



City Research Online

City St George's, University of London

Citation: Ceryan, O., Duenyas, I. & Sahin, O. (2018). Dynamic Pricing and Replenishment with Customer Upgrades. *Production and Operations Management*, 27(4), pp. 663-679. doi: 10.1111/poms.12816

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/20877/>

Link to published version: <https://doi.org/10.1111/poms.12816>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Dynamic Pricing and Replenishment with Customer Upgrades

Oben Ceryan*

Cass Business School, City, University of London, UK
oben.ceryan@city.ac.uk

Izak Duenyas

Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI
duenyas@umich.edu

Ozge Sahin

Carey Business School, Johns Hopkins University, Baltimore, MD
ozge.sahin@jhu.edu

To appear, Production & Operations Management

We study a joint implementation of price- and availability-based product substitution to better match demand and constrained supply across vertically differentiated products. Our study is motivated by firms that utilize dynamic pricing as well as customer upgrades, as ex-ante and ex-post mechanisms, respectively, to mitigate inventory mismatches. To gain insight into how offering product upgrades impacts optimal price selection, we formulate a multiple period, nested two-stage model where the firm first sets prices and replenishment levels for each product while the demand is still uncertain, and after observing the demand, decides how many (if any) of the customers to upgrade to a higher quality product. We characterize the structure of the optimal upgrade, pricing and replenishment policies and find that firms having greater flexibility to offer product upgrades can restrain their reliance on dynamic pricing, enabling them to better protect the price differentiation between the products. We also show how the quality differential between the products or changes in the replenishment cost structures influence the optimal policy. Using insights gained from the optimal policy structure, we construct a heuristic policy and find that it performs well across various parameter values. Finally, we consider an extension in which the firm dynamically sets upgrade fees in each period. Our results overall help further our understanding of the intricate relationship among a firm's decisions on pricing, replenishment, and product upgrades in an effort to better match demand and constrained supply.

Keywords: Dynamic pricing, Customer upgrades, Revenue management, Joint pricing and inventory control

* Corresponding author

1. Introduction

Firms offering vertically differentiated products may occasionally choose to substitute a higher quality product to satisfy demand for a lower quality product. This downward substitution, also referred to as firm-driven customer upgrades, can potentially reduce inventory costs and diminish customer dissatisfaction by better aligning supply and demand. Bassok et al. (1999) discuss several instances where such substitution takes place. For example, in semiconductor manufacturing, a high speed chip or a larger capacity memory can be used to satisfy demands for lower speed or memory devices. In the steel industry, high strength steel may be used as a substitute for lower strength steel. Similarly, cold-rolled steel, which is generally used for higher surface quality applications such as exposed automobile and appliance panels may be used to satisfy demand for hot-rolled steel that is generally used in non-critical surface applications such as automotive frames or wheels.

The joint stocking and upgrading decisions a firm faces are often also further intricately linked with dynamically adjusted prices. Consider again the steel industry, where prices are strongly influenced by iron ore commodity prices. As stated in a recent report by McKinsey&Company (2013), “it does not take much for iron ore to flip from scarcity to surplus supply.” Indeed, according to a Commodity Market Monthly by IMF (2013), the price of iron ore soared 17 percent in January 2013 and was up more than 50 percent the previous four months. This price increase was then followed up by more than 40 percent decline by the end of 2014. Overall, based on this paper’s authors’ calculations using data sourced from the World Bank, the coefficient of variation for average monthly iron ore price has been 19 percent over the three year period ending on December 2014, with a mean month-on-month absolute price change of 6 percent. As a result, steelmakers dynamically set prices for their products taking into account the volatile iron ore commodity prices, along with other factors such as the current state of demand from industries they cater to (i.e., the construction, appliance, and the automobile industries), their current level of stock, and their available production capacity. Thus, a steelmaker facing shortages of lower-strength steel may readjust prices while considering whether it will also be more profitable to offer its high-strength steel to be substituted to meet some of the demand for lower-strength steel.

The applications of firm driven upgrades with dynamic price adjustments and replenishments extend beyond the manufacturing industries. As other examples, consider a retailer carrying multiple versions of a product with different perceived quality levels such as a smaller vs larger storage sizes for electronic devices, store-brand vs national-brand products, or regular-quantity vs larger-quantity packages of the same brand. When faced with shortages for the lower quality item, the retailer may offer a ‘rain check’ allowing the customer to backlog the item at the prevailing price,

or upgrade them to a higher quality item. In fact, many retailers' terms and conditions state that, at the sole discretion of the retailer and at no cost to the customer, a product that is of equal or greater value (i.e., larger quantity or an updated product) may be substituted for an out-of-stock product. The practice is also in line with consumer protection laws against 'bait and switch', an insincere offer to sell one item in order to induce the buyer to purchase another as described in Friedman (2013).

Dynamic pricing, in which prices are set to partially influence the demand may be seen as an ex-ante tool to provide price-based product substitutions. In the same light, product upgrades are often implemented after demand is observed, and the practice may be viewed as an ex-post tool in the form of a one-way, availability-based product substitution. Our goal in this paper is to study the simultaneous application of these two tools to better match supply and demand. Specifically, we would like to answer the following main questions that naturally arise when pricing, replenishment, and upgrades are jointly considered: (1) How should a firm decide on extending product upgrade offers in any given period? (2) How should the firm set prices and replenishment levels for its products in each period? (3) What is the impact of a firm's willingness to offer upgrades on its pricing and replenishment decisions?

To do so, we formulate the firm's production, pricing, and upgrade problem as a multiple period, two-stage, finite-horizon stochastic dynamic program. At the beginning of each period, the firm reviews the current inventory levels and decides on the replenishment quantities and the optimal prices to be applied within the period. The optimal replenishment quantities are constrained by the firm's limited capacity for the current period. We assume the demands for both products are correlated through a linear, additive, stochastic demand model. After the demands for both products are realized at the end of each period, the firm has the option to upgrade part of the demand for the lower quality item to the higher quality item.

Our first contribution in this paper is characterizing the structure of the optimal upgrade, replenishment, and pricing policies. We show that the second-stage optimal upgrade policy is defined by a protection level on the higher quality item, where the protection level depends on the intermediate inventory positions of the products through their sum. Further, we provide monotonicity results on how this threshold changes with the total intermediate inventory. Taking into consideration the structure of the optimal upgrade policy, we then address the first-stage replenishment and pricing decisions. We find that the optimal replenishment policy for the products follow state-dependent modified base-stock levels, where the base-stock levels are decoupled for some initial inventories, and are decreasing with the inventory of the other item otherwise. In each period, each item that

requires replenishment is brought either up to its base-stock level or it is replenished to the full extent of its capacity if base stock level can not be achieved by the available capacity. The capacity limitations give rise to various regions of price surcharges, list prices, and price discounts. We provide results on how prices for the items change with respect to the starting inventories.

Second, we analytically explore the impact of offering product upgrades to the firm's pricing and replenishment decisions. Particularly, we find that offering product upgrades in instances where there are inventory imbalances between the products may allow the firm to restrain its reliance on dynamic pricing as the sole mechanism to counteract the imbalance. Consequently, the optimal price difference between the products in such instances is closer to their respective list prices, enabling the firm to have a more consistent price positioning between the products.

Third, we study how the quality differentiation between the products or changes in replenishment costs influence the optimal policy. We show that an increase in the quality difference between the products leads the firm to increase the base-stock level for the higher quality product, lower the base-stock level for the lower quality product and apply list prices that are further apart. Regarding replenishment costs, we find that the firm's pick of replenishment level and prices leads it to offer fewer subsequent customer upgrades if the replenishment cost for the higher quality product increases and to offer more upgrades if the cost for the lower quality product increases. We also provide sensitivity results for a correlated cost structure with a change in an underlying cost parameter driving the replenishment costs for both products.

In addition to these main theoretical contributions, we utilize the insights gained from the optimal policy structure to construct an easier to implement heuristic policy that would be valuable in practice. Through a numerical study, we compare the profits obtained by the heuristic policy with the optimal profit across a variety of parameters and show that the heuristic policy consistently performs well.

Lastly, we extend our model to a setting where the firm sets an upgrade fee in each period that results in a proportion of customers being interested in paying the additional fee to receive an upgrade. We show that the firm will charge more for an upgrade if the availability of the higher quality product is lower, and charge less for an upgrade when the number of customers who were unable to get the lower quality product is larger.

2. Related Literature

The problem we investigate in this paper is mainly related to two major areas, stocking under availability-based product substitutions and dynamic pricing with replenishment. There has been considerable prior interest in availability-based product substitutions, generally classified into

‘consumer-driven’ and ‘firm-driven’ substitutions. In consumer-driven substitution models, a customer who arrives to find their first choice product no longer available will determine which (if any) of the other available products they would purchase. The area of consumer-driven product substitutions is well studied and there exists a vast literature on stocking levels incorporating choice models. Mahajan and van Ryzin (2001) show that under substitution the firm should stock relatively more of popular variants and less of unpopular variants. Nagarajan and Rajagopalan (2008) observe partially decoupled optimal inventory decisions if a fixed proportion of customers who do not find their first choice available switches to purchase the other product. There has also been recent interest in making pricing decisions in conjunction with inventory decisions. Karakul and Chan (2008) and Transchel (2011) study single-period models where the firm is a price taker for the lower quality product and sets the price for the higher quality product. Assuming a fixed substitution rate, Karakul and Chan (2008) show that substitution leads to higher prices and safety stock for the higher quality product and lower safety stock for the lower quality product. Transchel (2011) allows the substitution rate to be a function of the price of the higher quality product and compares prices for centralized and decentralized decision making approaches. Tomlin and Wang (2008) consider the pricing, inventory and downconversion problem for a single period, vertically differentiated two-product model and characterize optimal recourse prices. There is also a related stream of work on pricing and upselling, in which a customer is required to make a side payment to receive a higher value product. For examples of recent work on upselling, we refer the reader to Aydin and Ziya (2008), Gallego and Stefanescu (2009), and Cui et al. (2017).

In contrast to consumer-driven models, in firm-driven substitutions, it is the firm that chooses to implement downward substitutions for items that experience shortages. In one of the earliest works, Pasternack and Drezner (1991) examine a single-period model with two substitutable products where substitution takes place if one product has excess inventory while the other faces shortages. They study the effects of substitutability on products’ optimal stocking levels. For the case of unidirectional substitutions, they show that the optimal stocking quantity for the product that can be used as a substitute is higher whereas the stocking quantity for the other product is lower compared to a setting without substitutions. Bassok et al. (1999) study a single-period model for an arbitrary number products that differ in quality and have full downward substitutability, that is, a customer demand for any particular product can be satisfied by any of the higher quality products. They show that inventory ordering follows a base-stock policy where the base-stock level for a product is non-increasing with the starting inventory of other products. In a closely related work, Netessine et al. (2002) study a similar setting as in Bassok et al. (1999) with an emphasis on the impact of demand correlation and with product upgrades restricted to at most one level.

The collection of works we mentioned so far have considered single period models. Recent contributions to this body of work include extensions to settings with multiple periods of allocation for an initial inventory or capacity. Shumsky and Zhang (2009) consider a setting in which a firm first makes an initial purchase of various types of capacities that can be used to satisfy its own demand as well as demand for the capacity one level below. They assume that the products can be ranked by their contribution margins and show that the allocation decisions in each period follow a rationing policy described by protection limits on the higher level capacity. Their work is extended in a subsequent study by Yu et al. (2015) that allows multi-level downward substitution. In both papers, product prices are exogenous which allows a monotonic ranking of contribution margins. In contrast, our paper considers the additional dynamic pricing and replenishment decisions in a multiple period setting together with the upgrade decisions.

Regarding the dynamic pricing literature, extensive reviews have been provided by Elmaghraby and Keskinocak (2002), Bitran and Caldentey (2003), Chan et al. (2004), and Chen and Simchi-Levi (2012). Here, we limit our discussion to joint pricing and replenishment decisions that are closest to our setting. One of the earliest works to consider a multi-period, joint pricing and inventory control problem is by Federgruen and Heching (1999). For a single product case, they show that the optimal policy can be characterized by a base-stock, list price pair. Specifically, in periods when it is optimal to order, the inventory is replenished up to a base-stock level and a list price is charged. For periods starting with excess inventory, no ordering takes place and a discount on the list price is applied. Zhu and Thonemann (2009) introduce a second, price-substitutable product to the multi-period, joint inventory and pricing control problem where they model the price substitution between products through a linear additive demand-price relationship. They show that the main findings of Federgruen and Heching (1999) does indeed extend to the two-product case and the optimal policy for each product is a base-stock, list price type, i.e., whenever an order is placed for a product, its inventory is brought up to a base-stock level and the product is charged its list price. The two dimensional case also brings additional insights: the base-stock level for each product is state dependent and decreases with the starting inventory of the other product. Ceryan et al. (2013) study the role of flexible capacity in the joint pricing and inventory control problem by considering a general capacity portfolio composed of product-dedicated and flexible resources. They show that limited capacity introduces a price-surcharge component to the optimal policy structure, and more importantly, they find that the availability of a flexible resource helps maintain stable price differences across products over time. In relation to the existing work on dynamic pricing and replenishment literature, in this paper we incorporate a nested second stage upgrade

decision at the end of each period in a multiple period capacitated setting as an additional means of reducing the mismatch between the demand and supply of each product type.

To summarize, our main contributions to the related body of work through this study are as follows: With respect to the existing dynamic pricing literature, we not only consider ex-ante price-based substitution, but also incorporate a nested, second-stage availability-based substitution via product upgrades. Through this generalization, we are able to explicitly characterize the impact of product upgrades on the optimal pricing and replenishment policy. As compared to the multi-period, firm-driven product substitution literature, we are endogenizing product prices in each period as an additional ex-ante form of partial product substitution and are also incorporating capacitated dynamic replenishment decisions. To our knowledge, we are also the first ones to explicitly incorporate quality differentiation in a consumer utility model in a multi-period problem to investigate how optimal dynamic pricing, replenishment and upgrade decisions are influenced by the quality differentiation between the products.

3. Problem Formulation

We model the dynamic pricing, replenishment and upgrade decisions for a firm that offers two products that are vertically differentiated by their quality level. Following Mussa and Rosen (1978), Bresnahan (1981), and more recently Mantin et al. (2014), we consider a consumer utility model where the surplus a consumer with valuation v receives from purchasing a product j , $j=\{1,2\}$, with price p_j and quality level q_j is assumed to be of the form $vq_j - p_j$. We index the products such that product type-1 indicates the higher quality product and product type-2 refers to the lower quality product, i.e., $q_1 > q_2$. Consumers purchase the product that provides them with the highest surplus. A consumer with valuation v_{12} will be indifferent between purchasing product type-1 or product type-2 if $v_{12}q_1 - p_1 = v_{12}q_2 - p_2$, i.e., if $v_{12} = (p_1 - p_2)/(q_1 - q_2)$. Further, consumers with valuations $v > v_{12}$ will prefer product type-1 over product type-2 and those with valuations $v < v_{12}$ will prefer product type-2 over product type-1. In addition, suppose that there exists products in the market, other than the ones offered by the firm of interest, with price and quality levels $(\underline{p}, \underline{q})$ and $(\bar{p}, \bar{q})$ such that $\underline{p} < p_j < \bar{p}$ and $\underline{q} < q_j < \bar{q}$. Assuming that v is distributed uniformly with density δ on $[0, \bar{v}]$ where $\delta\bar{v}$ is the total market size, and that all products are viable in the market, one can show that the expected demand for the products are given by the following:

$$d_1(p_1, p_2) = \delta \left[\frac{\bar{p}}{(\bar{q} - q_1)} - \left(\frac{1}{\bar{q} - q_1} + \frac{1}{q_1 - q_2} \right) p_1 + \frac{1}{(q_1 - q_2)} p_2 \right] \quad (1)$$

$$d_2(p_1, p_2) = \delta \left[\frac{\underline{p}}{(q_2 - \underline{q})} + \frac{1}{(q_1 - q_2)} p_1 - \left(\frac{1}{q_1 - q_2} + \frac{1}{q_2 - \underline{q}} \right) p_2 \right] \quad (2)$$

We further incorporate additive uncertainty into the demand-price function by letting $D_1^t(p_1, p_2, \epsilon_1^t)$ and $D_2^t(p_1, p_2, \epsilon_2^t)$ denote the current period demands with $D_1^t(p_1, p_2, \epsilon_1^t) = d_1(p_1, p_2) + \epsilon_1^t$ and $D_2^t(p_1, p_2, \epsilon_2^t) = d_2(p_1, p_2) + \epsilon_2^t$ where ϵ_1^t and ϵ_2^t refer to independent random variables having continuous probability distributions with zero mean and nonnegative support on product demands. (We discuss the impact of correlated random variables in Section 6.)

Before we introduce the firm's decisions and the sequence of events, we would like to make two additional remarks about the demand model. First, notice that as the quality difference between the products ($q_1 - q_2$) increases, the cross-price sensitivity coefficients decrease. This characteristic of the model is in line with empirical findings such as that of Sethuraman et al. (1999), where they show that brands that are closer to each other in terms of price and quality exhibit stronger cross-price effects. Second, for any given p_1 , let $\bar{p}_2(p_1)$ be the price that would make the expected demand for product type-2 vanish. Solving for p_2 in (2) with $d_2(p_1, p_2) = 0$, we find $\bar{p}_2(p_1) = \left(\frac{\underline{p}}{q_2 - \underline{q}} + \frac{p_1}{q_1 - q_2} \right) / \left(\frac{1}{q_1 - q_2} + \frac{1}{q_2 - \underline{q}} \right) < p_1$ where the inequality follows from $\underline{p} < p_1$. Thus, given there is a quality difference between the products, the expected demand for the lower quality product vanishes before its price equals the price of the higher quality product, in line with vertical differentiation.

At the beginning of each period t of a finite planning horizon of length T , the firm reviews the current inventory positions x_1^t and x_2^t for the higher quality and lower quality product, respectively. It first simultaneously decides on (i) the prices, p_1^t and p_2^t to charge during the period for the products that will influence their respective demands D_1^t and D_2^t observed within that period, and (ii) the optimal replenishment quantities implied by replenish-up-to levels y_1^t and y_2^t , which are constrained by limited replenishment capacities K_1 and K_2 . After the demand in period t is realized, the firm observes its remaining intermediate inventories for both products. It then also has the option to offer some customers product upgrades, hence satisfying their original demand for the lower quality product by supplying them with an upgrade to the higher quality product.

We formulate the firm's replenishment and pricing problem as well as its upgrade decisions through a multiple period, nested two-stage model. Letting $V^t(x_1^t, x_2^t)$ denote the expected discounted profit-to-go function under the optimal policy starting at state (x_1^t, x_2^t) with t periods remaining until the end of the horizon, the problem can be expressed as a stochastic dynamic program satisfying the following recursive relations:

Stage 1:

$$V^t(x_1^t, x_2^t) = \max_{\substack{y_1^t, p_1^t \\ x_i^t \leq y_i^t \leq x_i^t + K_i}} \underbrace{R(p_1^t, p_2^t)}_{\text{E[Revenue]}} \underbrace{-c_1^t \cdot (y_1^t - x_1^t) - c_2^t \cdot (y_2^t - x_2^t)}_{\text{Replenishment cost}} + \underbrace{\text{E}_{D_1^t, D_2^t} [G^t(y_1^t - D_1^t, y_2^t - D_2^t)]}_{\substack{\text{Intermediate inventory} \\ \text{E[Second-stage profit-to-go for period } t]}} \quad (3)$$

Stage 2:

$$G^t(w_1^t, w_2^t) = \max_{\substack{u^t \\ u^t \geq 0}} \underbrace{-h_1(w_1^t - u^t) - h_2(w_2^t + u^t)}_{\text{Holding \& shortage cost after upgrades}} + \underbrace{\beta V^{t-1}(w_1^t - u^t, w_2^t + u^t)}_{\text{First-stage profit-to-go for period } t-1} \quad (4)$$

In (3), the term $R(p_1^t, p_2^t)$ represents the expected revenue in period t , which can be expressed as $R(p_1^t, p_2^t) = p_1^t \cdot d_1^t(p_1^t, p_2^t) + p_2^t \cdot d_2^t(p_1^t, p_2^t)$ where, as introduced earlier, $d_1^t(p_1^t, p_2^t)$ and $d_2^t(p_1^t, p_2^t)$ refer to the mean demand for product type-1 and product type-2, respectively. The terms $c_1^t \cdot (y_1^t - x_1^t)$ and $c_2^t \cdot (y_2^t - x_2^t)$ correspond to replenishment costs, where c_i^t (with $c_1^t \geq c_2^t$) and $y_i^t - x_i^t$ refer, respectively, to the unit replenishment cost and the current period replenishment quantity for product type- i . The current period replenishment quantity for product type- i is limited by the available replenishment capacity for that product, K_i , as indicated by the constraints $x_i^t \leq y_i^t \leq x_i^t + K_i$. Finally, the expected profit-to-go term $E_{D_1^t, D_2^t}[G^t(y_1^t - D_1^t, y_2^t - D_2^t)]$, with its arguments as the intermediate inventory positions for the items after replenishment and demand realizations is obtained through the second-stage optimal upgrade problem as described next.

In (4), w_1^t and w_2^t correspond to intermediate inventory positions for the higher quality and lower quality products, respectively. In other words, (w_1^t, w_2^t) is a particular realization of $(y_1^t - D_1^t, y_2^t - D_2^t)$. When the firm decides to upgrade u^t of the customers that initially requested the lower quality product, the final inventory positions for the products after the upgrade decision can be expressed by $w_1^t - u^t$ for the higher quality product and $w_2^t + u^t$ for the lower quality product. The constraint $u^t \geq 0$ guarantees that the upgrade quantity is nonnegative, implying unidirectional substitutions of the higher quality product for the demand for the lower quality product and not vice versa. After the upgrade decision is made, the firm incurs linear holding and backorder costs on ending inventories. To facilitate the analysis of the optimal upgrade policy, we also define $u_1^t = w_1^t - u^t$ and $u_2^t = w_2^t + u^t$ to represent final inventory positions after upgrades. The holding and backorder cost for product type- i is then denoted by $h_i^t(u_i^t)$, which is defined as $h_i^t(u_i^t) := h_i^{t+} u_i^{t+} + h_i^{t-} u_i^{t-}$ for $i = 1, 2$ where h_i^{t+} and h_i^{t-} represent the unit holding and backorder cost, respectively, and $u_i^{t+} := \max(0, u_i^t)$, $u_i^{t-} := \max(0, -u_i^t)$. We let β denote the discount factor and the terminal value function be represented as $V^0(x_1^0, x_2^0) = -c_1^0 x_1^{0-} - c_2^0 x_2^{0-}$. Such a terminal value function may especially be appropriate for perishable items belonging to product categories such as technologically rapidly advancing products in a manufacturing setting or baked goods in a retail setting. In other applications such as the non-perishable steel example, excess inventory at the end of the planning horizon may continue to preserve their value. We would like to point out that a simply modified version of the terminal value function of the form $V^0(x_1^0, x_2^0) = c_1^0 x_1^0 + c_2^0 x_2^0$, which

rewards (penalizes) any excess (shortage) based on the replenishment cost would also preserve all our results.

Before we proceed to our discussion of the optimal policy structure, we would like to make two remarks regarding our formulation for the second-stage upgrade decision given in (4). First, we allow for upgrades even if both products have positive intermediate inventories, which can indeed be optimal in the presence of capacity limitations as we will discuss following the structural results on the optimal policy. Second, our formulation permits any unsatisfied outstanding demand for product type-2 from earlier periods to be eligible for upgrades.

4. Structure of the Optimal Pricing, Replenishment and Upgrade Policy

In order to characterize the structure of the optimal policy in period t , we first focus on the second-stage upgrade problem and describe the optimal upgrade policy. By incorporating this optimal upgrade policy, we then proceed to show the structure of the optimal production and pricing decisions for the first stage. Throughout the paper, we use ‘increasing’ and ‘decreasing’ in their weak sense, i.e., non-decreasing and non-increasing, respectively. Where applicable, we will denote strict monotonicities by the terms ‘strictly increasing’ and ‘strictly decreasing’.

4.1. Optimal Customer Upgrade Policy

Consider any intermediate inventory position (w_1^t, w_2^t) in the beginning of the second-stage upgrade problem. As described earlier in the problem formulation, this intermediate inventory position reflects the inventory of each product after it is augmented by the current period production quantity and depleted by the current period observed demand. We introduce and let $w^t := w_1^t + w_2^t$ denote the total intermediate inventory position. The below result summarizes the structure of the optimal upgrade policy.

THEOREM 1. (Optimal Upgrade Policy) *The optimal upgrade policy is defined by a protection level $r^t(w^t)$ on the higher quality product. Specifically, it is optimal to upgrade $u^t = (w_1^t - r^t(w^t))^+$ customers by satisfying their initial demand for the lower quality product by upgrading them to the higher quality product. Further, $r^t(w^t)$ is increasing with respect to w^t , and $r^t(w^t) - w^t$ is decreasing with respect to w^t .*

Proof: Proofs of all results are provided in the Online Appendix.

Theorem 1 states that the main component of the optimal upgrade policy is a *protection level* on the higher quality product, $r^t(w^t)$, which depends on the intermediate inventory position of the

products only through their sum, w^t . This protection level on the higher quality product increases with the total intermediate inventory level w^t , but this increase is at most as much as the increase in w^t itself.

We would like to note that upgrades may occur even if the inventory position of the higher quality product is negative. In this context, an upgrade indicates that the demand will be satisfied through a higher quality product in future periods. In other words, a customer currently in the backlog of the lower quality product may be offered a place in the backlog of the higher quality product instead if it is more profitable for the firm to do so, e.g., in instances for which the firm can clear the higher quality backlog more quickly than the lower quality backlog. Our model thus allows for more flexibility in the firm's upgrade decision, as when such a shift is not profitable for the firm, it still has the option not to upgrade any customers. Another interesting point is that when the intermediate inventory positions w_1^t and w_2^t are both positive and both products require replenishments in the subsequent period facing no capacity limitations, upgrades will not occur provided that $h_1^+ - h_2^+ < \beta(c_1 - c_2)$ in a stationary parameter setting, i.e., the benefit of an upgrade in terms of the current savings in holding costs is lower than the discounted replenishment cost increase in the subsequent period. (A common convention in the literature is to consider the unit holding cost as a fraction of the purchase/replenishment cost, i.e., $h_i^+ = \alpha c_i$ for some $0 < \alpha < 1$. Consequently, as long as the fraction α is less than the per period discount factor β , which is almost guaranteed in most practical settings, the firm will not upgrade if both products have positive inventory.) However, as our formulation in (4) allows and Theorem 1 implies, offering upgrades even if both products have positive intermediate inventories can be optimal when there are capacity limitations. That is, even if the lower quality product has positive intermediate inventory, offering upgrades in the current period may be more profitable for the firm compared to starting the next period in a less favorable inventory position requiring stronger price adjustments due to limited replenishment capacity. Our results in the next section will show the role capacity limitation plays in restricting replenishment of a product up to a desired inventory level and its consequences on pricing.

4.2. Optimal Replenishment Policy

Having characterized the second-stage optimal upgrade policy in period t for any particular intermediate inventory position, we now turn our attention to the first-stage pricing and replenishment decisions at the beginning of period t . The pricing and replenishment decisions are made simultaneously and their derivation is through a joint analysis. For expositional clarity though, we present




Figure 1 Optimal replenishment policy structure

the structures of the replenishment and pricing policies separately, starting with the optimal replenishment policy that is described by Theorem 2 below.

THEOREM 2. (Optimal Replenishment Policy) *For $i, j = 1, 2, j \neq i$, the optimal replenishment policy for product type- i is defined by a partially decoupled state-dependent base-stock level $y_i^{*t}(x_j^t)$. For any starting inventory pair (x_1^t, x_2^t) , it is optimal to replenish $\min((y_i^{*t}(x_j^t) - x_i^t)^+, K_i)$ units of product type- i . Further, let (x_1^{ot}, x_2^{ot}) be defined such that $x_1^{ot} = y_1^{*t}(x_2^{ot})$ and $x_2^{ot} = y_2^{*t}(x_1^{ot})$. Then, the base-stock level for each product type- i , $y_i^{*t}(x_j^t)$, is independent of the inventory of the other product, x_j^t , for $x_j^t - K_j < x_j^t \leq x_j^{ot}$, and strictly decreasing with the inventory of the other product, x_j^t , otherwise.*

As indicated by Theorem 2, the replenishment decision for each product follows a modified state-dependent base-stock policy, where the base-stock level for product type- i , $y_i^{*t}(x_j^t)$, is a function of the inventory level of the other product, x_j^t . No replenishment takes place for a product if its inventory level is above its current period base-stock, i.e., if $x_i^t \geq y_i^{*t}(x_j^t)$. If a product requires replenishment, its inventory is brought up to its base-stock level $y_i^{*t}(x_j^t)$ if its capacity K_i permits. Otherwise, if its capacity is limiting, the product is replenished by the full extent of its capacity.

Figure 1 conceptually illustrates the structure of the optimal replenishment policy given in Theorem 2. The initial inventory state space is collectively partitioned into nine regions based on whether i) each product type- i already has sufficient inventory, i.e., $x_i^t \geq y_i^{*t}(x_j^t)$, ii) is understocked

yet the capacity is adequate to bring its inventory to its base-stock level, i.e., $y_i^{*t}(x_j^t) - K_i \leq x_i^t < y_i^{*t}(x_j^t)$, or iii) is understocked and its limited capacity prevents the firm from replenishing it up to the desired base-stock level, i.e., $x_i^t < y_i^{*t}(x_j^t) - K_i$. As the structure of the replenishment decisions for three of these regions can be symmetrically described, we only exemplify and highlight six distinct cases. Initial inventory positions labeled A-F, with their corresponding inventory positions after replenishment indicated by the prime symbol ('), exemplify these six distinct cases: (A) neither product is replenished, (B) only one product (type-2) is replenished through adequate capacity to its base-stock level, (C) only one product (type-1) is replenished by the full extent of its available capacity, (D) both products are replenished through their adequate capacities up to their base-stock levels, (E) one product (type-1) is replenished up to its base-stock level and the other product (type-2) is replenished by the full extent of its available capacity, and (F) both products are replenished by the full extent of their available capacities.

Notice in Figure 1 case (D) that, for $x_2^{\circ t} - K_2 < x_2^t \leq x_2^{\circ t}$, the base-stock level for product type-1, $y_1^{*t}(x_2^t)$, is independent of x_2^t and equals $x_1^{\circ t}$. Similarly, the base-stock level for product type-2 is independent of x_1^t and equals $x_2^{\circ t}$ for $x_1^{\circ t} - K_1 < x_1^t \leq x_1^{\circ t}$. If on the other hand, product type-2 is understocked such that its available capacity is not sufficient to bring its inventory to its desired base-stock level as in case (E), then the base-stock level for product type-1 will be higher than the independent base-stock level $x_1^{\circ t}$ and will increase further as the starting inventory level of the product type-2 decreases. Similarly, when product type-2 requires no further replenishment, the base-stock level for product type-1 is less than $x_1^{\circ t}$ and decreases further as the starting inventory level of product type-2 increases.

There are two drivers that contribute to a higher base-stock level for product type-1 when the inventory position of product type-2 is low. First, when the firm has fewer inventories for the lower quality type-2 product, it will be more likely to offer upgrades once demand is realized. Hence, the overall requirement on the amount of the higher quality type-1 product will be larger, increasing its base-stock level. Second, due to the substitutable nature of the products and as we will see next, the firm may also choose to alter the prices to shift more demand from the lower quality product to the higher quality product, thereby further increasing the quantity needed of the higher quality product. Similar reasoning applies for the monotonicity of the base-stock level for product type-2.

4.3. Optimal Pricing Policy

As we set forth in the model formulation, the firm selects prices for each product at the beginning of each period while simultaneously determining the replenishment decisions. This ex-ante choice of prices may partially shift demand from one product to the other as well as elevate or suppress

demand for both products. While determining the prices, the firm takes into account the possibility of subsequent upgrades and this flexibility to realign demand through prices allows the firm to improve profits further than what would have been possible by solely relying on upgrades to mitigate inventory imbalances.

Our next focus is the characterization of the optimal pricing policy. To do so, we first introduce a ‘list price’ for each product for each period. The description of the optimal pricing policy is in reference to these list prices and requires the same state-space segmentation presented earlier for the replenishment policy.

THEOREM 3. (Optimal Pricing Policy) *Let $p_1^{\circ t} = \frac{\bar{p}(q_1 - q) + p(\bar{q} - q_1)}{2(\bar{q} - q)} + \frac{c_1^t}{2}$ and $p_2^{\circ t} = \frac{\bar{p}(q_2 - q) + p(\bar{q} - q_2)}{2(\bar{q} - q)} + \frac{c_2^t}{2}$ denote the ‘list price’ in period t for product type-1 and product type-2, respectively.*

*For any initial inventory pair (x_1^t, x_2^t) at the beginning of period t with the corresponding base-stock levels $y_1^{*t}(x_2^t)$ and $y_2^{*t}(x_1^t)$, it is optimal to apply product type- i its list price, i.e. $p_i^{*t}(x_1^t, x_2^t) = p_i^{\circ t}$, if $y_i^{*t}(x_j^t) - K_i \leq x_i^t \leq y_i^{*t}(x_j^t)$, a discount, i.e. $p_i^{*t}(x_1^t, x_2^t) < p_i^{\circ t}$, if $x_i^t > y_i^{*t}(x_j^t)$, or a surcharge, i.e., $p_i^{*t}(x_1^t, x_2^t) > p_i^{\circ t}$ if otherwise ($x_i^t < y_i^{*t}(x_j^t) - K_i$). Furthermore, the price of each product is decreasing with the inventory of either product.*

Theorem 3 states that the optimal pricing policy for each product consists of applying list prices, price surcharges, or price discounts across nine subregions of the state-space segmentation (based on whether each product has inventory beyond its current-period desired base-stock level, has adequate capacity to reach its base-stock level, or faces capacity limitations, as described in Theorem 2 and illustrated in Figure 1). Specifically, when a product requires replenishment and its capacity is adequate, it is optimal to apply its list price. It is interesting to note that, while the quality levels of both products impact the demand for either product as described in the demand model given in (1) and (2), we find that the profit maximizing optimal list price for each product depends only on its own quality level in reference to the quality and price levels of the outside options. We provide a more comprehensive discussion on the impact of quality difference between the products in Section 5.2.

The list prices maximize the firm’s expected revenue minus replenishment costs, i.e., its net revenue. Any initial inventory position that can guarantee the firm to replenish both units to their desired base-stock level allows the firm to price the products at their list prices in order to gain this highest net revenue. If the initial inventory position does not allow this, for example, when either product faces shortages and their capacity is not sufficient to bring their inventory to a desired level, then the firm needs to alter the prices and sacrifice its net revenue to prevent excessive shortage

costs that it would otherwise face. Hence, outside the list price region, a price surcharge is applied to a product if it experiences inventory shortages such that its capacity is not sufficient to bring its inventory to its desired base-stock level and the magnitude of the price surcharge gets smaller as the initial inventory level of either product is higher. Analogously, a price discount will be given to a product if it has inventory in excess of its current period base-stock level. The magnitude of the price discount gets larger as the initial inventory level of either product is higher.

A direct implication of the optimal pricing policy is that in situations where inventory for product type-1 is high while product type-2 inventory is inadequate, the firm will apply a surcharge to product type-2 to discourage some customers initially requesting product type-2 and instead prompt them to purchase product type-1. This demand realignment occurs even when the firm can also offer upgrades to some product type-2 customers in order to alleviate the inventory imbalance. Therefore, it is important to shed light on the interplay of product upgrades and pricing. Our next focus, and one of our main objectives in this paper, is precisely the identification of the impact of customer upgrades on the pricing and replenishment policy, which we discuss in the following.

5. Sensitivity of the Optimal Policy

In this section, we first provide insights into how offering customers product upgrades influences a firm's optimal pricing and replenishment decisions. Then, we study how the optimal policy changes with the degree of quality differentiation between products as well as with changes in replenishment costs.

5.1. Impact of Customer Upgrades on Pricing and Replenishment

Since prices are set at the beginning of each period and influence the demand over that period, the adjustment of prices may be thought of as enabling ex-ante price-driven partial product substitutions. In contrast, product upgrades are initiated after demand is realized and thus enable ex-post availability-driven one-way product substitutions. When considered jointly, it can be anticipated that the possibility of subsequent product upgrades has some impact on product price and replenishment decisions made in the beginning of a period. Specifically, one might conjecture that the option to offer upgrades might reduce the firm's reliance on the use of pricing to mitigate inventory imbalances. In this section, we show that this conjecture is indeed true by providing insights into how optimal prices and replenishment are affected by a firm's implementation of a product upgrade strategy.

To do so, we first introduce an extension to our model, in which the firm implements an upper limit on the amount of customers that it may subsequently upgrade. We let $\bar{u}^t$ denote a current

period upgrade limit and modify the constraint $u^t \geq 0$ in our original formulation in (4) to $0 \leq u^t \leq \bar{u}^t$. At one extreme, when $\bar{u}^t = 0$, the problem reduces to a dynamic pricing and capacitated replenishment problem for two substitutable products without upgrades. At the other extreme, setting $\bar{u}^t = \infty$ results in the original problem formulation we studied in the preceding section. Thus, the higher the level of $\bar{u}^t$, the greater flexibility the firm has to use upgrades as an ex-post substitution mechanism. We are interested in how higher levels of $\bar{u}^t$ affect the pricing and thus the ex-ante product substitution.

For this setting, the optimal upgrade policy can now be stated as $u^{t*} = \min((w_1^t - r^t(w^t))^+, \bar{u}^t)$. We also note that all our preceding results, including the monotonicity of $r^t(w^t)$ with respect to w^t , and the structure of the optimal pricing and replenishment policies continue to hold in this extension.

Next, we explore analytically the sensitivity of the current period optimal pricing and replenishment decisions with respect to an increase in the upgrade limit, $\bar{u}^t$.

THEOREM 4. (Impact of Upgrades on Pricing and Replenishment) *Suppose the upgrade limit for the second stage of period t , $\bar{u}^t$, increases. Then,*

- (a) *the optimal price for the higher quality product in the beginning of period t increases and the optimal price for the lower quality product decreases,*
- (b) *the base-stock level for the higher quality product increases, whereas the base-stock level for the lower quality product decreases.*

The main highlight of Theorem 4 part (a) is that allowing more upgrades while managing shortages for the lower quality product leads the firm to select prices that are set further apart and that better protect the vertical differentiation between the products. To see why, consider a setting where the firm offers very limited upgrades. When it faces shortages for the lower quality product, it has only limited ability to mitigate this shortage through upgrades. Thus, the firm needs to rely more on pricing to alleviate the shortage. Through a price surcharge on the lower quality product, that may also be accompanied by a price discount on the higher quality product, it suppresses demand for the lower quality product and shifts some demand towards the higher quality product. We now contrast this with a setting where the firm allows more customers to be upgraded. With further upgrades allowed, the firm can now alleviate more of the shortage through upgrades and thus relies less on price changes. Consequently, it does not need to apply as large a price surcharge for the lower quality product or offer as deep a discount on the higher quality product, hence keeps the prices closer to the list prices.

Part (b) of Theorem 4 states that base-stock levels increase for the higher quality product and decrease for the lower quality product as more upgrades are allowed. This follows from the greater flexibility to satisfy unmet demand for the lower quality product through upgrades to the higher quality product. This result also extends the earlier similar finding by Pasternack and Drezner (1991) for a single-period setting with exogenous prices to a multi-period setting, in which the firm also selects optimal prices in each period.

5.2. Sensitivity to Quality Differentiation and Replenishment Cost Parameters

We now identify how the degree of quality differentiation between the products and the changes in replenishment costs influence the optimal policy. In the following, for expositional clarity and tractability, we limit our focus to changes in a single period and to initial inventory states for which the optimal policy is to apply list prices to both products and to replenish the products up to their base-stock levels.

We first study the impact of quality differentiation. The result below provides sensitivity results with respect to an increase in the quality level, q_i , of either product with the assumption that the outside option price $\underline{p}$ is *low enough*, in particular $\underline{p} \leq \left(\frac{q_2 - \underline{q}}{q_1 - q_2} \right)^2 (c_1^t - c_2^t) + c_2^t$.

THEOREM 5. (Sensitivity to Quality Differential) *Suppose the current-period quality level for product type- i , q_i , increases. Then, the optimal list-price and base-stock level for product type- i increases whereas the optimal base-stock level for product type- j ($j \neq i$) decreases with no change in its list price. Further, the expected number of customers subsequently receiving upgrades does not change with a change in only the current-period quality level q_i .*

As one can observe through an immediate inspection of (1) and (2), an increase in the quality level of a product strengthens that product's demand and weakens the demand for the other product. Theorem 5 indicates that this strengthened demand prompts the firm to charge a higher price and increase the base stock level for a product in response to an increase in its quality. Consequently, a weakened demand for the other product causes its base-stock level to decrease. In other words, an increase in the quality differential between the products leads the firm to increase the base-stock level for the higher quality product, lower the base-stock level for the lower quality product and apply list prices that are further apart. (We would like to note that all results except that the base-stock level for the lower quality product is increasing in its own quality level continue to hold even when the outside option does not satisfy the condition stated above.) Interestingly, we also find that the resulting expected demands caused by any changes in the current period quality levels are fully compensated by adjustments in the base-stock levels, thus resulting in no change in the

expected number of customers that would subsequently receive upgrades. It is important to note that this result holds *ceteris paribus*, and that one might expect that a change in the quality levels would have cost consequences as well.

We consider how changes in the replenishment cost structure influences the optimal policy next. We explore (a) changes in each of the individual replenishment cost parameters c_i^t , and (b) a change in an underlying cost parameter c^t , such that $c_1^t = \gamma c^t$ and $c_2^t = c^t$ where $\gamma > 1$ is a cost differential parameter. The latter setting is especially relevant for industries where replenishment cost of products with varying quality levels are heavily influenced by the price of an underlying component, such as the iron ore commodity price in steel manufacturing. (For other examples on how price fluctuations in commodity markets impact inventory policy, see Berling and Martinez-de-Albeniz, 2011.) The following results summarize our findings.

THEOREM 6. (Sensitivity to Replenishment Cost Parameters)

(a) *Suppose the current-period replenishment cost for product type- i , c_i^t , increases. Then, the optimal list price for product type- i increases whereas the list price for product type- j ($j \neq i$) does not change. The optimal base-stock level for product type- i decreases while the optimal base-stock level for product type- j increases. Further, the expected number of customers subsequently receiving upgrades decreases with c_1^t and increases with c_2^t .*

(b) *Let $c_1^t = \gamma c^t$ and $c_2^t = c^t$, where $\gamma > 1$. Suppose the current period underlying cost c^t increases. Then, the optimal list price for both products increase where the increase in the list price for product type-1 is larger than that for product type-2. The optimal base-stock level for product type-1 decreases. For product type-2, there exists a critical replenishment cost differential threshold γ_r^t , such that the base-stock level for product type-2 decreases if $\gamma < \gamma_r^t$ and increases if $\gamma > \gamma_r^t$. Further, there also exists a critical cost differential threshold γ_u^t for upgrades, such that the expected number of customers subsequently receiving upgrades decreases with c^t if $\gamma > \gamma_u^t$.*

Theorem 6 part (a) shows that the firm applies a higher price for a product when that product's replenishment cost increases, which in turn, decreases the expected demand for the product and increases the expected demand for the other product. Consequently, the firm selects a lower base-stock level for the product experiencing the cost increase and increases the base-stock level for the other product. These optimal pricing and replenishment decisions lead the firm to offer fewer subsequent customer upgrades if the replenishment cost for the higher quality product increases and to offer more upgrades if the replenishment cost of the lower quality product increases. That is, when the higher quality product's cost increases, decreased availability of the higher quality

product accompanied by an increased availability of the lower quality product necessitates fewer upgrades.

The results for the correlated replenishment cost structure we outline in Theorem 6 part (b) show that as an underlying cost increases, the optimal list price charged for both products increase. We also find that the optimal base-stock level for the higher quality product decreases. When the cost differential between products is low, the base-stock level for the lower quality product decreases as well. On the other hand, when the cost differential is substantially high, the firm prefers to rely less on upgrades and thus increases the base-stock level for the lower quality product.

As the preceding two results in this section correspond to the list price region, we also would like to comment on how optimal prices are influenced by the quality differentiation and replenishment costs beyond this region. Through numerical tests, and as an example, we observe that an increase in the quality of the lower quality product continues to result in an increase in its own price beyond the list price region as well. Even though the price of the higher quality product remained constant in the list price region, we observe that it weakly decreases outside this region. Regarding sensitivity with respect to the replenishment costs, numerical studies indicate that both prices weakly increase beyond the list price region when (1) the replenishment costs are independent and either one of them increases, or (2) the replenishment costs are correlated and an underlying replenishment cost driving the replenishment costs for both products increases.

6. Numerical Study

In the following, we first numerically demonstrate the value obtained by being able to offer product upgrades instead of just being able to adjust prices. We then look into the impact of demand correlation between the higher and lower quality products. Finally, we describe a heuristic policy that can be computed efficiently and test its performance against the optimal policy.

6.1. Value of Upgrades

As a follow up to our findings regarding the impact of upgrades on the optimal policy, we now evaluate the implications of upgrades on a firm's profit. To do so, we numerically compute the optimal profit obtained with and without upgrades and compare the two profit values to evaluate the profit improvement through offering upgrades. We represent our findings utilizing a problem in which the quality levels for the higher and lower quality products are $q_1 = 1$ and $q_2 = 0.7$. We normalize the quality and price pairs for the outside options for a simpler representation of the demand-price model similar to the one considered by Mantin et al. (2014). Specifically, we let $\underline{q} = 0$, $\underline{p} = 0$, and assume $\bar{q}$ and $\bar{p}$ large such that $\bar{q} \gg q_1$ and $\bar{p} \gg p_1$ for any possible p_1 and that

$\bar{p}/\bar{q} = \bar{v}$. We let the valuation upper bound $\bar{v} = 100$ and the distribution density $\delta = 0.2$. Regarding the cost parameters, we assume replenishment costs as $c_1 = 40$ and $c_2 = 20$, and the holding and backorder costs as $h_i^+ = 0.5c_i$ and $h_i^- = 4h_i^+$ for $i = 1, 2$. We set replenishment capacities at $K_1 = 5$ and $K_2 = 6$. We set the demand uncertainty to be uniformly distributed in the range $[-\epsilon, \epsilon]$ with $\epsilon = 2$. (As a side note, these parameters lead to list prices of 70 and 45 for the higher and lower quality products, corresponding to list price demands to be uniformly distributed on $[1.3, 5.3]$ and $[1.8, 5.8]$, respectively.) Finally, we assume a discount rate $\beta = 0.8$ and set the horizon length as $T = 5$.

In the following, we first vary the holding and shortage costs and the uncertainty range in order to gain insights into how these parameters influence the additional profit gained through upgrades. Then, we will comment on a secondary set of parameters where the products have closer quality levels and replenishment costs. For each parameter set, we compute the optimal profit starting from a range of 25 initial inventory positions corresponding to $x_1 = \{-2, -1, 0, 1, 2\}$ and $x_2 = \{-2, -1, 0, 1, 2\}$. The profit values reported in Table 1 for each parameter set correspond to the average profit across the range of the starting inventory positions. Table 1 indicates that the benefit of being able to offer upgrades can be substantial. The parameters we have tested for these particular instances resulted in profit improvements between 2.9% and 17.2%. Cases 1, 2, and 3 indicate that the value of upgrades increases when an upgrade helps the firm avoid incurring expensive holding cost for the higher quality item. Comparing cases 4, 1 and 5 indicates that as the holding cost for the lower quality item increases, the value of being able to offer upgrades also increases. This is due to the fact that when faced with higher holding cost, a firm that can offer upgrades can reduce the replenishment level for the lower quality product and rely more

Case	h_1^+	h_2^+	h_1^-	h_2^-	ϵ	Optimal profit without upgrades	Optimal profit with upgrades	% Profit improvement
1	20	10	80	40	2	412.6	440.5	6.8
2	10	10	80	40	2	453.0	482.0	6.4
3	30	10	80	40	2	372.4	420.6	13.5
4	20	5	80	40	2	432.7	449.7	4.3
5	20	15	80	40	2	392.7	428.4	9.4
6	20	10	40	40	2	454.1	467.7	2.9
7	20	10	120	40	2	385.8	428.7	11.4
8	20	10	80	20	2	433.5	462.8	7.4
9	20	10	80	50	2	405.7	428.8	5.7
10	20	10	80	60	2	399.0	424.2	6.3
11	20	10	80	40	1	513.9	538.2	4.7
12	20	10	80	40	3	298.6	350.1	17.2

Table 1 Value of upgrades across various holding and shortage costs and uncertainty levels

on upgrades. In a similar fashion, cases 6, 1, and 7 demonstrate that a higher shortage cost for the higher quality product also improves the relative value of upgrades as it prompts the firm to increase the replenishment level for the higher quality product and in turn increase its availability for a possible upgrade. Our numerical test also indicates that the value of upgrades may increase or decrease with the shortage cost for the lower quality product. Lastly, as can be expected and observed in Cases 11, 1, and 12, the relative value of upgrades increases when demand uncertainty is higher. The parameter sets we test where the products' quality levels and replenishment costs are closer to each other also resulted in a significant benefit through upgrades. For example, upgrades for an instance with $q_1 = 1$, $q_2 = 0.9$, $c_1 = 30$, and $c_2 = 24$ while all other parameters remain as previously set raises the optimal profit from 620.4 to 667.4, thus improving the profit by 7.6%. This is due to the fact that when the two products cost roughly the same, upgrading customers from one to the other does not result in a significant cost to the firm while giving it significant flexibility.

6.2. Impact of Demand Correlation

Various circumstances may lead to correlations between the demand for the higher quality product and the demand for the lower quality product. For example, a rapid economic growth (or downturn) may lead to an increase (or decrease) in the demand for both high quality and low quality steel, indicating positive demand correlation. On the other hand, items that exhibit a higher degree of substitutability may also likely result in negatively correlated demand uncertainties. Hence, it is important to highlight how potential demand correlations impact the firm's decisions and the benefits it receives from offering upgrades.

To examine the impact of correlation, we consider five problem instances where the correlation coefficient between the demand for the two different product types, denoted by ρ , ranges from -1 to 1 with increments of 0.5 and solve the stochastic dynamic program described in (3) - (4) for each of these instances. We then record the optimal pricing, replenishment, and upgrade decisions, as well as the optimal value function over the planning horizon. To gain further insights, we also generate and examine sample paths that follow the corresponding recorded optimal policy. This numerical study allows us to explore whether the structure of the optimal policies prevail under various correlation instances. In addition, it also provides insights into how demand correlation impacts the value of upgrades. To shed light on the optimal policy structure in the presence of demand correlation, we consider a problem instance similar to the once described previously with parameters set as $q_1 = 1$, $q_2 = 0.6$, $\bar{v} = 100$, $\delta = 0.2$, $c_1 = 50$, $c_2 = 20$, $h_i^+ = 0.25c_i$ and $h_i^- = 4h_i^+$ for $i = 1, 2$, $K_1 = 4$ and $K_2 = 5$. We again set the demand uncertainty to be uniformly distributed in the




Figure 2 Histograms for the proportion of demand for the lower quality product that is met through upgrades across the planning horizon for instances where the demands for the two types of products are (a) perfectly positively correlated, $\rho = 1$, (b) moderately positively correlated, $\rho = 0.5$, (c) independent, (d) moderately negatively correlated, $\rho = -0.5$, and (e) perfectly negatively correlated, $\rho = -1$.

range $[-2, 2]$, assume a discount rate $\beta = 0.8$, and set the horizon length as $T = 5$. After recording the optimal policy over the 5-period horizon, and in order to address the impact of correlation on upgrades, we also run 10,000 sample paths for each problem instance starting with the initial state $(0, 0)$ and following the optimal policy thereafter.

Our numerical study suggests that the optimal policy structure for the correlated demand instances remain similar to the optimal policy structure derived for the independent demand instance. That is, the optimal pricing policy is defined by regions of list prices, price discounts, and price surcharges, the replenishment policy is of a modified base-stock type, and the upgrade policy follows a protection threshold on the higher quality product with similar monotonicities. (As an example, a visualization indicating the similarity of the pricing policy for the higher and lower quality products across different demand correlations can be found in the Online Appendix.)

Even though the policy structure remains similar, the nature of the demand correlation greatly influences how often the firm resorts to upgrades. Figure 2 displays the relative frequencies of the proportion of demand for the lower quality product that is met through upgrades across the planning horizon for five instances where demand uncertainties are (a) perfectly positively correlated, $\rho = 1$, (b) moderately positively correlated, $\rho = 0.5$, (c) independent, (d) moderately negatively correlated, $\rho = -0.5$, and (e) perfectly negatively correlated, $\rho = -1$. Whereas virtually none of the customers for the lower quality product receives an upgrade offer in the perfectly positive correlated case, an average of 6.7% of the overall demand for the lower quality product across the horizon receive an upgrade offer for the independent demand case and 11.9% receive upgrades if demands are perfectly negatively correlated. The direction of this finding is intuitive as upgrades are less likely to be utilized when both products have sufficient inventory or when both products face shortages, and more likely to be utilized when there is an inventory imbalance

between the products with the higher quality product having excess inventory and the low quality product facing a shortage.

Consequently, we find that the value of upgrades also increases when demands are negatively correlated. For example, based on the above parameters, when demand between the products are perfectly positively correlated, the optimal profit with and without upgrades are both 323.7, indicating no profit improvements due to upgrades being allowed. On the other extreme, when product demands are perfectly negatively correlated, upgrades improve the profits from 324.5 to 339.0, i.e., by 4.5%.

6.3. A Heuristic Policy

As computing the optimal upgrade, replenishment and pricing policies require solving a two-stage, multi-period dynamic program, and is rather computationally expensive, we would like to consider a heuristic policy that utilizes insights gained through the optimal policy characterization. The starting point for the heuristic will be a reduced problem that mimics an unlimited capacity, stationary version of the problem to determine an initial set of base-stock levels for the products by taking upgrading into account. We then consider capacity limitations and make adjustments to the replenishment levels and prices accordingly.

Before we begin, we would like to comment on an observation for the stationary setting with no capacity limitations. Whenever both items need replenishment, the firm brings inventories back to base-stock levels and charges list prices. If the initial inventory of either product is higher than its base-stock level, it will take a number of transient periods until both inventories fall below their base-stock levels and the firm will apply this list price, base-stock policy thereafter. (See for example, Zhu and Thonemann, 2009). Therefore, in Step 1 of the heuristic, we determine the initial base-stock levels for infinite capacity based on the list prices p_1° and p_2° as given in Theorem 3 and let d_1° and d_2° represent the corresponding expected demand for the products. As a reminder for the notation, we let $f_1(\epsilon_1)$ and $f_2(\epsilon_2)$ denote zero-mean, independent probability density functions for the stochastic terms ϵ_1 and ϵ_2 in the price-demand relationships, with $(\epsilon_1, \epsilon_2) \in \mathcal{E}^2$ where $\mathcal{E}^2 = [\underline{\epsilon}_1, \bar{\epsilon}_1] \times [\underline{\epsilon}_2, \bar{\epsilon}_2]$. We also let $F_1(\epsilon_1)$ and $F_2(\epsilon_2)$ denote the corresponding cumulative distribution functions. We describe the heuristic for instances with $h_1^+ + h_2^- \geq \beta(c_1 - c_2)$ first, followed by a description on modifications for instances with $h_1^+ + h_2^- < \beta(c_1 - c_2)$.

Step 1. We obtain an initial set of base-stock levels $\hat{y}_1^\circ$ and $\hat{y}_2^\circ$ considering a single-period reduced problem in which a firm with no capacity restrictions and no initial inventory determines optimal base-stock levels for two products with expected demands d_1° and d_2° by taking upgrading into account. Specifically, we consider the problem of minimizing a single-period expected cost function

$C(y_1, y_2)$ that consists of replenishment costs c_i per unit of product type- i , holding and shortage costs h_i^+ and h_i^- after demand realization and any subsequent upgrades, and a discounted cost βc_i for any negative inventory (imitating the replacement cost to return to the original zero inventory position) or a reward $-\beta c_i$ for any remaining positive inventory for product type- i , $i = 1, 2$. For brevity, we relegate the explicit representation of the objective function $C(y_1, y_2)$ to the Online Appendix. The optimal base-stock levels for this reduced problem are obtained by simultaneously solving for $\hat{y}_1^\circ$ and $\hat{y}_2^\circ$ in the following:

$$\begin{aligned} F_1(\hat{y}_1^\circ - d_1^\circ) &= \frac{h_1^- - (1-\beta)c_1 + (h_1^+ + h_2^- - \beta(c_1 - c_2)) \int_{\epsilon_1}^{\hat{y}_1^\circ - d_1^\circ} \left(1 - F_2(\hat{y}_1^\circ + \hat{y}_2^\circ - d_1^\circ - d_2^\circ - \epsilon_1)\right) f_1(\epsilon_1) d\epsilon_1}{h_1^+ + h_1^-} \\ F_2(\hat{y}_2^\circ - d_2^\circ) &= \frac{h_2^- - (1-\beta)c_2 + (h_1^+ + h_2^- - \beta(c_1 - c_2)) \int_{\epsilon_1}^{\hat{y}_1^\circ - d_1^\circ} \left(F_2(\hat{y}_1^\circ + \hat{y}_2^\circ - d_1^\circ - d_2^\circ - \epsilon_1) - F_2(\hat{y}_2^\circ - d_2^\circ)\right) f_1(\epsilon_1) d\epsilon_1}{h_2^+ + h_2^-} \end{aligned} \quad (5)$$

Step 2. (a) Next, we determine adjusted base-stock levels $\hat{y}_1(x_2)$ and $\hat{y}_2(x_1)$ as well as price surcharges considering capacity limitations. For ease of implementation, we do not consider price discounts or base-stock adjustments for excess inventory as it would eventually be drawn and remain below base-stock levels after a number of transient periods. Hence, we set $\hat{y}_1(x_2) = \hat{y}_1^\circ$ if $x_2 \geq \hat{y}_2^\circ - K_2$ and $\hat{y}_2(x_1) = \hat{y}_2^\circ$ if $x_1 \geq \hat{y}_1^\circ - K_1$. When $x_2 < \hat{y}_2^\circ - K_2$, the lower quality product can only be replenished up to $x_2 + K_2$, and the corresponding state dependent base-stock for the higher quality product is determined by solving for $\hat{y}_1(x_2)$ below:

$$F_1(\hat{y}_1(x_2) - d_1^\circ) = \frac{h_1^- - (1-\beta)c_1 + (h_1^+ + h_2^- - \beta(c_1 - c_2)) \int_{\epsilon_1}^{\hat{y}_1(x_2) - d_1^\circ} \left(1 - F_2(\hat{y}_1(x_2) + x_2 + K_2 - d_1^\circ - d_2^\circ - \epsilon_1)\right) f_1(\epsilon_1) d\epsilon_1}{h_1^+ + h_1^-}$$

Similarly, for $x_1 < \hat{y}_1^\circ - K_1$, $\hat{y}_2(x_1)$ is obtained by solving the following:

$$F_2(\hat{y}_2(x_1) - d_2^\circ) = \frac{h_2^- - (1-\beta)c_2 + (h_1^+ + h_2^- - \beta(c_1 - c_2)) \int_{\epsilon_1}^{x_1 + K_1 - d_1^\circ} \left(F_2(x_1 + K_1 + \hat{y}_2(x_1) - d_1^\circ - d_2^\circ - \epsilon_1) - F_2(\hat{y}_2(x_1) - d_2^\circ)\right) f_1(\epsilon_1) d\epsilon_1}{h_2^+ + h_2^-}$$

To summarize, the replenishment level set by the heuristic for the higher and lower quality products are $y_1 = x_1 + \min((\hat{y}_1(x_2) - x_1)^+, K_1)$ and $y_2 = x_2 + \min((\hat{y}_2(x_1) - x_2)^+, K_2)$, respectively.

(b) We now determine price surcharges for the products if their initial inventory cannot be brought up to their respective base-stock levels. Charging list prices allows the firm to retain the highest revenue at the expense of a potentially less favorable inventory position resulting in an increased cost. On the other hand, applying a price surplus reduces expected revenue while

Case	h_1^+	h_2^+	h_1^-	h_2^-	ϵ	Optimal Profit	Heuristic Profit	% difference
1	20	10	80	40	1	538.2	536.0	0.4
2	10	10	80	40	1	564.4	560.5	0.7
3	30	10	80	40	1	522.5	521.3	0.2
4	20	5	80	40	1	543.8	541.6	0.4
5	20	15	80	40	1	532.6	529.7	0.5
6	20	10	40	40	1	549.4	541.2	1.5
7	20	10	120	40	1	538.2	532.3	1.1
8	20	10	80	20	1	552.5	549.8	0.5
9	20	10	80	60	1	530.7	528.0	0.5
10	20	10	80	40	2	440.5	432.2	1.9
11	10	10	80	40	2	482.2	479.2	0.6
12	30	10	80	40	2	422.6	416.1	1.5
13	20	5	80	40	2	451.1	440.9	2.3
14	20	15	80	40	2	429.8	423.0	1.6
15	20	10	40	40	2	467.4	465.6	0.4
16	20	10	120	40	2	429.6	426.7	0.7
17	20	10	80	20	2	465.4	457.4	1.7
18	20	10	80	60	2	424.2	416.9	1.7

Table 2 Performance of the heuristic policy across varying holding and shortage costs and uncertainty levels

facilitating a reduction in overall cost through a more favorable inventory position. We set the adjusted price to the price level that would balance revenue loss with cost improvement.

Let $\Delta d_1(p_1, p_2) := d_1(p_1, p_2) - d_1^\circ$ and $\Delta d_2(p_1, p_2) := d_2(p_1, p_2) - d_2^\circ$ denote, respectively, the change in expected demand for the higher and lower quality product if the firm applies prices (p_1, p_2) . Similarly, $\Delta R(p_1, p_2) := R(p_1, p_2) - R(p_1^\circ, p_2^\circ)$ denotes the change in revenue. Further, let $\Delta C(p_1, p_2) := C(y_1 - \Delta d_1(p_1, p_2), y_2 - \Delta d_2(p_1, p_2)) - C(y_1, y_2) + c_1 \Delta d_1(p_1, p_2) + c_2 \Delta d_2(p_1, p_2)$ denote the change in overall cost due to the particular price selection. (Note: The terms $c_i \Delta d_i(p_1, p_2)$ are corrections to the cost function reflecting the change in inventory position is due to demand suppression and not due to additional replenishment.) The adjusted prices (p_1, p_2) are determined by solving $\max_{p_1, p_2} \Delta R(p_1, p_2) - \Delta C(p_1, p_2)$.

Next, we test the performance of the heuristic policy. As in our initial numerical study, we set $q_1 = 1$, $q_2 = 0.7$, $\bar{v} = 100$, $\delta = 0.2$, $c_1 = 40$, $c_2 = 20$, $h_i^+ = 0.5c_i$ and $h_i^- = 4h_i^+$ for $i = 1, 2$, $K_1 = 5$, $K_2 = 6$, $\beta = 0.8$, and $T = 5$, and systematically increase and decrease various problem parameters and comprehensively test for a lower and higher demand uncertainty setting across the range of initial starting inventories as described in Section 6.1.

Table 2 reports the optimal profit along with the profit achieved by the heuristic policy and its percent difference from the optimal profit. The average difference between the profit obtained by the optimal policy and the profit obtained by the heuristics across all instances is 1.0% with a maximum optimality gap of 2.3%. Thus, we observe that the heuristic policy performs well

compared to the optimal policy across a wide range of parameter values.

Before we conclude this section, we would like to add a final note on instances where $h_1^+ + h_2^- < \beta(c_1 - c_2)$. As we mentioned earlier, this case does not result in upgrades in an infinite capacity setting. With capacity limitations however, upgrades may still be economical for a sufficiently high level of shortage for the lower quality product. We handle such situations as follows. First, we again start by finding initial base-stock levels $\hat{y}_1^o$ and $\hat{y}_2^o$ corresponding to an infinite capacity setting that minimizes overall cost, which is now separable in each product as no upgrades are offered in the infinite capacity setting. That is, $\hat{y}_i^o$ satisfies: $F_i(\hat{y}_i^o - d_i^o) = \frac{h_i^- - (1-\beta)c_i}{h_i^+ + h_i^-}$ for $i=1,2$. We then identify the smallest N such that $h_1^+ + Nh_2^- > \beta(c_1 - c_2)$, where $N - 1$ can be interpreted as an approximation for the number of periods the firm would be willing to carry a demand shortage without offering upgrades. We can then follow Step 2 analogously to determine adjustments to base-stock levels and prices for any $x_2 < \hat{y}_2^o - NK_2$. As one example, consider an instance with $h_i^+ = 0.1c_i$ and $h_i^- = 4h_i^+$ with all remaining parameters as defined earlier, which does not lead to immediate upgrades. The optimal profit for this instance is 595.8 while the modified heuristic policy results in a profit of 593.4, with an optimality difference of 0.4%.

7. An Extension to Incorporate Upgrade Fees

We now would like to extend our analysis to a setting where the firm sets an upgrade fee p_u^t at each period, resulting in only a proportion $\pi(p_u^t)$ of the customers being interested in paying the price differential to receive the upgrade. This extension may be considered to be close to consumer-driven substitution as it is the consumer that decides on whether to upgrade based on the upgrade fee announced by the firm. The model also retains the features of a firm-driven substitution as well in the sense that the firm decides on the expected number of upgrades by adjusting the upgrade fee. We make two assumptions for tractability. First, we assume that customers arriving in each period are myopic in the sense that they do not take into account the possibility of being offered an upgrade while they are making their initial purchase decisions. This can be considered reasonable as the upgrade fee is not constant but is chosen by the firm in each period after the customers had made their purchase decisions, making it difficult for a customer to compare their payoffs and act strategically. Second, we assume that the proportion of customers who accept an upgrade offer is a decreasing linear function of the upgrade fee $p_u^t \in [\underline{p}_u, \bar{p}_u]$, where an upgrade fee of $\underline{p}_u$ results in all customers accepting the offer, and an upgrade fee of $\bar{p}_u$ discourages all customers, leading to no upgrades. Thus, $\pi(p_u^t) = \frac{\bar{p}_u - p_u^t}{\bar{p}_u - \underline{p}_u}$. (In fact, this is the form upgrade probability takes when one considers the valuations for customers who initially prefer the lower quality product, and

for those who prefer to pay the upgrade fee to receive the higher quality product. For example, if the firm charges list prices p_1^o, p_2^o for the products, one can show that $\bar{p}_u = p_2^o - p_1^o$ and $\underline{p}_u = (p_2^o - \underline{p})(q_1 - q_2)/(q_2 - \underline{q})$.

Overall, the pool of potential customers eligible for upgrades in each period consists of the realized demand for the lower quality product, D_2^t . Thus by selecting an upgrade fee p_u^t , the firm expects $u^t = \pi(p_u^t)D_2^t$ customers to upgrade. We further incorporate an additive uncertainty term on the number of customers that are willing to pay for an upgrade and let $u^t + \zeta^t$ express the actual number of people who upgrade where ζ^t is a zero-mean random variable with a probability density function $f(\zeta^t)$. (To aid our analysis, we assume D_2^t in any period is positive, and that ζ^t is bounded on the interval $[-u^t, D_2^t - u^t]$, resulting in actual upgrades to be distributed on $[0, D_2^t]$. For continuity at the boundaries, we also assume $f(\zeta^t)$ has zero density at the boundaries and that the uncertainty vanishes as $\pi(p_u^t)$ approaches zero or one.) Similar to our original formulation, we will describe the problem in terms of the target upgrade quantity, u^t , as the decision variable. For any target upgrade quantity u^t effectively set by the firm, we can compute the corresponding upgrade fee through $p_u^t(u^t) = \bar{p}_u - \frac{u^t}{D_2^t}(\bar{p}_u - \underline{p}_u)$. We now provide the revised formulation:

Stage 1:

$$V^t(x_1^t, x_2^t) = \max_{\substack{y_i^t, p_i^t \\ x_i^t \leq y_i^t \leq x_i^t + K_i}} R(p_1^t, p_2^t) - c_1^t \cdot (y_1^t - x_1^t) - c_2^t \cdot (y_2^t - x_2^t) + E_{D_1^t, D_2^t}[G^t(y_1^t - D_1^t, y_2^t - D_2^t, D_2^t)] \quad (6)$$

Stage 2:

$$G^t(w_1^t, w_2^t, D_2^t) = \max_{\substack{u^t \\ 0 \leq u^t \leq D_2^t}} E_{\zeta^t} \left[\underbrace{\left(\bar{p}_u - \frac{u^t}{D_2^t}(\bar{p}_u - \underline{p}_u) \right) (u^t + \zeta^t)}_{\text{Upgrade Revenue}} - h_1(w_1^t - u^t - \zeta^t) - h_2(w_2^t + u^t + \zeta^t) + \beta V^{t-1}(w_1^t - u^t - \zeta^t, w_2^t + u^t + \zeta^t) \right] \quad (7)$$

Compared to our original formulation given by (3) and (4), the second stage problem described in (7) now includes an additional term for the revenue generated by upgrades, where $\bar{p}_u - \frac{u^t}{D_2^t}(\bar{p}_u - \underline{p}_u)$ corresponds to the upgrade fee and $u^t + \zeta^t$ is the upgrade quantity. Notice that, along with the intermediate inventory levels, we are now also passing the information on the realization of the demand for the lower quality product as the potential pool of customers eligible for an upgrade.

The result below describes the optimal upgrade policy in terms of the upgrade fee charged by the firm.

THEOREM 7. *The optimal upgrade fee, $p_u^{*t}(w_1^t, w_2^t, D_2^t) \in [\underline{p}_u, \bar{p}_u]$, decreases with w_1^t and increases with w_2^t and D_2^t .*

As indicated by Theorem 7, the optimal upgrade fee charged by the firm decreases with the intermediate inventory level for the higher quality product and increases with the intermediate inventory level for the lower quality product. In other words, when the intermediate inventory level for the higher quality product increases, the firm reduces the upgrade fee to encourage more customers to upgrade. When fewer higher quality products are available, the firm charges more for the upgrades. Similarly, when the intermediate inventory level for the lower quality product increases (i.e., fewer shortages), the firm does not need to offer as many upgrades and thus increases the upgrade price. Likewise, a lower intermediate inventory level for the lower quality product (more shortages) prompts the firm to encourage more customers to upgrade through a lower upgrade fee in order to reduce backlogs for the product. (Note that the total fee an upgrading customer pays, i.e., the price for the lower quality product plus the upgrade fee, may not necessarily be less than the price of the higher quality product on every sample path.) In addition, as in our original model, a particular target for the upgrade quantity corresponds to a target protection limit on the higher quality product. Theorem 7 also implies that this target protection level on the higher quality product is now a function of both intermediate inventory levels individually. (We formally show this in the proof of Theorem 7.) It is important to note that this result is different from our earlier findings for the original problem which showed that the protection limit on the higher quality product is a function of the intermediate inventory levels only through their sum and not individually.

For the first stage decisions, we find that the partially decoupled state-dependent base-stock policy structure of the original model remains optimal in this setting as well. Regarding pricing, as opposed to the original model, the list prices p_1^{ot} and p_2^{ot} in this modified setting no longer have closed form solutions. In addition, even though the general structure of the optimal pricing policy for the higher quality product prevails, we find that the list price region for the lower quality product may narrow. For example, consider the initial inventory states where the higher quality product has excess inventory while the lower quality product is understocked yet it can be replenished to its base stock level. As opposed to applying the list price as in the original model, the firm may now give a price discount to the lower quality product as well. Knowing that it can encourage some customers to upgrade, this enables the firm to capture more demand overall, and use revenue generating upgrades to help bring down excess inventory for the higher quality product. For brevity, we relegate stating the formal results to the Online Appendix.

8. Conclusions

Our focus in this paper has been the joint implementation of dynamic pricing and customer upgrades as means of price- and availability-based product substitutions to better match demand and constrained supply across vertically differentiated products. Specifically, we study a multi-period model where the firm first sets prices and replenishment levels for each product while demand is still uncertain, and after observing the demand, decides whether it should offer any customers an upgrade to a higher quality product. We find that the optimal upgrade policy is defined by a protection level on the higher quality item, the optimal replenishment follows a partially decoupled, modified base-stock policy, and the pricing policy consists of various regions of price surcharges, list prices, and price discounts based on the initial inventory positions of the products. We then focus on the impact of offering product upgrades on the choice of prices and replenishment levels and show that offering upgrades enables the firm to select prices that better protect the list-price differentiation between the products. We also investigate how the optimal policy is influenced by the quality differentiation between the products and replenishment costs. We show that an increase in the quality difference between the products prompts the firm to select list prices that are further apart and to increase the base-stock level for the higher quality product while decreasing the base-stock level for the lower quality product. We also find that the firm relies less on upgrades when the replenishment cost for the higher quality product increases or the replenishment cost for the lower quality product decreases. The insights gained through the characterization of the optimal policy structure further allows us to construct an easily implementable heuristic policy that performs well compared to the optimal policy across various parameter values. Finally, studying an extension where the firm can also dynamically set an upgrade fee, we find that the firm charges more for an upgrade if the availability of the higher quality product is lower, and charges less for an upgrade when the number of customers who were unable to get the lower quality product is larger. We believe our results overall further our understanding of the intricate relationship among a firm's decisions on pricing, replenishment, and product upgrades in an effort to better match demand and constrained supply.

References

- Aydin, G., S. Ziya. 2008. Pricing Promotional Products Under Upselling. *Manufacturing & Service Operations Management*, 10 (3), 360-376.
- Bassok, Y., R. Anupindi, R. Akella. 1999. Single-Period Multiproduct Inventory Models with Substitution. *Operations Research*, 47 (4), 632-642.

- Berling, P., V. Martinez-de-Albeniz. 2011. Optimal Inventory Policies when Purchase Price and Demand Are Stochastic. *Operations Research*, 59 (1), 109-124.
- Bitran, R.G., R. Caldentey. 2003. Pricing Models for Revenue Management. *Manufacturing & Service Operations Management*, 5 (3), 203-229.
- Bresnahan, T.F. 1981. Departures from Marginal Cost Pricing in the American Automobile Industry. *Journal of Econometrics*, 17, 201-227.
- Ceryan, O., O. Sahin, I. Duenyas. 2013. Dynamic Pricing of Substitutable Products in the Presence of Capacity Flexibility. *Manufacturing & Service Operations Management*, 15 (1), 86-101.
- Chan, L.M.A., Z.J.M. Shen, D. Simchi-Levi, J.L. Swann. 2004. Coordination of Pricing and Inventory Decisions: A Survey and Classification. D. Simchi-Levi, D. S. Wu, Z. J. Shen, eds. *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, Kluwer Academic Publisher, Boston, 335-392.
- Chen, X., D. Simchi-Levi. 2010. Coordinated Pricing and Inventory Management. O. Ozer and R. Phillips, eds. *Handbook of Pricing Management*, Oxford University Press, Oxford, 784822.
- Cui, Y., I. Duenyas, O. Sahin. 2017. Pricing of Conditional Upgrades in the Presence of Strategic Consumers, *Management Science*, Articles in Advance.
- Elmaghraby, W., P. Keskinocak. 2002. Dynamic Pricing: Research Overview, Current Practices and Future Directions. *Management Science*, 49 (10), 1287-1309.
- Federgruen, A., A. Heching. 1999. Combined Pricing and Inventory Control under Uncertainty. *Operations Research*, 47 (3), 454-475.
- Friedman, D.A. 2013. Explaining "Bait-and-Switch" Regulation. *William & Mary Business Law Review*, 4 (2), 575-638.
- Gallego, G., C. Stefanescu. 2009. Upgrades, Upsells and Pricing in Revenue Management. *Working Paper*, Columbia University.
- IMF Commodities Team. February 13, 2013. *International Monetary Fund Commodity Market Monthly*, Washington, D.C.
- Karakul, M., L.M.A. Chan. 2008. Analytical and Managerial Implications of Integrating Product Substitutability in the Joint Pricing and Procurement Problem. *European Journal of Operational Research*, 190, 179-204.
- Mahajan, S., G. van Ryzin. 2001. Stocking Retail Assortments Under Dynamic Consumer Substitution. *Operations Research*, 49 (3), 334-351.

- Mantin, B., H. Krishnan, T. Dhar. 2014. The Strategic Role of Third-Party Marketplaces in Retailing. *Production and Operations Management*, 23 (11), 1937-1949.
- McKinsey&Company. 2013. Scarcity and Saturation. Steel and the challenges of volatile raw materials, flat margins, and overcapacities *McKinsey on Metals & Mining*, 8, 5-39.
- Mussa, M., S. Rosen. 1978. Monopoly and Product Quality. *Journal of Economic Theory*, 18 (2), 301-317.
- Nagarajan, M., S. Rajagopalan. 2008. Inventory Models for Substitutable Products: Optimal Policies and Heuristics. *Management Science*, 54 (8), 1453-1466.
- Netessine, S., G. Dobson, R.A. Shumsky. 2002. Flexible Service Capacity: Optimal Investment and the Impact of Demand Correlation. *Operations Research*, 50 (2), 375-388.
- Pasternack, B. A., Z. Drezner. 1991. Optimal Inventory Policies for Substitutable Commodities with Stochastic Demand. *Naval Research Logistics*, 38, 221-240.
- Sethuraman R., V. Srinivasan, D. Kim. 1999. Asymmetric and Neighborhood Cross-Price Effects: Some Empirical Generalizations. *Marketing Science*, 18 (1), 23-41.
- Shumsky, R.A., F. Zhang. 2009. Dynamic Capacity Management with Substitution. *Operations Research*, 57 (3), 671-684.
- Tomlin, B., Y. Wang. 2008. Pricing and Operational Recourse in Coproduction Systems. *Management Science*, 54 (3), 522-537.
- Transchel, S. 2011. Joint Pricing and Inventory Decisions Under Stockout-Based and Price-Dependent Substitution. *Working Paper*, The Pennsylvania State University.
- Yu, Y., X. Chen, F. Zhang. 2015. Dynamic Capacity Management with General Upgrading. *Operations Research*, 63 (6), 1372-1389.
- Zhu, K., U. W. Thonemann. 2009. Coordination of Pricing and Inventory Control Across Products. *Naval Research Logistics*, 56 (2), 175-190.