \mathbb{E}_{\{\mathbf{D}^1, \dots, \mathbf{D}^T\}} \sum_{s=1}^N \sum_{t=1}^T [\alpha_{ss} \min(x_s^t, d_s^t) - g_s(\tilde{d}_s^t + d_s^t)] - \mathbf{C} \cdot \mathbf{X}^1 \right\} \quad (17) \\ \text{s.t. } x_s^{t+1} &= (x_s^t - d_s^t)^+, \quad \tilde{d}_s^{t+1} = \tilde{d}_s^t + (d_s^t - x_s^t)^+, \\ x_s^1 &= (\mathbf{X}^1)_s, \quad d_s^t = (\mathbf{D}^t)_s, \\ & \quad s = 1, \dots, N, \quad t = 1, \dots, T. \end{aligned}$$

Note that although the two benchmark models (CB and NV) are similar to the static and independent newsvendor models used in SZ, because of the backlogging assumption, the firm has to allocate capacity in each period in our model, rather than accumulate the demand for the entire selling season and then allocate the capacity as in SZ. In §8, we use Monte Carlo simulation to generate demand sample paths and use the sample average approximation method to solve the two benchmark models.

8. Numerical Studies

In this section, we conduct numerical studies to derive insights into the capacity management problem. First, we test the performance of the RCEC heuristic proposed in the previous section. After that, by using the heuristic and benchmark models, we investigate the importance of the allocation mechanism and the capacity sizing decision. For simplicity, we focus on integral demands.

8.1. Performance of RCEC

Because of the complexity of the problem, we use extensive numerical experiments to test the performance of the heuristics. These experiments are conducted using MATLAB R2013a on an Intel Core i7-2600 desktop with 12 Gb RAM. We focus on the RCEC heuristic because it will be used later for further numerical investigation.

The first set of experiments has $N = 4$ and $T = 3$. For such a problem size, we can use backward induction to evaluate the firm’s optimal profit $\Pi(\mathbf{X})$ given in (1). Later we will also discuss the performance of the RCEC for larger problem sizes where it is difficult to evaluate $\Pi(\mathbf{X})$ directly. Given an initial capacity $\mathbf{X} \in \mathfrak{N}_+^N$, define the performance measure

$$\Delta_{\text{opt}} = \left| \frac{\Pi_{\text{RCEC}}(\mathbf{X}) - \Pi(\mathbf{X})}{\Pi(\mathbf{X})} \right| * 100\%, \quad (18)$$

i.e., the percentage of profit loss by using $\Pi_{\text{RCEC}}(\mathbf{X})$ rather than $\Pi(\mathbf{X})$.

Table 1. The percentage profit loss (Δ_{opt}) of RCEC relative to the optimal solution.

Mean	Std.	Median	90%-percentile	95%-percentile	99%-percentile	Max.
0.40	1.13	0.15	0.77	1.27	5.58	17.99

To calculate $\Pi(\mathbf{X}) = \Theta^1(\mathbf{X}, \mathbf{0})$, we use the Monte Carlo method and consider a comprehensive range of scenarios, which capture different fluctuation patterns of demand means along the selling horizon (i.e., variation of $\mathbb{E}[\mathbf{D}^t]$ from $t = 1$ to T), different correlations between classes of demands in each period (i.e., $\text{Corr}(d_i^t, d_j^t)$ for all $1 \leq i \leq j \leq N$), different demand distributions (i.e., Normal distribution with integer rounding and Poisson distribution), and various economic parameters (i.e., revenue (r_1, \dots, r_N) , goodwill cost (g_1, \dots, g_N) , usage cost (u_1, \dots, u_N) , and capacity cost (c_1, \dots, c_N)). Furthermore, to ensure the robustness of the results, we also test a number of different initial capacities \mathbf{X} used in (18), which consist of both realistic and extreme scenarios. In total there are 4,212 experiments in this numerical study. A full description of the setup of the numerical study is lengthy and thus given in the appendix.

The statistics for the Δ_{opt} value are reported in Table 1. It can be seen that the RCEC performs very well in this numerical study. Among all the experiments tested, the 90th percentile of the profit loss is 0.77%, and the average is 0.40%.

Next we test the performance of the RCEC in larger problems. Specifically, we consider problems with $N = 5$ products and up to $T = 60$ periods. Given such sizes, it is extremely time-consuming to evaluate the optimal revenue function $\Pi(\mathbf{X})$. Instead, we use $\Pi_{\text{CB}}(\mathbf{X})$ from the crystal ball (CB) model defined in (16) as the benchmark for comparison. Recall that $\Pi_{\text{CB}}(\mathbf{X})$ is an upper bound of the optimal revenue $\Pi(\mathbf{X})$ for any \mathbf{X} , and the following relationship holds: $\Pi_{\text{CB}}(\mathbf{X}) \geq \Pi(\mathbf{X}) \geq \Pi_{\text{RCEC}}(\mathbf{X})$. Define

$$\Delta_{\text{CB}} = \left| \frac{\Pi_{\text{RCEC}}(\mathbf{X}) - \Pi_{\text{CB}}(\mathbf{X})}{\Pi_{\text{CB}}(\mathbf{X})} \right| * 100\%.$$

Then Δ_{CB} is an upper bound of Δ_{opt} , the percentage profit loss of the RCEC (i.e., $\Pi_{\text{RCEC}}(\mathbf{X})$) relative to the optimal revenue (i.e., $\Pi(\mathbf{X})$).

Similar experiment design has been used as Table 1 except that now we consider 5 products with several different T values. This allows us to examine up to 4 levels of upgrading. Also by varying T we can study the impact of the number of periods (or the frequency of upgrading decisions) on the problem. Specifically, T takes values from a set $\{3, 15, 30, 60\}$. For each T , there are 13,260 experiments in total in this numerical study. To save space, we provide a detailed description in the appendix.

We summarize the statistics of Δ_{CB} for different T 's in Table 2. The table shows that the value of Δ_{CB} is increasing in the number of periods, T . The RCEC ignores the

Table 2. The percentage profit loss (Δ_{CB}) of RCEC relative to the CB solution.

T	Mean	Std.	Median	90%-percentile	Max.
3	0.14	0.38	0.00	0.34	6.74
15	1.52	2.51	0.23	4.83	12.06
30	2.37	3.36	0.42	5.37	23.37
60	3.52	3.82	0.59	8.21	39.36

randomness of the demand in future periods (recall that the mean demand is used). Thus, compared to $\Pi_{\text{CB}}(\mathbf{X})$, more demand information is lost as T increases. Table 2 also indicates that the value of Δ_{CB} is small in general: Even for $T = 60$, Δ_{CB} is 8.21% at the 90th percentile, and the average is about 3.52%. This observation has two implications. First, since Δ_{CB} is the upper bound of Δ_{opt} , we know that Δ_{opt} is generally small in the tested examples. This means that for the 5-product numerical experiments, the RCEC also performs well. Second, the observation implies that the difference between $\Pi_{\text{CB}}(\mathbf{X})$ and $\Pi(\mathbf{X})$ is small. In other words, the value of advance demand information is not significant for many cases. Such a result is in line with some of the findings reported in the literature. For instance, SZ finds from numerical study that when the optimal upgrading policy is used, the firm's expected revenue is consistently within 1% of the revenue in a static model (i.e., the crystal ball model). Similarly, Acimovic and Graves (2015) find in a dynamic order fulfillment setting that the crystal ball model improves the performance of the proposed heuristic by 2%, i.e., the performance difference between the crystal ball model and the true optimum is smaller than 2%.

We now compare the performances of the RCEC and the CEC. Define the ratio

$$\eta = \frac{\Pi_{\text{RCEC}}(\mathbf{X})}{\Pi_{\text{CEC}}(\mathbf{X})}$$

to measure the relative performances of the two heuristics. So a ratio higher (lower) than 1 implies that the RCEC outperforms (underperforms) the CEC. We calculate the ratio for the problem instances used in the numerical study underlying Table 2 (i.e., $N = 5$ and $T = \{3, 15, 30, 60\}$). The statistics of the ratio values are summarized in Table 3 since the results are consistent across different T 's. Meanwhile, as we mentioned earlier, we also compare the actual computation times of the CEC and the RCEC heuristics in these instances, i.e., the total computation times for deriving the capacity allocations and obtaining the values of Π_{CEC} and Π_{RCEC} for any given \mathbf{X} . Specifically, we use MOSEK toolbox for MATLAB version 7 to solve the linear program in (13) in the CEC heuristic, and we apply the binary search to solve the protection levels in (8) while replacing Θ^{t+1} by $\bar{\Theta}_{\text{RCEC}}^{t+1}$ in (14). Similarly, we define

$$\eta_{\text{time}} = \frac{\text{Time for solving } \Pi_{\text{RCEC}}(\mathbf{X})}{\text{Time for solving } \Pi_{\text{CEC}}(\mathbf{X})},$$

whose statistics are also reported in Table 3.

Table 3. Comparison of RCEC and CEC.

	Mean	Std.	Min.	25%-percentile	Median	Max.
η	1.00	0.04	0.02	1.00	1.00	4.81
η_{time}	0.09	0.05	0.00	0.05	0.08	1.13

We observe that the CEC may outperform the RCEC in some instances (e.g., the ratio can be as low as 2%);⁶ however, for the majority of the examples, the RCEC performs better than the CEC. More importantly, the reduction of computation time from CEC to RCEC is substantial: all else being equal, the average time for solving a test instance using the RCEC is only 9% of that using the CEC.

Why does the RCEC exhibit a better overall performance? We offer the following plausible explanation. In both the CEC and RCEC heuristics, we replace the future random demands by their means in each period. Such an approximation clearly will change our original problem and result in suboptimal solutions. In the RCEC, the optimal protection level is determined by comparing two values: The first is the upgrading value from using the product in the current period; the second is the expected marginal value of the product if it is saved to the next period. For illustration, consider the upgrading of demand j using capacity i in period t . The latter value is defined as $\Theta_{\text{RCEC}}^{t+1}(\mathbf{X}^{t+1} + \mathbf{e}_i, \mathbf{D}^{t+1} + \mathbf{e}_j | \mu^{t+1}, \dots, \mu^T) - \Theta_{\text{RCEC}}^{t+1}(\mathbf{X}^{t+1}, \mathbf{D}^{t+1} | \mu^{t+1}, \dots, \mu^T)$, where \mathbf{e}_s ($s = i, j$) is the unit vector with 1 in position s . The mean demand approximation may introduce biases into the two revenue functions. However, since the expected marginal revenue is defined as the difference between the two revenue functions, these biases may be canceled out to some degree. In other words, the inaccuracies introduced by certainty equivalence control might be reduced in the RCEC heuristic. Note that such a cancellation effect does not exist in the traditional CEC heuristic. Therefore, the RCEC generally outperforms the CEC. In addition, the RCEC is more attractive than the CEC in terms of computational time in our numerical study.

One may also use the deflected linear decision rule (DLDR) method proposed in Chen et al. (2008) to approximate Θ^t in the PSR algorithm. Let Θ_{DLDR}^t be the revenue collected by using $\mathbf{Y}_{\text{DLDR}}^t$'s in the remaining sales horizon, and denote $\Pi_{\text{DLDR}}(\mathbf{X}) = \Theta_{\text{DLDR}}^1(\mathbf{X}, \mathbf{0})$ as the expected revenue under the DLDR heuristic. We evaluate $\Pi_{\text{DLDR}}(\mathbf{X})$ in the numerical study described above and find that $\Pi_{\text{DLDR}}(\mathbf{X})$ and $\Pi_{\text{RCEC}}(\mathbf{X})$ are almost identical in all the problem instances.

Based on the results in Tables 1 and 2, we conclude that the RCEC performs very well in a wide range of problem situations. In addition, the RCEC greatly reduces the computational complexity of the original problem. Therefore, in the rest of the paper, we will use the RCEC to solve the dynamic capacity management problem. Although the RCEC only provides an approximation of the optimal policy, we expect the insights to carry over to the optimal

policy because the performance of the RCEC is close to optimal in most of our numerical studies.

8.2. Value of Optimal Upgrading

Given the efficiency and effectiveness of the RCEC heuristic, we are ready to derive more insights into the problem using numerical studies. There are a couple of natural questions we would like to address. First, what is the value of using multistep upgrading? Second, what is the value of using the optimal capacity? Both questions are important from a practical standpoint because managers need to know how complex an upgrading structure should be used and how to determine the initial capacity. This subsection focuses on the first question, and the second will be addressed in the next subsection.

Let $\Pi_{\text{RCEC}}^k(\mathbf{X})$ be the revenue function given initial capacity \mathbf{X} and k -level upgrading (i.e., product i can be used to satisfy class j demand only if $i \leq j \leq i + k$). Note that when $k = 0$, no upgrading is allowed, and $\Pi_{\text{RCEC}}^0(\mathbf{X}) = \Pi_{\text{NV}}(\mathbf{X})$, where $\Pi_{\text{NV}}(\mathbf{X})$ is the optimal revenue in the newsvendor model in (17). Define

$$\Delta_{\text{RCEC}}^k = \frac{\Pi_{\text{RCEC}}^k(\mathbf{X}) - \Pi_{\text{RCEC}}^{k-1}(\mathbf{X})}{\Pi_{\text{RCEC}}^{k-1}(\mathbf{X})} * 100\%, \quad k = 1, 2, 3, 4,$$

which measures the percentage profit gain from one additional level of upgrading under the RCEC.⁷

We evaluate the values of Δ_{RCEC}^k using the same parameters as those for Table 2 except the initial capacities. Intuitively, upgrade is more valuable when the capacity is unbalanced, i.e., there is excess capacity for some products while there is shortage for the others. Such unbalance may occur even if the initial capacities are optimally set, because demand may fluctuate because of seasonality and trend while capacities are determined for the long term. Thus, when choosing the initial capacity we use the following procedure. Start with the optimal capacity under the RCEC, denoted as \mathbf{X}_{RCEC} (the computation of \mathbf{X}_{RCEC} will be discussed in the next subsection); then set the capacity for one product (say, product j) to 0 while adding capacity $(\mathbf{X}_{\text{RCEC}})_j$ to a higher-quality product; finally, scale the entire capacity vector by different multipliers. Mathematically, for $1 \leq i < j \leq 5$, we consider all initial capacity \mathbf{X} , whose components are given by

$$\begin{aligned} (\mathbf{X})_i &= \lambda((\mathbf{X}_{\text{RCEC}})_i + (\mathbf{X}_{\text{RCEC}})_j), \quad (\mathbf{X})_j = 0, \\ (\mathbf{X})_s &= \lambda(\mathbf{X}_{\text{RCEC}})_s, \quad \forall s \in \{1, 2, 3, 4, 5\} \setminus \{i, j\}, \end{aligned}$$

where $\lambda \in \{0.9, 1, 1.1\}$. There are 10 combinations of the initial capacities for each λ and parameter set; one example is $\mathbf{X} = ((\mathbf{X}_{\text{RCEC}})_1 + (\mathbf{X}_{\text{RCEC}})_2, 0, (\mathbf{X}_{\text{RCEC}})_3, (\mathbf{X}_{\text{RCEC}})_4, (\mathbf{X}_{\text{RCEC}})_5)$. A full list of the initial capacities are given in the appendix. We believe such a design captures the possible capacity scenarios that may happen over time as the firm allocates products to satisfy realized demand, especially those with unbalanced capacities. Moreover, the

Table 4. The value of using multistep upgrading (Δ_{RCEC}^k).

T	Upgrading level k	Mean	Median	90%-percentile
3	1	29.75	20.64	51.63
	2	5.71	2.09	15.21
	3	1.45	0.11	4.99
	4	0.25	0.01	0.28
15	1	25.96	20.25	47.44
	2	4.86	2.99	12.88
	3	0.79	0.04	2.70
	4	0.07	0	0.09
30	1	20.38	19.88	45.63
	2	3.89	1.63	11.70
	3	0.67	0.02	2.02
	4	0.05	0	0.07
60	1	16.77	12.57	32.70
	2	2.88	1.07	10.23
	3	0.40	0	2.1
	4	0.03	0	0.04

mean of total demand over the selling horizon remains the same for different $T \in \{3, 15, 30, 60\}$, which implies that less demand information is available within each period for larger T . The numerical results for different T values are given in Table 4.

There are several observations from Table 4. First, Table 4 suggests that the firm’s profit increases in the number of upgrading levels k and the marginal value decreases in k . In particular, we can see that most of the gains are achieved with a single level of upgrading, though more can be achieved with multistep upgrading. In some cases, the value from two or more levels of upgrading can be quite high. For instance, with $T = 3$, the benefit of moving from one-step upgrading to two-step upgrading can be as high as 15.21% at the 90th percentile (i.e., for at least 10% of the scenarios, the value is more than 15.21%). This result implies that single-step upgrading may not capture the full benefit of upgrading, and multistep upgrading is needed in certain cases.

Second, Table 4 indicates that the value of multistep upgrading decreases in T . That is, using more upgrading levels will be less beneficial for the firm when there are more time periods in the selling horizon. Close scrutiny reveals that there is a key contributing factor to this interesting observation. A large T value means there are more time periods, which allows “chain allocation” to be more likely to happen. To see this, first consider $T = 1$. In this case, under single-step upgrading, product 1 cannot be used to satisfy demand 3. However, with $T = 2$, it is possible that product 2 is used to satisfy demand 3 in period 1; then, in the second period, product 1 is used to satisfy demand 2. These two allocations essentially mean that product 1 is used to satisfy demand 3. Evidence of such a chain allocation has been observed in our numerical study. For instance, we examine the upgrading patterns for scenarios with $T = 3$ and $T = 15$ under the full upgrading structure. With $T = 3$, the average numbers of upgrades

Table 5. The value of using multistep upgrading (Δ_{RCEC}^k) under optimal initial capacity.

Upgrading level k	Mean	Median	90%-percentile
1	1.34	0.87	3.61
2	0.40	0.21	1.14
3	0.14	0.05	0.42
4	0.05	0.01	0.17

from product 1 to demand 2 (i.e., $\sum_{t=1}^T y_{12}^t$), product 2 to demand 3 (i.e., $\sum_{t=1}^T y_{23}^t$), and product 1 to demand 3 (i.e., $\sum_{t=1}^T y_{13}^t$) are 10.6, 13.3, and 5.1, respectively, whereas the corresponding numbers with $T = 15$ are 12.9, 15.0, and 4.0. Note that there are more upgrades from product 1 to demand 3 with $T = 3$ than with $T = 15$ even though there are fewer chances for upgrading with $T = 3$. The chain allocation is analogous to multistep upgrading; the only difference is that it can be better executed when there are more time periods. Therefore, multistep upgrading is less valuable since it can be implemented even under single-step upgrading, but in a different way.

Finally, the numerical experiments suggest that multistep upgrading is most valuable when the initial capacity is unbalanced. For example, for $T = 3$, when the optimal initial capacity \mathbf{X}_{RCEC} is used, the incremental value of moving from 2-level to 3-level upgrading is 0.04% on average; however, for initial capacity $\mathbf{X} = ((\mathbf{X}_{\text{RCEC}})_1, (\mathbf{X}_{\text{RCEC}})_2 + (\mathbf{X}_{\text{RCEC}})_5, (\mathbf{X}_{\text{RCEC}})_3, (\mathbf{X}_{\text{RCEC}})_4, 0)$, the counterpart value is 5.10%. This indicates that the multistep upgrading is quite important because unbalanced capacity may arise over time, even if the problem starts with the optimal initial capacity.

What is the benefit of using more upgrading levels if the optimal initial capacities are used? To answer this question, let $\mathbf{X}_{\text{RCEC}}(k)$ ($k = 0, 1, \dots, 4$) be the optimal initial capacities obtained from the RCEC heuristic with k -level upgrading, and redefine

$$\Delta_{\text{RCEC}}^k = \frac{\Pi_{\text{RCEC}}^k(\mathbf{X}_{\text{RCEC}}(k)) - \Pi_{\text{RCEC}}^{k-1}(\mathbf{X}_{\text{RCEC}}(k-1))}{\Pi_{\text{RCEC}}^{k-1}(\mathbf{X}_{\text{RCEC}}(k-1))} * 100\%,$$

$k = 1, 2, 3, 4$

which is the percentage profit gain from one additional level of upgrading under the RCEC if the corresponding optimal initial capacities are used. Using the same set of parameters as in Table 4, we obtain the numerical results given in Table 5.

As one may expect, the values of using multistep upgrading are much smaller in Table 5 because the initial capacities have been accordingly adjusted, and this lowers the benefit of using more levels of upgrading. However, the value of multistep upgrading should not be overlooked either: the additional profit by moving from one-step to two-step upgrading captures about 30% of that from no upgrading to one-step upgrading.⁸ Our numerical analysis also indicates that the more variable the demand, the higher

the value of upgrading. We omit the details because similar observations have already been made in Bassok et al. (1999) and SZ.

8.3. Capacity Sizing Decision vs. Allocation Mechanism

The profit of the upgrading problem hinges upon both the initial capacity and the allocation mechanism. This raises an interesting question: Which decision is more important, capacity sizing or allocation mechanism? This is a practical question because the firm may wish to focus limited resources on improving the decision that has a bigger impact on profit. To shed some light on this question, we measure the importance of each decision using the profit loss when a suboptimal decision is applied rather than the optimal one. Next, we describe the suboptimal decisions that will be used.

In our problem, it is time-consuming to derive the optimal initial capacity even if we can efficiently solve the optimal allocation decision by the RCEC heuristic. So we consider two simple alternatives. The first alternative is to use the optimal capacity \mathbf{X}_{CB} in the crystal ball model. The crystal ball model is called static model in SZ, who find that \mathbf{X}_{CB} yields nearly optimal revenue for the firm in their single-step upgrading model. To check whether the result carries over to our general upgrading model, define

$$\Delta_{\mathbf{X}_{CB}} = \left| \frac{\Pi_{RCEC}(\mathbf{X}_{CB}) - \Pi_{RCEC}(\mathbf{X}_{RCEC})}{\Pi_{RCEC}(\mathbf{X}_{RCEC})} \right| * 100\%$$

to measure the performance of the crystal ball capacity \mathbf{X}_{CB} . With the same parameters used for Tables 2–4, we evaluate $\Delta_{\mathbf{X}_{CB}}$ for 624 examples and summarize the results in Table 6 (the first row). It can be seen that $\Delta_{\mathbf{X}_{CB}}$ is generally negligible in the numerical study: The average revenue difference is 0.02% and the maximum is 1.10%.

Note that since the true optimal capacity is unknown, we use \mathbf{X}_{RCEC} , the optimal initial capacity under RCEC, as the benchmark for the comparison. It is worth mentioning that despite its excellent performance, the firm's total revenue function $\Pi_{RCEC}(\mathbf{X})$ is not necessarily concave in \mathbf{X} under the RCEC heuristic. A nonconcave example is provided in the appendix. Thus if needed, we resort to the following approach to determine \mathbf{X}_{RCEC} . We start with the optimal capacity from the crystal ball model and apply the grid search to find the optimal initial capacity: For each search step, we move from $\mathbf{X} = (x_1, x_2, \dots, x_N)^T$ to the point in set $\{\mathbf{X}' = (x'_1, x'_2, \dots, x'_N)^T : \max_{i \in \{1, 2, \dots, N\}} |x'_i - x_i| = 1\}$, which gives the largest function value of Π_{RCEC} . That is, we conduct a local search along the grid; such an approach has also been used in SZ.

An even simpler alternative is to use the newsvendor capacity \mathbf{X}_{NV} , i.e., the optimal capacity under no upgrading. Similarly, in the same numerical study, we define

$$\Delta_{\mathbf{X}_{NV}} = \left| \frac{\Pi_{RCEC}(\mathbf{X}_{NV}) - \Pi_{RCEC}(\mathbf{X}_{RCEC})}{\Pi_{RCEC}(\mathbf{X}_{RCEC})} \right| * 100\%$$

Table 6. Capacity decision vs. allocation mechanism.

	Mean	Std.	Median	90%-percentile	Max.
$\Delta_{\mathbf{X}_{CB}}$	0.02	$5.32 * 10^{-2}$	0	0.04	1.10
$\Delta_{\mathbf{X}_{NV}}$	0.30	$2.89 * 10^{-1}$	0.26	0.68	1.62
Δ_{greedy}	5.60	5.68	8.03	12.81	13.75

and present the statistics of $\Delta_{\mathbf{X}_{NV}}$ in Table 6 (the second row). We can see that $\Delta_{\mathbf{X}_{NV}}$ is greater than $\Delta_{\mathbf{X}_{CB}}$ in general, but it offers reasonably good performance as well. The average and maximum revenue differences are 0.30% and 1.62%, respectively. In particular, the number at the 90th percentile is 0.68%, which means that the newsvendor capacity performs quite well for the majority of the scenarios. From the above observations, one can see that these simple alternatives to the optimal capacity perform reasonably well. Therefore, as long as the optimal upgrading policy is used, the value of using the optimal capacity seems to be very small in our problem setting.

Next, we consider the impact of using a suboptimal allocation policy. We first use greedy upgrading as the suboptimal policy. Under such a policy, unmet demands are upgraded in the order characterized by the PSR policy; however, there is no rationing and surplus capacities are used for upgrading as much as possible. It serves as a reasonable suboptimal policy because it is intuitive and straightforward to implement in practice. Furthermore, the RCEC heuristic incorporates greedy upgrading to simplify its computation. Specifically, let $\Pi_{\text{greedy}}(\mathbf{X})$ be the expected profit using greedy upgrading given initial capacity \mathbf{X} . We define

$$\Delta_{\text{greedy}} = \left| \frac{\Pi_{RCEC}(\mathbf{X}_{RCEC}) - \Pi_{\text{greedy}}(\mathbf{X}_{RCEC})}{\Pi_{RCEC}(\mathbf{X}_{RCEC})} \right| * 100\%$$

as the profit loss because of greedy upgrading. The same parameters for $\Delta_{\mathbf{X}_{CB}}$ and $\Delta_{\mathbf{X}_{NV}}$ have been used, and the statistics of Δ_{greedy} are presented in Table 6 (the third row). The average profit loss because of greedy upgrading is 5.598%, which is much larger than those for $\Delta_{\mathbf{X}_{CB}}$ and $\Delta_{\mathbf{X}_{NV}}$. In addition to greedy upgrading, we also test suboptimal allocation policies that involve only k -step ($k = 0, \dots, N - 2$) upgrading. The magnitudes of profit losses are still generally much larger than those for $\Delta_{\mathbf{X}_{CB}}$ and $\Delta_{\mathbf{X}_{NV}}$. To save space, the detailed results are presented in the appendix.

The above numerical results based on the RCEC heuristic suggest that the benefit of choosing an effective allocation mechanism outweighs that of choosing an accurate initial capacity. Based on these observations, in practice, the firm may decide the initial capacity by using simple approximations (e.g., either the NV or CB model) and focus on optimally allocating the capacity during the sales horizon.

9. Conclusion

This paper studies a firm's capacity investment and allocation problem in a dynamic setting with stochastic demand. There are multiple demand classes, which can be satisfied by multiple classes of capacities. Demand arrives in discrete time periods, and the firm needs to make capacity allocation decisions in each period before observing future demand. A general upgrading structure is considered, which is broad enough to cover a wide range of practical upgrading situations. One may also view this as an inventory management problem with one-way dynamic substitution.

We first show that for any given initial capacity, a Parallel and Sequential Rationing (PSR) policy is optimal for the firm. Under the PSR policy, the firm can make upgrading decisions in each period sequentially rather than jointly, which greatly reduces the complexity of the capacity allocation problem. Despite the well-structured PSR policy, the dynamic allocation problem is still subject to the curse of dimensionality. Thus we propose a Refined Certainty Equivalence Control (RCEC) heuristic that improves over the traditional CEC methodology by exploiting the property of the PSR policy. Through extensive numerical experiments, we find that the RCEC heuristic is highly efficient and yields nearly optimal revenue for the firm. With the help of the RCEC heuristic, we conduct numerical studies to derive managerial insights about the dynamic capacity management problem. Our numerical studies indicate that the multistep upgrading could be significantly valuable, especially when the capacities are not balanced (either because of suboptimal initial investment or unexpected demand realizations over time). We find that using simple approximations (e.g., the newsvendor model without substitution and the crystal ball model with perfect demand forecast) for the initial capacities leads to negligible profit loss, whereas the negative impact of using a suboptimal allocation (e.g., greedy upgrading) could be quite significant. In this sense, the allocation mechanism plays a more important role in our problem than the capacity sizing decision.

There are several interesting directions for future research. First, it is worthwhile exploring models with general nonstationary model parameters. The PSR policy remains optimal as long as the profit margin is monotonically decreasing over time. However, with general nonstationary model parameters, the optimal policy is still unknown. It is possible that the parallel allocation is not used to the maximum extent, and the upgrading decision may depend on the demands and capacities below the class being upgraded if a higher profit margin is expected in some later period. Second, it is a challenge to analyze models with lost sales. The backorder assumption used in this paper is critical for the optimal PSR allocation policy. It is not clear how the optimal policy looks under the lost-sales assumption. Third, it would be interesting to take pricing decisions into account, i.e., the firm may adjust prices over

time depending on the evolution of demand and remaining capacity levels. Finally, recently there has been a growing interest in studying opportunistic consumer behavior in operations problems. In our upgrading setting, a consumer may intentionally choose the product that is sold out, hoping to receive a free upgrade later. It would be interesting to investigate how such a behavior may affect the firms' operational strategies (e.g., upgrading policies).

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/opre.2015.1446>.

Acknowledgments

The authors thank the department editor Chung-Piaw Teo, the anonymous associate editor and reviewers for insightful comments and suggestions that have significantly improved the paper. Specifically, §6 was suggested by the review team and §8.3 was suggested by Chung-Piaw Teo. This research of the second author is partly supported by National Science Foundation [Grants CMMI-1030923, CMMI-1363261 and CMMI-1538451] and National Natural Science Foundation of China [Grants 71228203 and 71520107001].

Endnotes

1. The backorder assumption is used mainly for tractability. Notice that an unmet demand could be upgraded in any subsequent period, so it is reasonable to assume that the customers are willing to wait for potential upgrades, i.e., unsatisfied demands can be backlogged.
2. This counter-intuitive example remains valid for any goodwill cost g_2 if the length T satisfies $\alpha_{11} - (T - 2)g_2 > \alpha_{12} > \alpha_{11} - (T - 1)g_2$ and $\mathbf{D}^2 = \dots = \mathbf{D}^{T-1} = (0, 0)$ and $\mathbf{D}^T = (1, 0)$.
3. To deal with the dimensionality issue, SZ propose a series of bounds to approximate the optimal protection levels. For instance, when computing the protection level for product i , one may consider only the capacity for $i - 1$, while assuming the products above $i - 1$ to be either ∞ (this gives a lower bound of the protection level) or 0 (this gives an upper bound). It has been found that under the single-step upgrading assumption, these bounds are very tight and yield nearly optimal revenue for the firm. However, such bounds do not work well in our model, where general upgrading is allowed.
4. We have tested the heuristic without the greedy upgrading and found that the performance is almost identical. That is, the use of greedy upgrading in this heuristic can significantly reduce the computational complexity but has a negligible impact on the revenue performance.
5. Since future demands are known, there exists a period s ($t + 1 \leq s \leq T$) in which capacity i will be depleted. From the expression in (15), a marginal change of p only affects the greedy upgrading in period s because both capacity i and backorder demand j change simultaneously in p . In particular, capacity i is used to sequentially satisfy demands

from class i to j in period s . As p increases, the additional units of capacity i will be used to satisfy demands from lower classes that have smaller profit margins. Thus, the partial derivative is a decreasing step function of p .

6. The differences between RCEC and CEC seem to be quite large in extreme cases. The reason is that the optimal expected profits in these cases are very close to zero and thus a small perturbation in Π_{RCEC} and Π_{CEC} may drastically change the ratio η .

7. Note that Δ_{RCEC}^k may be negative since Π_{RCEC}^k may be different from the optimal value. However, we only observed 15 cases out of 74,880 (i.e., 0.02%) in our numerical study.

8. In our numerical study, upgrade constitutes 1.92% of the total satisfied demands on average when the optimal initial capacity is used, and 26.59% when the suboptimal initial capacities are adopted. If the firm uses frequent upgrading to satisfy customer demand (e.g., the initial capacity is poorly decided), customers may learn about the upgrading pattern and become opportunistic. That is, a class i customer may intentionally ask for product j ($i < j$), hoping that she will be upgraded when product j is out of stock. Incorporating such a behavior is out of the scope of this paper and therefore left for future research.

References

- Acimovic J, Graves SC (2015) Making better fulfillment decisions on the fly in an online retail environment. *M&SOM* 17(1):34–51.
- Archibald TW, Sassen SAE, Thomas LC (1997) An optimal policy for a two depot inventory problem with stock transfer. *Management Sci.* 43(2):173–183.
- Axsäter S (2003) A new decision rule for lateral transshipments in inventory systems. *Management Sci.* 49(9):1168–1179.
- Bassok Y, Anupindi R, Akella R (1999) Single-period multiproduct inventory models with substitution. *Oper. Res.* 47(4):632–642.
- Bertsekas DP (2005a) *Dynamic Programming and Optimal Control*, Vol. 1 (Athena Scientific, Belmont, MA).
- Bertsekas DP (2005b) Dynamic programming and suboptimal control: A survey from adp to mpc. *Eur. J. Control* 11(4–5):310–334.
- Bish EK, Wang Q (2004) Optimal investment strategies for flexible resources, considering pricing and correlated demands. *Oper. Res.* 52(6):954–964.
- Brenner U (2008) A faster polynomial algorithm for the unbalanced hitchcock transportation problem. *Oper. Res. Lett.* 36(4):408–413.
- Chen X, Sim M, Sun P, Zhang J (2008) A linear decision-based approximation approach to stochastic programming. *Oper. Res.* 56(2):344–357.
- Chod J, Rudi N (2005) Resource flexibility with responsive pricing. *Oper. Res.* 53(3):532–548.
- Chou MC, Chua GA, Teo C-P, Zheng H (2010) Design for process flexibility: Efficiency of the long chain and sparse structure. *Oper. Res.* 58(1):43–58.
- Chou MC, Chua GA, Teo C-P, Zheng H (2011) Process flexibility revisited: The graph expander and its applications. *Oper. Res.* 59(5):1090–1105.
- de Véricourt F, Karaesmen F, Dallery Y (2001) Assessing the benefits of different stock-allocation policies for a make-to-stock production system. *M&SOM* 3(2):105–121.
- de Véricourt F, Karaesmen F, Dallery Y (2002) Optimal stock allocation for a capacitated supply system. *Management Sci.* 48(11):1486–1501.
- Deshpande V, Cohen MA, Donohue K (2003) A threshold inventory rationing policy for service-differentiated demand classes. *Management Sci.* 49(6):683–703.
- Ding Q, Kouvelis P, Milner JM (2006) Dynamic pricing through discounts for optimizing multiple-class demand fulfillment. *Oper. Res.* 54(1):169–183.
- Fine CH, Freund RM (1990) Optimal investment in product-flexible manufacturing capacity. *Management Sci.* 36(4):449–466.
- Gallego G, Stefanescu C (2009) Upgrades, upsells and pricing in revenue management. Accessed November 18, 2015, <http://ssrn.com/abstract=1334341>.
- Gallego G, van Ryzin GJ (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* 45(1):24–41.
- Ha AY (1997) Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Res. Logist.* 44(5):457–472.
- Ha AY (2000) Stock rationing in an $M/E_k/1$ make-to-stock queue. *Management Sci.* 46(1):77–87.
- Hsu A, Bassok Y (1999) Random yield and random demand in a production system with downward substitution. *Oper. Res.* 47(2):277–290.
- Jordan WC, Graves SC (1995) Principles on the benefits of manufacturing process flexibility. *Management Sci.* 41(4):577–594.
- Leachman R (1987) Preliminary design and development of a corporate-level production system for the semiconductor industry. Working Paper ORC 86-11. Operations Research Center, University of California, Berkeley, CA.
- Liu Q, van Ryzin G (2008) On the choice-based linear programming model for network revenue management. *M&SOM* 10(2):288–310.
- Murty KG (1983) *Linear Programming* (John Wiley & Sons, New York).
- Netessine S, Dobson G, Shumsky RA (2002) Flexible service capacity: Optimal investment and the impact of demand correlation. *Oper. Res.* 50(2):375–388.
- Parlar M, Goyal S (1984) Optimal ordering decisions for two substitutable product with stochastic demands. *OPSEARCH* 21:1–15.
- Pasternack BA, Drezner Z (1991) Optimal inventory policies for substitutable commodities with stochastic demand. *Naval Res. Logist.* 38(2):221–240.
- Robinson LW (1990) Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Oper. Res.* 38(2):278–295.
- Rockafellar RT (1996) *Convex Analysis* (Princeton University Press, Princeton, NJ).
- Savin SV, Cohen MA, Gans N, Katalan Z (2005) Capacity management in rental businesses with two customer bases. *Oper. Res.* 53(4):617–631.
- Shumsky RA, Zhang F (2009) Dynamic capacity management with substitution. *Oper. Res.* 57(3):671–684.
- Simchi-Levi D, Chen X, Bramel J (2014) *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management*, 3rd ed. (Springer, New York).
- Steinhardt C, Gönsch J (2012) Integrated revenue management approaches for capacity control with planned upgrades. *Eur. J. Oper. Res.* 223(2):380–391.
- Talluri KT, van Ryzin GJ (2004a) Revenue management under a general discrete choice model of consumer behavior. *Management Sci.* 50(1):15–33.
- Talluri KT, van Ryzin GJ (2004b) *The Theory and Practice of Revenue Management* (Springer, New York).
- Topkis DM (1968) Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Sci.* 15(3):160–176.
- Topkis DM (1998) *Supermodularity and Complementarity* (Princeton University Press, Princeton, NJ).
- Van Mieghem JA (1998) Investment strategies for flexible resources. *Management Sci.* 44(8):1071–1078.
- Van Mieghem JA (2003) Capacity management, investment, and hedging: Review and recent developments. *M&SOM* 5(4):269–302.
- Van Mieghem JA, Rudi N (2002) Newsvendor networks: Inventory management and capacity investment with discretionary activities. *M&SOM* 4(4):313–335.

- Wall M (2003) Manufacturing flexibility. *Automotive Industries* 183(10): 44–45.
- Xu H, Yao DD, Zheng S (2011) Optimal control of replenishment and substitution in an inventory system with nonstationary batch demand. *Production Oper. Management* 20(5):727–736.
- Yao DD, Zheng S (2003) Substitutable inventory: Single- and multiperiod models. Shanthikumar JG, Yao DD, Henk Z, eds. *Stochastic Modeling and Optimization of Manufacturing Systems and Supply Chains*, Vol. 63 (Kluwer Academic Publishers, Boston), 177–201.
- Zhang D (2011) An improved dynamic programming decomposition approach for network revenue management. *M&SOM* 13(1):35–52.
- Zhang D, Cooper WL (2005) Revenue management for parallel flights with customer-choice behavior. *Oper. Res.* 53(3):415–431.

Yueshan Yu is a research scientist at Amazon.com, Inc. He obtained his Ph.D. in operations and manufacturing management from Olin Business School, Washington University in St. Louis.

Xin Chen is a professor of industrial engineering at the University of Illinois at Urbana–Champaign. His research interests lie in revenue management, supply chain management, data analytics, and optimization. He received the INFORMS Revenue Management and Pricing Section Prize in 2009 and is the coauthor of the book “The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management (Second Edition and Third Edition).”

Fuqiang Zhang is a professor of operations and manufacturing management at Olin Business School, Washington University in St. Louis. His main research interests are in supply chain management, consumer analytics in operations management, and sustainable operations. His research has appeared in *Management Science*, *M&SOM*, *Operations Research*, and *Production and Operations Management*.