POMS PRODUCTION AND OPERATIONS MANAGEMENT

PRODUCTION AND OPERATIONS MANAGEMENT POMS

Vol. 31, No. 3, March 2022, pp. 941–962
ISSN 1059-1478 | EISSN 1937-5956 | 22 | 3103 | 0941
DOI 10.1111/poms.13589

Dynamic Pricing with Money-Back Guarantees

Yan Liu*
International Institute of Finance, School of Management, University of Science and Technology of China, Hefei, 230026, China, bizliuy@ustc.edu.cn

Ningyuan Chen
Rotman School of Management, University of Toronto, Toronto, Canada M5S 3E6, Canada, ningyuan.chen@utoronto.ca

Money-back guarantees (MBGs), which allow customers to return products that do not meet their expectations, are widely used in the retail industry. In this study, we study a retailer's MBG policy with dynamic pricing of limited inventory. A key decision for the retailer is to decide whether to offer MBGs. When the product can be returned instantly, we find that the optimal MBG policy is a simple threshold policy: given the inventory level, it is optimal to offer an MBG if and only if the remaining selling time is longer than a threshold. Moreover, the threshold is decreasing with the inventory level. We also address the problem of dynamic pricing with positive return times. Due to the complexity, we analyze the associated fluid model, which has an infinite number of constraints. We consider a series of relaxations that have a nested structure and use the Lagrangian approach to explicitly solve these relaxed problems. This allows us to develop an iterative approach that is guaranteed to solve the fluid model in finite iterations. Our numerical analysis shows that the deterministic solution is asymptotically optimal for the stochastic system.

Key words: revenue management; dynamic pricing; money-back guarantee; product return
History: Received: November 2019; Accepted: September 2021 by Dan Zhang, after 2 revisions.
*Corresponding author.

1. Introduction

According to the market research company Statista,[1] the value of the apparel market in the United States was $368 billion in 2019. Unfortunately, analysts estimate that nearly 20% of purchases were returned. A crucial reason for this high return rate is money-back guarantee (MBG) policies, under which customers can return products for any reason and get a full refund. For example, all top 10 online clothing retailers on the Top 500 List of the Internet Retailer website offer MBGs for clothing (McWilliams 2012). The return rate is approximately 9% for brick-and-mortar retailers but can be as high as 20% for online retailers (Orendorff 2019). Therefore, return management is extremely important for the retail industry. Meanwhile, due to seasonality and limited selling periods, the retail industry is well-known for using dynamic pricing to manage the sales process (Gallego and Van Ryzin 1994). This study aims to characterize firms' optimal dynamic pricing and return policies in managing the sale of a fixed initial capacity over a finite sales horizon.

Product returns are driven by a mismatch between consumers and the product, which can lead to consumers' dissatisfaction with the product. MBGs are particularly important for products sold over the Internet or through mail-order catalogs, where customers cannot enjoy the benefits of the traditional touch-and-feel shopping experience. A purchase with an MBG gives customers the best of both worlds: the benefit of enjoying a product with a good match and the flexibility to return should the consumer not like the product. Reciprocally, MBGs can benefit the retailer by allowing him to charge a high price for the right consumer, and he can resell the items that are returned by unmatched consumers. Nevertheless, a consumer's benefit is mitigated by the cost of returning the product. A rational consumer recognizes the possibility of a product misfit when considering purchasing the product. On the other hand, the retailer faces the risk of not being able to sell the (returned) products, which is a key characteristic of revenue management. Therefore, an important question for the retailer is whether to offer MBGs.

The major difficulty for the firm when adopting an MBG is that returned products blur stock information and complicate the decisions. For brick-and-mortar retailers, McWilliams (2012) finds that the minimal MBG duration for sampled products is 30 days. Online retailers impose much shorter return times; the main e-tailers (e.g., JD.com, Amazon) typically specify 7 days for returns. It takes time for the retailer

to deliver the product to the consumer and for the consumer to send it back in case of a misfit. In China, delivery times are about 1 day for deliveries within the same city and 3 days for nationwide deliveries. Therefore, the total return time for Internet or mail orders can be 2 weeks. The return time might be negligible for durable products but is important for perishable products with relatively short sales seasons (e.g., 2 to 3 months as in Gallego and Van Ryzin 1994). In addition to managing the return process, the retailer needs to decide the price for the product, which is typically done dynamically in practice.

In this study, we propose a dynamic programming model in which the seller jointly decides the MBG strategy (i.e., whether to offer an MBG) and the pricing policy. The seller has a limited inventory at the beginning of a short sales season. At the point of purchase, the buyer does not know whether the product fits well. If the product is not a good match and the seller offers an MBG, the buyer will return the product to the seller for a full refund. The consumer has heterogeneous willingness-to-pay (WTP); hence the purchase decision depends on both the price and the presence of MBG strategy. To shed managerial insights into the MBG strategy, we first consider the instant-return case where the misfit products are returned immediately. This can be viewed as an approximation of reality when the allowed return period is negligible compared to the sales horizon. For practical reasons, we also consider a number of benchmarks—including MBG only and MBG with switch-off policies—where the retailer has less flexibility in adjusting the MBG policy. To address positive return times, the model needs to track return information, leading to a high-dimensional state. For tractability, we study the corresponding fluid model and develop an iterative algorithm. Accordingly, this study makes the following contributions.

First, we fully characterize the optimal pricing and MBG policy for the instant-return model. Given the current stock, it is optimal for the seller to allow for free returns when the length of the remaining season is above a threshold. The threshold decreases with the stock level; the seller should shorten the MBG period with more inventory. In addition to the monotonicity of the marginal value of inventory with respect to the time-to-go and inventory (Gallego and Van Ryzin 1994), we identify a new monotonicity result: the marginal value of inventory increases in the fit probability. This means that the seller should extend the MBG period if the consumer's fit probability increases. Showing the monotonicity of marginal value of inventory with respect to a parameter is nonstandard in the revenue management literature. For example, Kuo et al. (2011) resort to explicit expression of the optimal solution and hence require a uniform WTP distribution. Our paper adopts a new approach that applies to a general WTP distribution.

Second, we formulate the deterministic problem with positive return times. Unlike the traditional fluid model in revenue management, the deterministic formulation involves the return process. Thus, the constraint of non-negative inventory has to be imposed at any time and the number of constraints is infinite. To solve this model, we sequentially and carefully select a subset of constraints and relax others. This gives us a sequence of optimization problems which are relaxations of the deterministic problem and have a nested structure. We use the Lagrangian approach to explicitly solve these relaxed problems and develop an iterative algorithm that is guaranteed to solve the fluid model in finite iterations. We are not aware of any previous papers in the literature that face a non-trivial fluid model as ours and the methodology can be used for dynamic pricing problems with a large number of inventory constraints.

The solution to the fluid model provides a number of useful managerial insights. When the initial stock level is high, free returns should not be allowed. When the initial stock level is low, the seller may need to divide the selling horizon into three episodes and implement different pricing/return policies. In the first episode, a relatively low price should be set and almost all inventory is sold. In the second episode, the seller should set a high price and rely on reselling the returned products from the first episode to earn profits. In the third episode approaching the end of the horizon, MBGs are no longer allowed and the seller clears all returned products. This structure makes the most use of the MBG policy and the initial inventory and performs very well in numerical examples.

The rest of the paper is organized as follows. Section 2 provides a brief review of related literature. Section 3 formulates the firm's dynamic pricing and MBG problem with instant returns. Section 4 studies the dynamic pricing and MBG problem with positive return times. We consider both the deterministic and stochastic versions. Section 5 numerically analyzes the value of MBG. Section 6 considers the extensions of the random return time and the restocking cost. Section 7 concludes and suggests future research directions. All proofs can be found in the appendix.

2. Literature Review

MBG policies have been widely studied in the marketing and economics literature. Mann and Wissink (1988) find that MBGs can reduce the risk of a mismatch between customer needs and the product. Moorthy and Srinivasan (1995) show that sellers can use MBGs to signal the high quality of their products. Davis et al. (1995) show that MBGs, even with full

refunds, can increase a seller's profit, as well as the social welfare, when the salvage value for the seller is higher than the transaction cost for the buyer to return the product. To mitigate the effect of moral hazard, Davis et al. (1995) suggest sellers charge restocking or handling fees. For capacity-constrained service providers, Xie and Gerstner (2007) find that MBGs (i.e., service cancelations) are profitable due to the multiple selling opportunities of the product. Instead of charging a hassle cost, the authors recommend that service providers decrease customers' cost of cancelation. In addition to consumers' return, Li and Liu (2021) study distribution channels where retailers can further return the unsold inventories to manufacturers.

The operations management community has increasingly recognized the importance of studying consumer returns. To facilitate the operation of a large catalog or Internet mail-order retailer, Mostard and Teunter (2006) develop a newsvendor problem where customers return unsatisfactory products and the firm can resell the returns. Su (2009) builds a stylized model to incorporate the market size uncertainty and the valuation uncertainty. He focuses on the impact of refund policies on supply chain performance and coordination issues. Shulman et al. (2009) investigate the pricing and restocking decisions of a seller offering two horizontally differentiated products where the seller has the option to provide information to customers to eliminate product misfit. Akcay et al. (2013) consider the optimal order quantity and return policy when the seller can resell the returned items as inferior products. Gao and Su (2016) consider the firm's information provision problem where the consumer strategically chooses between online and offline channels. Gallino and Moreno (2018) use a randomized field experiment to study how fit information (e.g., virtual fitting rooms) affects customers' online purchasing and return behavior. Sun and Tyagi (2020) explore the incentive to provide fit information to consumers to reduce their product-fit uncertainty in a distribution channel.

There is extensive literature on dynamic pricing with limited inventories and finite selling horizons, pioneered by Gallego and Van Ryzin (1994) and Bitran and Mondschein (1997). The main question in this literature is how the seller can maximize the expected revenue by adjusting the product price based on the remaining time of the selling season and the on-hand inventory level. See Bitran and Caldentey (2003) and Elmaghraby and Keskinocak (2003) for reviews of early studies. Shen and Su (2009) provide an excellent review on customer behavior in revenue management, including consumers' valuation uncertainty in our model. Recent research has extended the dynamic pricing problem to many business applications, such as multiple products (Zhang and Cooper 2005), upselling of additional products (Aydin and Ziya 2008), sampling-based pricing (Cohen et al. 2018), add-on pricing (Song and Li 2018), multi-unit demand pricing (Gallego et al. 2020), and trade-in program pricing (Xiao and Zhou 2020).

We consider the MBG (or free returns) problem, which is closely related to cancelations, delayed purchases, and returns in revenue management. Subramanian et al. (1999) formulate a dynamic programming problem for the airline seat allocation problem that incorporates cancelations, overbooking, and no-shows. Similarly, Aydin et al. (2017) study the revenue management problem when an airline gives passengers the option to delay their final purchase decisions for a fixed time. Typically, both the purchase rate and the cancelation rate (or the delayed purchase rate) are taken to be exogenous. In contrast, we endogenize customers' purchase and return decisions through willingness-to-pay and fit uncertainty. Moreover, our work contributes to the literature by offering a simple threshold policy for whether to offer free returns, which is consistent with the practice in the fashion retail and hotel industries. While Aydin et al. (2017) model the deterministic problem as linear programming, we model the deterministic problem as an intensity control problem and use the Lagrangian approach to fully characterize the solution. Altug (2012) studies a dynamic return management problem with a fixed inventory. He assumes that a more flexible return policy can increase a consumer's valuation for the product but it leads to a higher return rate, and his research question is how to design an optimal return policy. We focus on the match between consumers and the product and our key question is whether to offer an MBG. Moreover, he does not consider positive return times, which is a main contribution of our work. Dai et al. (2019) investigate the network revenue management problem when customers may cancel their bookings or do not show up at the end of the horizon. Because the cancelation is modeled using an exponential clock for each customer, the fluid model does not have the piecewise property as ours does. Their model is more similar to the MBG models with a random (exponential or geometric) return time. See section 6.1 for more details.

In a recent paper, Hu et al. (2019) also study the firm's dynamic pricing policy when consumers can return the product. The authors assume the return time has an exponential distribution and focus on the firm's pricing decision in the stochastic system. There are two major differences between our work and theirs. First, the return time is fixed in our paper. This is inspired by the common practice of retailers (e.g., free returns within 30 days of delivery) and has been widely used to model product return behavior in the re-manufacturing literature (e.g., Debo et al. 2005,

Ferrer and Swaminathan 2006, Yan et al. 2017). As a result, the dynamic program suffers from the curse of dimensionality. Second, we study when MBGs should be offered, along with the pricing decision. Hu et al. (2019) find ignoring the returned products can lead to significant loss, but we show that the MBG option has an even deeper impact on the firm's profitability. Moreover, our methodology contribution and key managerial insights differ from those of Hu et al. (2019). In particular, we develop an iterative algorithm to solve the fluid model with infinite inventory constraints, and show that the firm should sell up the inventory first and then rely solely on selling the returned items. Consequently, our research makes a novel contribution to the revenue management literature.

3. The Model with Instant Returns

Consider a retailer that has made a sunk investment in the initial inventory (c) of a product which cannot be replenished during a predetermined sales horizon $[0, T]$. We use a discrete-time model as in Bitran and Mondschein (1997), where the sales horizon is divided into T periods, each of which is short enough so that there is at most one consumer arrival. For simplicity, we assume a time-homogeneous arrival probability λ in each period. We use *he* to refer to the retailer and *she* to refer to the customer.

3.1. Consumer Valuations

Denote the random reservation price of a customer for the product by V. Let $f(v)$ and $F(v)$ be the probability density function and the cumulative distribution function of V, respectively. We assume that $F(v)$ has an increasing failure rate (IFR), that is, $f(v)/\bar{F}(v)$ is increasing in $v \in [0, \infty)$, where $\bar{F}(v) \triangleq 1 - F(v)$. This is a common assumption in the revenue management literature (Ziya et al. 2004) to assure regularities.

When making purchasing decisions, customers are uncertain about whether the product fits their personal tastes. For example, when buying a shirt online, customers cannot view the actual size and thus it may turn out to be a mismatch upon delivery. We assume that the fit probability is β for all customers and that it is independent of everything else, including the valuation. This is a standard assumption also adopted by Shulman et al. (2009), Gao and Su (2016), and Gallino and Moreno (2018). If the product is priced at p, then a customer would take into account the fit probability, the reservation value, and the price. More precisely, if an MBG is offered (the MBG policy), her expected net utility is $\beta(V - p) + (1 - \beta)(-h)$. Here we use h to denote the hassle cost (or transaction cost) of returning the product. That is, when the product fits her taste (with probability β), the net utility is $V - p$; when the product does not fit (with probability $1 - \beta$), she returns the product and gets her money back,[2] incurring cost h during the process. Thus, she purchases the product if and only if $V \geq \tilde{p} := p + (1 - \beta)h/\beta$, where $\tilde{p}$ is the effective price perceived by the customer by considering the fitness and the MBG option—that is, the customer buys the product at the posted price if and only if the valuation is higher than or equal to the perceived price. If no money-back guarantee (NMBG) is offered, then her expected net utility is $\beta V - p$ and she purchases the product if and only if $V \geq \tilde{p} := p/\beta$, which is the associated perceived price with NMBG.

3.2. The Retailer's Problem

In each period t, the retailer determines whether to offer an MBG and the associated retail price p. If an MBG is offered, then the probability of a purchase is $\bar{F}(p + (1 - \beta)h/\beta)$, by the analysis above. Moreover, conditional on a sale, with probability $1 - \beta$ the item will be returned *instantly* in the case of a mismatch. We assume that the returned product is resalable afterwards and is regarded as the same as new items (Akcay et al. 2013, Mostard and Teunter 2006, Xie and Gerstner 2007). This is consistent with the current practice enforced by major online retailers (e.g., Amazon requires that consumers "return the eligible product(s) in new and unworn condition in the original packaging"; Joybuy requires that "items must be returned in a re-sellable condition"). For the NMBG policy, the probability of a purchase is $\bar{F}(p/\beta)$ and no return will be allowed.

Next we formulate the dynamic programming for the retailer, whose objective is to maximize the total expected revenue during the sales season. Let $V(t, n)$ be the optimal expected revenue-to-go in period t with n units of remaining inventory. The optimization problem follows from the Bellman equation:

$$\begin{aligned} V(t, n) = \max \Big\{ & \max_p \{\lambda \bar{F}(p/\beta)[p + V(t+1, n-1)] + [1 - \lambda \bar{F}(p/\beta)]V(t+1, n)\}, \\ & \max_p \{\lambda\beta \bar{F}(p + (1-\beta)h/\beta)[p + V(t+1, n-1)] + \lambda(1-\beta)\bar{F}(p + (1-\beta)h/\beta)V(t+1, n) \\ & + [1 - \lambda \bar{F}(p + (1-\beta)h/\beta)]V(t+1, n)\} \Big\}. \end{aligned} \tag{1}$$

with boundary conditions $V(t, 0) = 0$ for $t = 1, \ldots, T$ and $V(T + 1, n) = 0$ for all $n \leq c$. In the outer bracket, the optimal value function takes the maximum of (i) the optimal revenue with NMBG and (ii) the optimal revenue with MBG. In (i), the first term denotes the revenue-to-go in the event that a customer purchases the product in period t, and the second term is the revenue-to-go if no product is sold. In (ii), the first term represents a customer purchasing and keeping the product, the second term represents a customer purchasing but returning the product, and the third term represents no product being sold in that period. For easy reference, the main notations used throughout the paper are summarized in Table A1 in the appendix.

For any inventory n and time t, let $\Delta V(t, n) = V(t, n) - V(t, n - 1)$ be the marginal value of inventory. Rearranging Equation (1), we obtain

$$V(t, n) = \lambda \max(R_N(\Delta V(t+1, n)), \\ R_M(\Delta V(t+1, n))) + V(t+1, n),$$

where $R_N(\Delta) = \max_p \bar{F}(p/\beta)(p - \Delta)$ and $R_M(\Delta) = \max_p \beta\bar{F}(p + (1-\beta)h/\beta)(p - \Delta)$. Here, $R_M(\Delta)$ (resp. $R_N(\Delta)$) is the expected additional gain that the retailer can obtain by offering (resp. not offering) an MBG when the opportunity cost for selling the unit (the marginal value of inventory) is Δ. Using the definitions of perceived prices introduced in section 3.1, we have $\tilde{p} = p/\beta$ and thus $R_N(\Delta) = \max_{\tilde{p}} \bar{F}(\tilde{p})(\beta\tilde{p} - \Delta)$ for the NMBG policy. Similarly, for the MBG policy, $\tilde{p} = p + (1-\beta)h/\beta$ and thus $R_M(\Delta) = \max_{\tilde{p}} \bar{F}(\tilde{p})(\beta\tilde{p} - (1-\beta)h - \beta\Delta)$. It is easy to see that $R_M(\Delta) > R_N(\Delta)$ if and only if $\Delta > h$. As a result, it is optimal to offer an MBG when $\Delta > h$. Moreover, let $\tilde{p}^*(\Delta) = \arg\max_{\tilde{p}} \bar{F}(\tilde{p})(\beta\tilde{p} - \Delta)$. As $F(v)$ has an IFR, $\tilde{p}^*(\Delta)$ is uniquely determined (Ziya et al. 2004). Following the structures of $R_N(\Delta)$ and $R_M(\Delta)$, the optimal pricing policy for the retailer is given as follows.

Lemma 1. *For any (n, t), if $\Delta V(t + 1, n) > h$, it is optimal to offer an MBG with the perceived price $\tilde{p}^*(\beta\Delta V(t+1, n) + (1-\beta)h)$; otherwise, it is optimal to use NMBG with the perceived price $\tilde{p}^*(\Delta V(t+1, n))$.*

Lemma 1 makes it clear that the decision to offer an MBG is closely related to the marginal value of inventory. When the marginal value is relatively high, having customers return misfitting products is more valuable and thus it is optimal to offer an MBG. This is consistent with the literature (e.g., Akcay et al. 2013, Davis et al. 1995).

3.3. Properties of the Optimal Policy

We first characterize the structural properties of the optimal policy. The structural properties provide valuable insights into the computation and implementation of the optimal and heuristic policies.

Lemma 2. *The marginal value of inventory $\Delta V(t, n)$ is decreasing in n and t.*

In other words, the marginal value of inventory is diminishing when the end of the selling season approaches. Moreover, as more sales occur, the remaining inventory becomes more valuable. Such properties are consistent with those of the classic dynamic pricing policy in the literature (Gallego and Van Ryzin 1994). From the lemma, if we introduce $t_n \triangleq \min\{t : \Delta V(t+1, n) \leq h\}$, then for fixed n, the marginal value is less than or equal to h if and only if $t \geq t_n$. Combining this with Lemma 1 gives the following.

Proposition 1. *Given n, it is optimal for the retailer to offer an MBG if and only if $t < t_n$. Moreover, t_n is decreasing in n.*

Proposition 1 disentangles the complex interaction between the pricing and the decision of whether to offer an MBG. In particular, the policy has a simple thresholding structure: MBGs should be offered when the remaining sales season is sufficiently long; as the end of the selling season approaches, it is optimal to switch to NMBG, that is, free returns are no longer allowed. Moreover, t_n is decreasing in n, so the switch occurs earlier when the retailer has a higher stock level.

Next we investigate the impact of the fit probability β. Some online retailers create offline showrooms where customers can try on products before placing an order. These offline showrooms have been shown to reduce fulfillment costs (Bell et al. 2018). Another innovation that serves a similar purpose is the technology of virtual fitting rooms, which has been adopted by a growing number of prominent apparel retailers—QVC, Nordstrom, Amazon, among many others (Gallino and Moreno 2018). Such measures are aimed at increasing the fit probability. The next result shows how it affects the seller's MBG strategy.

Proposition 2. (i) *$\Delta V(t, n)$ is increasing in β*; (ii) *t_n is increasing in β.*

Proposition 2(i) shows that a higher fit probability leads to a higher marginal value of inventory. Technically, it indicates that the value function is supermodular in the fit probability and the stock level.

The super-modularity property seems to be intuitive but actually is not easy to satisfy. For example, $\Delta V(t, n)$ is not necessarily decreasing in the hassle cost.[3] The proposition implies that a higher fit probability leads to a higher expected profit, which coincides with the current endeavors of providing offline showrooms and developing virtual fitting room technology. As shown in Lemma 2, the opportunity cost determines whether to offer an MBG or not. Consequently, the seller's MBG strategy would be adjusted according to the change in the fit probability. Proposition 2(ii) indicates that MBGs should be implemented for a longer period when the seller increases the fit probability. This is because fewer products are returned due to the improved fit probability.

3.4. Benchmarks

One of the primary assumptions in section 3.2 is that the retailer can adjust the price and the MBG policy simultaneously and as frequently as needed. In practice, adjusting the MBG policy frequently may be costly and cause consumer backlash, while adjusting the price is widely used by many online retailers with automatic pricing and inventory tracking systems. In this section, we study variants of the main model in which the retailer has full flexibility in adjusting the price but has limited options in terms of the return policy. First, consider two benchmark models in which the firm must stick to the same return policy (either MBG or NMBG) throughout the whole sales horizon. Given the model setup and assumptions in section 3.1, let $V_M(t, n)$ and $V_N(t, n)$ denote the optimal expected revenue-to-go if the firm uses MBG and NMBG throughout the selling season, respectively. Then, we have the Bellman equations

$$\begin{aligned} V_N(t, n) = \max_p \{&\lambda\bar{F}(p/\beta)[p + V_N(t+1, n-1)] \\ &+[1 - \lambda\bar{F}(p/\beta)]V_N(t+1, n)\} \end{aligned} \tag{2}$$

and

$$\begin{aligned} V_M(t, n) = \max_p \{&\lambda\beta\bar{F}(p + (1-\beta)h/\beta)[p + V_M(t+1, n-1)] \\ &+\lambda(1-\beta)\bar{F}(p + (1-\beta)h/\beta)V_M(t+1, n) \\ &+[1 - \lambda\bar{F}(p + (1-\beta)h/\beta)]V_M(t+1, n)\} \end{aligned} \tag{3}$$

with boundary conditions $V_N(t, 0) = 0$ (resp. $V_M(t, 0) = 0$) for $t = 1, \ldots, T$ and $V_N(T+1, n) = 0$ (resp. $V_M(T+1, n) = 0$) for all $n \le c$. The terms inside the maximization in Equations (2) and (3) correspond to the two terms in Equation (1)—that is, the maximum revenue-to-go without and with MBG.

Pure-Return Policy. If the retailer is required to choose and stick to a return policy at the very beginning of the sales horizon, he would choose the one with higher revenue. That is,

$$V_P = \max\{V_M(0, c), V_N(0, c)\}. \tag{4}$$

In this benchmark, the retailer needs to commit to the return policy at the beginning. Consequently, the retailer in Equation (4) chooses whether to offer an MBG by selecting the one leading to more revenue, which is the maximum of Equations (2) and (3) initially.

MBG with Switch-off. We next examine a model in which the retailer can start with MBG but turn it off in the sales horizon. The final sale policy during the end of the sales season is almost an apparel industry standard that has been widely used by online and offline retailers—Ann Taylor, Gap, Zara, Uniqlo, among many others. We analyze the optimal stopping time to switch off the MBG policy. Under this policy, once MBG is switched off, it is no longer available and hence the revenue-to-go becomes identical to the value function V_N. Given the on-hand inventory n and time t, let $V_S(t, n)$ denote the optimal expected revenue-to-go for this policy. By the Bellman equation, for $t > 0$ and $n > 0$, we have

$$\begin{aligned} V_S(t, n) = \max\Big\{&V_N(t, n), \max_p\{\lambda\beta\bar{F}(p + (1-\beta)h/\beta) \\ &[p + V_S(t+1, n-1)] + [1 - \lambda\beta\bar{F} \\ &(p + (1-\beta)h/\beta)]V_S(t+1, n)\Big\} \end{aligned}$$

with boundary conditions $V_S(t, 0) = 0$ for $t = 1, \ldots, T$ and $V_S(T+1, n) = 0$ for all $n \le c$. We can characterize the optimal stopping time to switch off MBG.

Proposition 3. *For any inventory level $n > 0$, there exists a threshold $t_n^S \ge 0$ such that MBG is always offered if $t < t_n^S$.*

Proposition 3 implies that it is optimal to offer MBG at the beginning of the season, under which the marginal value of inventory remains high and hence the retailer can obtain higher revenue by selling the returned inventory to future customers. This is because using MBG has a certain "option" value for maintaining future flexibility. Therefore, the retailer should consider MBG only early in the selling season. There is a switching curve in the (t, n) space, as illustrated in Figure 1. If the retailer starts in the NMBG region, he should instantly use NMBG from then on. If he starts in the MBG region, he should use MBG

Figure 1 The Optimal Boundary for Money-Back Guarantees with Switch-Off [Color figure can be viewed at wileyonlinelibrary.com]

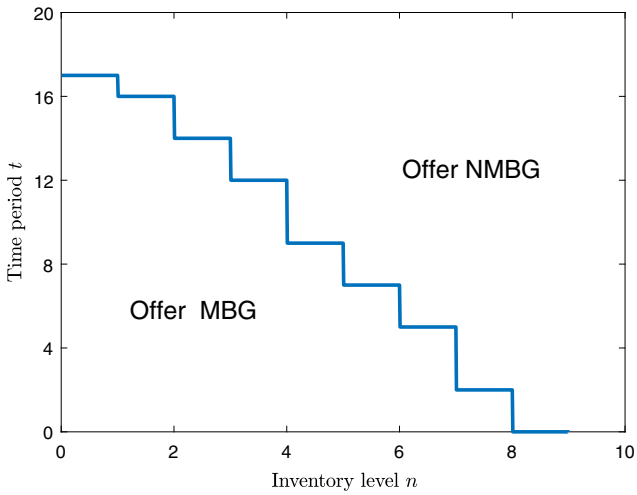

and switch to NMBG until the state (t, n) crosses the switching curve. Note that the switching thresholds t_n^S decrease in n, as predicted by Proposition 1.

3.4.1. Comparisons. To study the performance of the optimal policy and the benchmarks, we conduct a numerical study with 6250 scenarios covering a reasonable range of parameters. In particular, we consider $T = 50$ and $\lambda = 0.5$. The consumer valuations have a Gamma distribution, with mean $\mu = 15$ and standard deviation σ varying in the range $\sigma \in \{0.2\mu, 0.4\mu, \ldots, 1.0\mu\}$. The fit probability β is drawn from $\{0.5, 0.6, \ldots, 0.9\}$, which covers the empirical observations in Gallino and Moreno (2018). The hassle cost h is drawn from $\{1, 2, \ldots, 10\}$ such that the hassle-value ratio h/μ ranges from 6.7% to 67%. We vary the initial inventory level from 1 to 25.

Table 1 reports the quantiles of the revenue performance of the benchmarks relative to the optimal policy in all 6250 scenarios. Recall that we use the subscripts N, M, P, S, and F for NMBG only, MBG only, pure-return, switch-off, and the optimal policy for the flexible model, respectively. On average, NMBG only and MBG only perform relatively poorly, earning only, on average, 96.9% and 87.8% of the optimal policy model, respectively. This is primarily due to the poor performance of a few scenarios; the median performance of these two benchmarks is considerably better: 99.8% and 94.0%. The pure-return policy significantly improves the performance—the average revenue is 99.6%. However, there are a few scenarios in which the pure-return policy performs poorly—in the worst scenario, the pure-return policy earns only 92.6%. We find that it tends to perform badly for intermediate stock levels when the flexibility of choosing MBG or NMBG is most valuable. The switch-off policy performs nearly optimally—the average revenue is 99.9% of the optimal revenue, and the worst scenario still earns at least 98.1% of the revenue of the optimal policy.

Table 1 Relative Performance of Benchmarks (%)

	V_N/V_F	V_M/V_F	V_P/V_F	V_S/V_F
Minimum	58.3	25.8	92.6	98.1
5%	83.6	55.3	97.7	99.2
25%	96.4	81.6	99.6	99.8
50%	99.8	94.0	99.9	100.0
75%	100.0	99.1	100.0	100.0
95%	100.0	100.0	100.0	100.0
Maximum	100.0	100.0	100.0	100.0
Average	96.9	87.8	99.6	99.9

4. Positive Return Times

In practice, sellers usually allow for free returns within a grace period $\tau > 0$ after a purchase. Since customers do not have incentives to return the product before the end of the grace period, we assume a fixed return time τ. The setup is similar to that in section 3; in each period, the seller decides whether to offer an MBG and the associated price. When the arriving consumer makes a purchase, the seller generates revenue p. After τ periods, the seller needs to return the money to the buyer if he offers an MBG and the buyer finds a misfit. As the returned product is resalable, the inventory increases by one unit in that period. Moreover, buyers may still return purchased products after the end of the sales season; these products cannot be resold.

To formulate the dynamic program, note that in period t the seller needs the sales information of the previous τ periods to make the optimal decision. We record the MBG-sales information from periods $t - \tau$ to $t - 1$ as $\mathbf{s}^t = (s_i^t) \in \{0, 1\}^\tau$, where $s_i^t = 1$ represents a sale with MBG in period $t - i$ and $s_i^t = 0$ otherwise (i.e., a sale with NMBG or no sales). If $s_\tau^t = 1$, then the product sold in period $t - \tau$ is returned in period t with probability $1 - \beta$. We use a Bernoulli random variable $\mathbf{B}(s_\tau^t(1-\beta))$, where $s_\tau^t(1-\beta)$ is the success probability, to model the event of a return. Given the MBG-sales information $z \in \{0, 1\}$ in period t, the vector $\mathbf{s}$ is updated in period $t + 1$ as $\mathbf{s}^{t+1} = \Phi(\mathbf{s}^t, z)$, where $\mathbf{s}_1^{t+1} = z$ and $\mathbf{s}_i^{t+1} = \mathbf{s}_{i-1}^t$ for $i = 2, \ldots, \tau$.

In period t, given the on-hand stock n and the MBG-sales information $\mathbf{s}^t$ for the last τ periods, the seller needs to determine the MBG policy and the associated price. Let $V_t(n, \mathbf{s}^t)$ be the optimal expected revenue in period t with remaining stock n and sales history $\mathbf{s}^t$. We have

$$\begin{aligned} V_t(n, \mathbf{s}^t) = \max\Big\{ & \max_p \{\lambda \bar{F}(p/\beta)[p + \mathbb{E}V_{t+1}(n-1+\mathbf{B}(s_\tau^t(1-\beta)), \Phi(\mathbf{s}^t, 0))] \\ & +[1-\lambda\bar{F}(p/\beta)]\mathbb{E}V_{t+1}(n+\mathbf{B}(s_\tau^t(1-\beta)), \Phi(\mathbf{s}^t, 0))\}, \\ & \max_p \{\lambda \bar{F}(p+(1-\beta)h/\beta)[\beta p + \mathbb{E}V_{t+1}(n-1+\mathbf{B}(s_\tau^t(1-\beta)), \Phi(\mathbf{s}^t, 1))] \\ & +[1-\lambda\bar{F}(p+(1-\beta)h/\beta)]\mathbb{E}V_{t+1}(n+\mathbf{B}(s_\tau^t(1-\beta)), \Phi(\mathbf{s}^t, 0))\}\Big\} \end{aligned} \tag{5}$$

for $n \geq 1$ and $V_t(0, \mathbf{s}^t) = \mathbb{E}V_{t+1}(\mathbf{B}(s_\tau^t(1-\beta)), \Phi(\mathbf{s}^t, 0))$, with terminal conditions $V_{T+1}(n, \mathbf{s}^{T+1}) = 0$ for any $(n, \mathbf{s}^{T+1})$. The initial state is $\mathbf{s}^1 = \mathbf{0}$, where $\mathbf{0}$ is a τ-dimensional vector of zeros. Note that the expected revenue from an MBG buyer is βp, as the portion of $(1-\beta)p$ would be returned to the buyer later.

On the right-hand side of the Bellman equation, the seller compares the revenues of MBG/NMBG. In both cases, the first term represents the revenue-to-go if a customer arrives and purchases the product in period t, and the second term is the revenue-to-go if no product is sold. The Bernoulli random variable represents the potential return. That is,

$$\mathbb{E}V_{t+1}(n+\mathbf{B}(s_\tau^t(1-\beta)), \mathbf{s}^{t+1}) = \begin{cases} \beta V_{t+1}(n, \mathbf{s}^{t+1}) + (1-\beta)V_{t+1}(n+1, \mathbf{s}^{t+1}) & \text{if } s_\tau^t = 1; \\ V_{t+1}(n, \mathbf{s}^{t+1}) & \text{if } s_\tau^t = 0. \end{cases}$$

We have the following structural properties.

Proposition 4. *For any $t \geq T - \tau$, the optimal policy is to use the NMBG policy.*

Proposition 4 states that allowing customers to return items is not optimal when the remaining sales season is less than τ, consistent with the intuition. Any items sold during $[T-\tau, T]$ will be returned, if allowed, to the retailer after the end of the season. As these returned items cannot be resold, this policy incurs significant costs and thus MBG should not be offered.

Since the discrete formulation approximates reality by using short periods (at most one customer arrives within a period), T and τ are large measured by the number of periods. Therefore, the dimension of the state variable $\tau + 1$ makes the exact solution computationally infeasible. Next, we study the fluid approximation and derive heuristics for the problem.

4.1. Fluid Model

The fluid approximation has proved to be a powerful tool in the revenue management literature, in providing heuristics and performance bounds when the optimal policy is intractable. More precisely, consider the deterministic formulation of the problem: The length of the selling horizon and the initial inventory level are still T and c. Different from the stochastic model, the product and time periods are assumed to be infinitely divisible, and consumers arrive in a flow. That is, within an infinitesimal period of length dt, there are λdt customers arriving whose valuation of the product has distribution $F(v)$. Therefore, if the retailer offers an MBG at price p_t at t, then the (deterministic) purchase rate is $\lambda\bar{F}(p_t + (1-\beta)h/\beta)$. Moreover, the return rate at $t+\tau$ is $\lambda(1-\beta)\bar{F}(p_t + (1-\beta)h/\beta)$. If an MBG is not offered, then the purchase rate is $\lambda\bar{F}(p_t/\beta)$. Thus, we can formulate the deterministic optimization problem for the fluid model as

$$\begin{aligned} \bar{V}(T, c) = \max_{m_t \in \{0;1\}, p_t} & \Big\{ \int_0^T \lambda m_t p_t \bar{F}(p_t + (1-\beta)h/\beta)dt + \int_0^T \lambda(1-m_t)p_t\bar{F}(p_t/\beta)dt \\ & - \int_0^T \lambda m_t(1-\beta)p_t\bar{F}(p_t + (1-\beta)h/\beta)dt \Big\} \\ \text{subject to} \quad & \int_0^s \lambda m_t \bar{F}(p_t + (1-\beta)h/\beta)dt + \int_0^s \lambda(1-m_t)\bar{F}(p_t/\beta)dt \\ & - \int_0^{(s-\tau)^+} \lambda m_t(1-\beta)\bar{F}(p_t + (1-\beta)h/\beta)dt \leq c, \quad \forall 0 \leq s \leq T. \end{aligned}$$

Here $\bar{V}(T, c)$ denotes the optimal value for given T and c. The binary decision m_t captures whether an MBG is offered ($m_t = 1$) or not ($m_t = 0$) at t. The objective function is the total revenue earned during $[0, T]$. The constraint corresponds to the requirement of non-negative inventory at any time

s, which is the initial inventory minus the total number of items sold by s plus the number of items returned (from the items sold by $(s-\tau)^+$).

As in section 3, we replace the price p_t with the perceived price $\tilde{p}_t$ for the subsequent analysis. Thus, the optimization is equivalent to

$$\text{(DPR)} \quad \bar{V}(T,c) = \max_{m_t \in \{0;1\}, \tilde{p}_t} \left[\int_0^T \lambda m_t \bar{F}(\tilde{p}_t)[\beta \tilde{p}_t - (1-\beta)h]dt + \lambda(1-m_t)\beta \tilde{p}_t \bar{F}(\tilde{p}_t)dt \right]$$
$$\text{subject to} \int_0^s \lambda m_t \bar{F}(\tilde{p}_t)dt + \int_0^s \lambda(1-m_t)\bar{F}(\tilde{p}_t)dt$$
$$- \int_0^{(s-\tau)^+} \lambda m_t (1-\beta)\bar{F}(\tilde{p}_t)dt \le c, \quad \forall 0 \le s \le T. \tag{6}$$

Here "DPR" stands for the (D)eterministic formulation with (P)ositive (R)eturn time.

For most studies in the revenue management literature (e.g., Gallego and Van Ryzin 1994), the fluid approximation has one constraint. That is, the retailer only needs to make sure that the inventory is nonnegative at the end of the sales horizon T. In Problem (6), however, there are infinite constraints. This is because returned products make the constraint nonredundant at any time t. The presence of the constraints greatly complicates the fluid model.

Next, we characterize some structural properties of the solution to Problem (6).

Proposition 5. *For the DPR problem* (6), *we have* (i) *for any feasible policy* $\{\tilde{p}_t, m_t\}_{t=0}^T$, *we can find another policy with the MBG threshold structure*

$$m_t = \begin{cases} 1 & \text{for } t \le T-\eta; \\ 0 & \text{for } t > T-\eta \end{cases} \tag{7}$$

that generates no less revenue for some $\eta \ge 0$; (ii) *for any optimal policy, it always uses NMBG* (i.e., $m_t = 0$) *for* $t > T - \tau$.

Proposition 5(i) shows that it is sufficient to investigate policies that implement an MBG at the beginning of the horizon but an NMBG at the end. In a class-based revenue management problem, Dai et al. (2019) derive a similar threshold policy that accepts customers from a certain class first and rejects those from the same class later. This is because customers arriving earlier are more profitable than those arriving later. Nevertheless, our result is driven by inventory availability. Intuitively, arranging MBGs at the beginning guarantees the maximum number of returned items that can potentially be resold. Proposition 5(ii) states that the seller should not allow customers to return items when the remaining sales season is less than τ in the fluid model, which is analogous to Proposition 4 for the stochastic system.

Given Proposition 5, we can reduce the dimension of the pricing policy from $\{\tilde{p}_t, m_t\}_{t=0}^T$ to $\{\tilde{p}_t\}_{t=0}^T$ and $\eta \in [\tau, T]$, where η is the length of the NMBG duration, such that the MBG policy is specified by Equation (7). In particular, DPR in (6) can be transformed into

$$\text{(DPR)} \quad \bar{V}(T,c) = \max_{0 \le \eta \le T, \tilde{p}_t} \left[\int_0^{T-\eta} \lambda \bar{F}(\tilde{p}_t)[\beta \tilde{p}_t - (1-\beta)h]dt + \int_{T-\eta}^T \lambda \beta \tilde{p}_t \bar{F}(\tilde{p}_t)dt \right]$$
$$\text{subject to } I(s, \eta, \tilde{p}_t) \le c, \quad \forall 0 \le s \le T;$$
$$\eta \ge \tau, \tag{8}$$

where

$$I(s, \eta, \tilde{p}_t) = \int_0^s \lambda \bar{F}(\tilde{p}_t)dt - \int_0^{\min((s-\tau)^+, T-\eta)} \lambda(1-\beta)\bar{F}(\tilde{p}_t)dt$$

is net sales, that is, the number of items purchased minus returns at time s. Although the firm's optimal policy is simplified, there are still an infinite number of inventory constraints. In this section, we devote ourselves to the solution to Problem (8). We sequentially choose a finite set of time points, at which the inventory constraints are binding at its optimal solution, and relax the other constraints. This allows us to analyze a sequence of relaxed problems that gradually approximates and eventually allows us to analytically solve Problem (8). Indeed, if the optimal solution from the relaxed problem can be checked to

satisfy all inventory constraints, then it is optimal for Problem (8) as well.

In particular, we first consider a relaxed problem of DPR in (8) by considering only the inventory constraint at $s = T$ and the NMBG duration constraint $\eta \geq \tau$, which we refer to as "rDPR-0." If the solution to rDPR-0 violates some other inventory constraints for $s \in [0, T)$, we gradually add additional inventory constraints at $s = T - k\tau$ for $k \geq 1$ to rDPR-0, which we refer to as "rDPR-k." The time points are carefully selected so that for at least one k, the finite inventory constraints of rDPR-k guarantee non-negative inventory for all $s \in [0, T]$. As a result, the optimal solution to rDPR-k is feasible for DPR, which thus is also optimal.

4.1.1. Initial Step. rDPR-0 is a relaxed problem of DPR in (8) that has only two constraints: the inventory constraint at time T—that is, $I(s = T, \eta, \tilde{p}_t) \leq c$—and the NMBG duration constraint $\eta \geq \tau$. Given $\eta \geq \tau$, the net sales at T becomes $I(s = T, \eta, \tilde{p}_t) = \int_0^{T-\eta} \lambda\beta\bar{F}(\tilde{p}_t) + \int_{T-\eta}^{T} \lambda\bar{F}(\tilde{p}_t)dt$. Thus, rDPR-0 is equivalent to:

$$\text{(rDPR-0)} \quad \bar{V}_0(T, c) = \max_{0 \leq \eta \leq T, \tilde{p}_t} \int_0^{T-\eta} \lambda\bar{F}(\tilde{p}_t)[\beta\tilde{p}_t - (1-\beta)h] + \int_{T-\eta}^{T} \lambda\beta\tilde{p}_t\bar{F}(\tilde{p}_t)dt$$
$$\text{subject to} \int_0^{T-\eta} \lambda\beta\bar{F}(\tilde{p}_t) + \int_{T-\eta}^{T} \lambda\bar{F}(\tilde{p}_t)dt \leq c;$$
$$\eta \geq \tau. \tag{9}$$

To analyze and understand the optimal solution to Problem (9), we first consider two related deterministic problems where the return time is instant (i.e., $\tau = 0$) and either NMBG or MBG is available throughout the whole sales season, which corresponds to the fluid models for the two stochastic benchmarks in (2) and (3), respectively.

Fluid Model with NMBG. Given the benchmark with NMBG in (2), the corresponding fluid model is $\bar{V}_N(T, c) = \max_{\tilde{p}_t} \int_0^T \lambda\bar{F}(\tilde{p}_t)\beta\tilde{p}_t dt$ such that $\int_0^T \lambda\bar{F}(\tilde{p}_t)dt \leq c$, which is equivalent to the standard intensity control problem studied by Gallego and Van Ryzin (1994). A crucial observation made in the literature is that if $F(v)$ has an IFR, then the instantaneous revenue rate $\lambda\beta\tilde{p}_t\bar{F}(\tilde{p}_t)$ is a concave function of the purchase rate $\lambda\bar{F}(\tilde{p}_t)$ after we transform the price p_t into a function of the purchase rate using the inverse of $\bar{F}$. As a result of Jensen's inequality, given the inventory constraint, it is optimal to use a constant price (or equivalently, a constant purchase rate) over the horizon. Therefore, the optimal perceived price is, for $t \in [0, T]$, $\tilde{p}_t \equiv \tilde{p}^*(0)$ if $c \geq \lambda\bar{F}(\tilde{p}^*(0))T$ and $\tilde{p}_t \equiv \bar{F}^{-1}(\frac{c}{\lambda T})$ otherwise.

Fluid Model with MBG. If MBGs are always offered as in (3), the corresponding fluid model is $\bar{V}_M(T, c) = \max_{\tilde{p}_t} \int_0^T \lambda\bar{F}(\tilde{p}_t)[\beta\tilde{p}_t - (1-\beta)h]dt$ such that $\int_0^T \lambda\beta\bar{F}(\tilde{p}_t)dt \leq c$. This can be transformed into the classic dynamic pricing problem (Gallego and Van Ryzin 1994). Notice that the optimal perceived price to maximize the instantaneous revenue rate under MBG is constant $\tilde{p}^*((1-\beta)h)$. Thus, the optimal perceived price is $\tilde{p}_t \equiv \tilde{p}^*((1-\beta)h)$ if $c \geq \lambda\beta\bar{F}(\tilde{p}^*((1-\beta)h))T$ and $\tilde{p}_t \equiv \bar{F}^{-1}(\frac{x}{\lambda\beta T})$ otherwise.

Lemma 3. *For the fluid model of pure-return policy (either MBG or NMBG), the optimal (perceived) price is constant over the sales season.*

We leverage the result to analyze the optimal solution for rDPR-0. Notice that the rDPR-0 problem can be regarded as a two-stage problem: first deciding the length and the amount of inventory to sell in the first MBG period and the remaining NMBG period; then deciding the prices $\tilde{p}_t$ within each period. As indicated by Lemma 3, the optimal price must be constant during each period. More precisely, under the optimal policy of Problem (9), $\tilde{p}_t \equiv \tilde{p}_{0,1}$ for $t \in [0, T-\eta)$ and $\tilde{p}_t \equiv \tilde{p}_{0,0}$ otherwise, where the first subscript refers to rDPR-"0" and the second subscript represents the associated period. Notice that we count the periods in reverse order so 0 is closer to the end of the horizon T. As a result, rDPR-0 is equivalent to

$$\text{(rDPR}-0) \quad \bar{V}_0(\mathrm{T}, \mathrm{c}) = \max_{0 \leq \eta \leq \mathrm{T}, \tilde{\mathrm{p}}_{0,0}, \tilde{\mathrm{p}}_{0,1}} \left[(\mathrm{T}-\eta)\lambda\bar{F}(\tilde{p}_{0,1})[\beta\tilde{p}_{0,1} - (1-\beta)h] + \eta\lambda\beta\bar{F}(\tilde{p}_{0,0})\tilde{p}_{0,0}\right],$$
$$\text{subject to} \quad (\mathrm{T}-\eta)\lambda\beta\bar{F}(\tilde{p}_{0,1}) + \eta\lambda\bar{F}(\tilde{p}_{0,0}) \leq c;$$
$$\eta \geq \tau. \tag{10}$$

The optimal solution to rDPR-0 can be characterized as follows.

LEMMA 4. *Let* $\bar{c} \triangleq \lambda\bar{F}(\tilde{p}^*(h))T$ *and* $\underline{c} \triangleq [(1-\beta)\frac{\tau}{T}+\beta]\bar{c}$ *be two thresholds of the initial inventory. The optimal solution* $(\eta_0^*,\ \tilde{p}_{0,0}^*,\ \tilde{p}_{0,1}^*)$ *to rDPR-0* (10) *is*

(i) *if* $c \geq \bar{c}$, $\eta_0^* = T$, $\tilde{p}_{0,0}^* = \max(\tilde{p}^*(0),\ \bar{F}^{-1}(\frac{c}{\lambda T}))$, *and* $\tilde{p}_{0,1}^*$ *is irrelevant;*

(ii) *if* $\underline{c} \leq c < \bar{c}$, $\eta_0^* = \frac{c-\beta\bar{c}}{\bar{c}-\beta\bar{c}}T \geq \tau$ *and* $\tilde{p}_{0,0}^* = \tilde{p}_{0,1}^* = \tilde{p}^*(h)$;

(iii) *if* $c < \underline{c}$, $\eta_0^* = \tau$ *and* $(\tilde{p}_{0,0}^*,\ \tilde{p}_{0,1}^*)$ *uniquely solves*

$$\begin{aligned} \max_{(\tilde{p}_{0,0},\ \tilde{p}_{0,1})} \quad & \Big[(T-\tau)\lambda\bar{F}(\tilde{p}_{0,1})[\beta\tilde{p}_{0,1}-(1-\beta)h]+\tau\lambda\beta\bar{F}(\tilde{p}_{0,0})\tilde{p}_{0,0}\Big] \\ \text{subject to} \quad & (T-\tau)\lambda\beta\bar{F}(\tilde{p}_{0,1})+\tau\lambda\bar{F}(\tilde{p}_{0,0}) = c. \end{aligned} \tag{11}$$

Intuitively, when the initial inventory is relatively high, the NMBG duration constraint $\eta \geq \tau$ is redundant, as shown in Lemma 4(i) and 4(ii). In this case, the rDPR-0 problem balances the time and inventory allocated to the two sales modes (i.e., NMBG and MBG), which is similar to the fluid model in Hu et al. (2017). Furthermore, we utilize the Lagrangian method to solve the problem and derive its explicit solution. When the initial inventory is relatively low, both constraints for Problem (10) are binding at its optimal solution, as implied in Lemma 4(iii). Problem (11) can be easily translated into a concave maximization problem with a linear constraint regarding the purchase rates $(\bar{F}(\tilde{p}_0),\ \bar{F}(\tilde{p}_1))$ as decision variables.

We now link rDPR-0 to DPR. The optimal solution for rDPR-0 in Problem (9) uses MBG with $\tilde{p}_t^* = \tilde{p}_{0,1}^*$ for $t \in [0,\ T-\eta_0^*]$ and NMBG with $\tilde{p}_t^* = \tilde{p}_{0,0}^*$ for $t \in (T-\eta_0^*,\ T]$. Since rDPR-0 is a relaxation of DPR, the feasibility of rDPR-0's optimal solution for DPR implies the optimality. DPR imposes inventory constraints for $s \in [0,T]$, hence we need to analyze the underlying inventory process under the optimal policy of rDPR-0.

- If $c \geq \bar{c}$, by Lemma 4(i), the retailer uses NMBG and sets a constant price for the sales horizon. Accordingly, the inventory decreases at a constant rate $\lambda\bar{F}(\tilde{p}_{0,0}^*)$ for $t \in [0,T]$ and is non-negative at $s = T$. Thus, the optimal solution for rDPR-0 is optimal for DPR in (8).
- If $\underline{c} \leq c < \bar{c}$, by Lemma 4(ii), the optimal policy for rDPR-0 uses MBG for $t \in [0, T-\eta_0^*)$ and NMBG for $t \in [T-\eta_0^*, T]$ with the same perceived price $\tilde{p}^*(h)$. The return rate from the MBG phase is $(1-\beta)\lambda\bar{F}(\tilde{p}^*(h))$, which never exceeds the purchase rate. Therefore, the inventory is always decreasing during $[0, T]$. Since the inventory constraint is binding at T, it implies that the retailer's inventory is always non-negative, and this policy is feasible and optimal for DPR.
- If $c < \underline{c}$, the optimal policy of rDPR-0 uses MBG with a perceived price $\tilde{p}_{0,1}^*$ during $[0,\ T-\tau]$ and NMBG with $\tilde{p}_{0,0}^*$ during $[T-\tau,\ T]$. Accordingly, the inventory process has a constant sales rate, equal to the purchase rate minus the return rate (if any), in the following periods:

$$\frac{\partial I(s,\ \tau,\ \tilde{p}_t^*)}{\partial s} = \begin{cases} \lambda\bar{F}(\tilde{p}_{0,1}^*) & s \in [0, \min\{\tau, T-\tau\}); \\ \lambda\beta\bar{F}(\tilde{p}_{0,1}^*) & s \in [\min\{\tau, T-\tau\},\ T-\tau); \\ \lambda\bar{F}(\tilde{p}_{0,0}^*) & s \in [T-\tau,\ \max\{\tau,\ T-\tau\}); \\ \lambda\bar{F}(\tilde{p}_{0,0}^*)-\lambda(1-\beta)\bar{F}(\tilde{p}_{0,1}^*) & s \in [\max\{\tau,\ T-\tau\},\ T]. \end{cases}$$

The sales rate is piecewise constant in the three periods, depending on $\tau < T-\tau$ or $\tau \geq T-\tau$. The inventory is always decreasing for $t \in [0, \max\{\tau,\ T-\tau\})$. However, it might be increasing for $t \in [\max\{\tau,\ T-\tau\},\ T]$ as illustrated in Figure 2. This is a consequence of relaxing the inventory constraints of DPR. In particular, if the return rate during $[\max\{\tau,\ T-\tau\},\ T]$ is less than the purchase rate, that is, $(1-\beta)\bar{F}(\tilde{p}_{0,1}^*) \leq \bar{F}(\tilde{p}_{0,0}^*)$, then the inventory is always decreasing during $[0, T]$. If the return rate exceeds the purchase rate, that is, $(1-\beta)\bar{F}(\tilde{p}_{0,1}^*) > \bar{F}(\tilde{p}_{0,0}^*)$, then the inventory is strictly increasing for $s \in [\max\{\tau,\ T-\tau\},\ T]$, as shown in Figure 2. Since it reaches zero at T, the inventory level is negative for $s \in [\max\{\tau,\ T-\tau\},\ T)$. The explanation for this phenomenon is the interweaving inventory process: although the prices are piecewise constant with two pieces, because of the delayed return, the piecewise inventory process consists of three pieces. Imposing an inventory constraint at T does not automatically guarantee that the purchase rate exceeds the return rate in each piece. In this case, the optimal solution of rDPR-0 violates the inventory constraint of DPR in (8). To address this issue, we next gradually add inventory constraints to rDPR-0.

4.1.2. Recursive Steps: A Sequence of Relaxed DPR Problems with Selected Inventory Constraints. From the previous discussion, the optimal solution of rDPR-0 is infeasible for DPR if the return rate exceeds the purchase rate in the third period. That is, if $(1-\beta)\bar{F}(\tilde{p}_{0,1}^*) > \bar{F}(\tilde{p}_{0,0}^*)$, then the inventory is negative for $s \in [\max\{\tau,\ T-\tau\},\ T)$. Notice that the inventory process reaches its minimum at $\max\{\tau,\ T-\tau\}$, as shown in Figure 2. One would naturally think of imposing an additional inventory constraint at $s = \max\{\tau,\ T-\tau\}$ in rDPR-0.

If $\tau \geq T-\tau$ (as in Figure 2b), then after imposing an additional inventory constraint at τ, the optimal

Figure 2 A Possible Realization of the Inventory Process Associated with the Optimal Solution to rDPR-0. In Both Panels, in the First Period the Inventory Level is Driven by Purchases in the Money-Back Guarantees (MBG) Phase; in the Last Period, it is Driven by the Purchases in the No MBG (NMBG) Phase and the Returned Items from the MBG Phase. The Inventory in the Second Period of the Left Panel is Driven by Both Purchases and Returns in the MBG Phase. The Inventory in the Second Period of the Right Panel is Driven Only by the Purchases in the NMBG Phase. In Both Cases, the Inventory Constraints for DPR are Violated

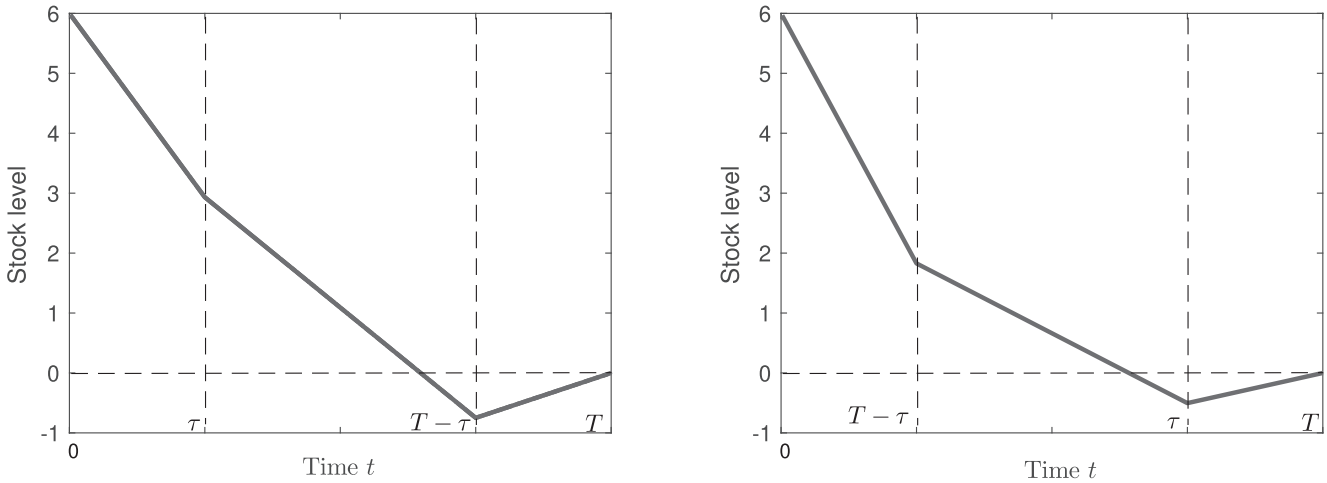

(a) Short return time ($\tau = 0.25T$) (b) Long return time ($\tau = 0.75T$)

solution to the new problem (as indicated by Lemma 3) is to offer constant prices during $[0, T-\tau]$, $[T-\tau, \tau]$, and $[\tau, T]$, respectively, where MBGs are offered in the first period. The additional constraint forces the inventory to be non-negative at τ, guaranteeing the optimal solution is feasible and optimal for DPR.

If $\tau < T - \tau$ (as in Figure 2a), or equivalently $2\tau < T$, then after imposing an additional constraint at $T - \tau$, the optimal solution is to offer MBGs during $[0, T-\tau]$. Moreover, the prices are constant during $[0, T-2\tau]$, $[T-2\tau, T-\tau]$, and $[T-\tau, T]$, respectively. The additional inventory constraint at $T - \tau$ implies that the purchase rate is equal to the return rate during $s \in [T-\tau, T]$, then the retailer's three-period pricing problem degenerates to a two-period pricing problem similar to rDPR-0. Consequently, this constraint guarantees non-negative inventory for $s \in [T-\tau, T]$, but it may lead to negative inventory for $s \in [\max\{\tau, T-2\tau\}, T-\tau)$ if the return rate exceeds the purchase rate in this period. In this case, we need to impose an additional inventory constraint at $\max\{\tau, T-2\tau\}$. The process continues until at some iteration, the return rate does not exceed the purchase rate and hence the optimal solution is feasible for DPR.

The discussion motivates us to consider a sequence of optimization problems. Let $K = \lfloor T/\tau \rfloor$ be the largest integer less than T/τ. For any $1 \le k \le K$, we consider the following sequence of relaxations of DPR:

$$\text{(rDPR-}k\text{)} \quad \bar{V}_k(T, c) = \max_{\eta, \tilde{p}_t} \left[\int_0^{T-\eta} \lambda \bar{F}(\tilde{p}_t)[\beta \tilde{p}_t - (1-\beta)h]dt + \int_{T-\eta}^{T} \lambda \beta \tilde{p}_t \bar{F}(\tilde{p}_t)dt \right]$$
$$\text{subject to} \quad I(s, \eta, \tilde{p}_t) \le c, \quad \forall s \in S_k;$$
$$\eta \ge \tau, \tag{12}$$

where

$$S_k = \begin{cases} \{T, \tau\}, & \text{if } K = 1; \\ \{T, T-\eta, T-\eta-\tau, \ldots, T-\eta-(k-2)\tau, T-\eta-(k-1)\tau\}, & \text{if } K \ge 2, k < K; \\ \{T, T-\eta, T-\eta-\tau, \ \ldots, T-\eta-(k-2)\tau\} \cup \{k\tau, (k-1)\tau, \ \ldots, \tau\}, & \text{if } K \ge 2, k = K. \end{cases}$$

Instead of imposing non-negative inventory constraints for all $s \in [0, T]$ as in DPR, the rDPR-k formulation imposes them only for $s \in S_k$. Compared to rDPR-0, rDPR-k gradually adds more constraints that guarantee non-negative inventory at selected time points. As a result, the set of inventory constraints has a nested structure—that is, $S_k \subset S_{k+1}$. It can be shown that in the optimal solution of rDPR-k, the firm offers MBGs during $[0, T - \tau]$, that is, $\eta = \tau$. Hence S_k is equivalent to[4]

$$S_k = \begin{cases} \{T, T-\tau,\ T-2\tau,\ \ldots,\ T-k\tau\}, & \text{if } k < K; \\ \{T,\ T-\tau,\ \ldots,\ T-(k-1)\tau\} & \\ \cup\{k\tau,\ (k-1)\tau,\ \ldots,\ \tau\}, & \text{if } k = K. \end{cases}$$

Lemma 5. *There always exists some $k \in \{0, 1, \ldots, K\}$ such that the optimal solution of rDPR-k in* (12) *is feasible for DPR in* (8).

Lemma 5 indicates that the time points in S_k are carefully selected so that for at least one k, the optimal solution of rDPR-k is feasible for DPR. Let κ be the smallest k such that the optimal solution of rDPR-k is feasible for DPR. The nested structure of inventory constraint sets enables us to develop an iterative algorithm to find κ by solving rDPR-k sequentially for $k = 0, 1, \ldots$. The initial step rDPR-0 has been discussed. We proceed to consider rDPR-k for $k \geq 1$.

For any $k \leq \min(\kappa, K-1)$, which implies that the optimal solution of rDPR-$(k-1)$ is not feasible for DPR, one can show that all its constraints are binding for rDPR-k at its optimal solution. Moreover, as indicated by Lemma 3, the optimal solution offers constant prices during $[0, T-(k+1)\tau]$ and $[T-(i+1)\tau, T-i\tau]$ for $0 \leq i \leq k$. Let $\tilde{p}_{k,i}$ and $\tilde{p}_{k,k+1}$ be the perceived prices for periods $[T-(i+1)\tau, T-i\tau]$ for $0 \leq i \leq k$ and $[0, T-(k+1)\tau]$, respectively, where the subscript k refers to rDPR-"k." Thus, Problem (12) becomes

$$\max_{\tilde{p}_{k,i}} \left\{ \lambda\bar{F}(\tilde{p}_{k,k+1})[\beta\tilde{p}_{k,k+1} - (1-\beta)h][T-(k+1)\tau] + \sum_{i=1}^{k} \lambda\bar{F}(\tilde{p}_{k,i})[\beta\tilde{p}_{k,i} - (1-\beta)h]\tau + \lambda\bar{F}(\tilde{p}_{k,0})\beta\tilde{p}_{k,0}\tau \right\}$$
$$\text{subject to} \quad \lambda\beta\bar{F}(\tilde{p}_{k,k+1})[T-(k+1)\tau] + \sum_{j=i+1}^{k} \lambda\beta\bar{F}(\tilde{p}_{k,j})\tau + \lambda\bar{F}(\tilde{p}_{k,i})\tau = c, \quad \forall i \in \{0,\ \ldots,\ k\}. \tag{13}$$

Note that the constraints in Problem (13) can be simplified as $\lambda\beta\bar{F}(\tilde{p}_{k,k+1})[T-(k+1)\tau] + \lambda\bar{F}(\tilde{p}_{k,k})\tau = c$ and $\bar{F}(\tilde{p}_{k,i-1}) = (1-\beta)\bar{F}(\tilde{p}_{k,i})$ for all $i \in \{1, \ldots, k\}$, which indicates that the seller first sells up all inventory by the time $T - k\tau$ and then sells returned items by keeping zero inventory. Problem (13) can be easily translated into a concave maximization problem with linear constraints in terms of the purchase rates. Thus, the optimal solution is uniquely determined, and denoted as $(\tilde{p}^*_{k,0}, \ldots, \tilde{p}^*_{k,k+1})$. The detailed description of the algorithm is given in Algorithm 1. Notice that the condition for the optimal solution of rDPR-k to be feasible for DPR can be simplified as $(1-\beta)\bar{F}(\tilde{p}^*_{k,k+1}) \leq \bar{F}(\tilde{p}^*_{k,k})$.

Algorithm 1 Determine κ

```
Input: β, T, h, τ, c, F(v)
for k = 0, ..., K − 1 do
    Solve rDPR-k in (13) and calculate (p̃_{k,0}, ..., p̃_{k,k+1})
    if (1 − β)F̄(p̃*_{k,k+1}) ≤ F̄(p̃*_{k,k}) then
        Terminate and return κ = k
    end if
end for
Return κ = K
```

Accordingly, if $\kappa < K$, the optimal solution of rDPR-κ is given above. If $\kappa = K$, similarly, all the constraints are binding for rDPR-K at its optimal solution, which offers constant prices during $[K\tau - i\tau, T - i\tau]$ for $0 \leq i \leq K$ and $[T - i\tau, (K+1)\tau - i\tau]$ for $1 \leq i \leq K$. Let $\tilde{p}_{K,2i}$ (resp. $\tilde{p}_{K,2i-1}$) be the perceived prices for $[K\tau - i\tau, T - i\tau]$ for $0 \leq i \leq K$ (resp. $[T - i\tau, (K+1)\tau - i\tau]$ for $1 \leq i \leq K$). Accordingly, the rDPR-K problem in (12) becomes

$$\begin{aligned}
\max_{\tilde{p}_{K,i}} \Bigg\{ & \sum_{i=1}^{K} \lambda \bar{F}(\tilde{p}_{K,2i})[\beta \tilde{p}_{K,2i} - (1-\beta)h](T - K\tau) + \lambda \bar{F}(\tilde{p}_{K,0})\beta \tilde{p}_{K,0}(T - K\tau) \\
& + \sum_{i=2}^{K} \lambda \bar{F}(\tilde{p}_{K,2i-1})[\beta \tilde{p}_{K,2i-1} - (1-\beta)h][(K+1)\tau - T] + \lambda \bar{F}(\tilde{p}_{K,1})\beta \tilde{p}_{K,1}[(K+1)\tau - T] \Bigg\} \\
\text{subject to} \quad & \lambda \bar{F}(\tilde{p}_{K,2K})[T - K\tau] + \lambda \bar{F}(\tilde{p}_{K,2K-1})[(K+1)\tau - T] = c; \\
& \bar{F}(\tilde{p}_{K,2i}) = (1-\beta)\bar{F}(\tilde{p}_{K,2(i+1)}) \quad \forall i \in \{0, 1, \ldots, K-1\}; \\
& \bar{F}(\tilde{p}_{K,2i-1}) = (1-\beta)\bar{F}(\tilde{p}_{K,2i+1}) \quad \forall i \in \{1, \ldots, K-1\}.
\end{aligned} \tag{14}$$

The optimal solution for Problem (14) is also uniquely determined, and denoted as $(\tilde{p}^*_{K,0}, \ldots, \tilde{p}^*_{K,2K})$.

4.1.3. Optimal Solution to DPR. We next link rDPR-κ to DPR. By definition, the optimal solution for rDPR-κ is feasible for DPR, which further implies the optimality since rDPR-κ is a relaxation of DPR. Combining all the previous results, we have

PROPOSITION 6. *The optimal solution* $(\eta^*, \tilde{p}^*_t)$ *to DPR* (8) *is*

(i) *if* $c \geq \bar{c}$, $\eta^* = T$ *and* $\tilde{p}^*_t = \max(\tilde{p}^*(0), \bar{F}^{-1}(\frac{c}{\lambda T}))$ *for* $t \in [0, T]$;
(ii) *if* $\underline{c} \leq c < \bar{c}$, $\eta^* = \frac{c - \beta\bar{c}}{\bar{c} - \beta\bar{c}} T$ *and* $\tilde{p}^*_t = \tilde{p}^*(h)$ *for* $t \in [0, T]$;
(iii) *if* $c < \underline{c}$, $\eta^* = \tau$. *Furthermore,*

(a) *when* $\kappa < K$, $\tilde{p}^*_t = \tilde{p}^*_{\kappa,\kappa+1}$ *for* $t \in [0, T - (\kappa+1)\tau]$ *and* $\tilde{p}^*_t = \tilde{p}^*_{\kappa,i}$ *for* $t \in [T - (i+1)\tau, T - i\tau]$ *and* $0 \leq i \leq \kappa$, *where* $(\tilde{p}^*_{\kappa,0}, \ldots, \tilde{p}^*_{\kappa,\kappa+1})$ *is the optimal solution for Problem* (13).
(b) *when* $\kappa = K$, $\tilde{p}^*_t = \tilde{p}^*_{\kappa,2i}$ *for* $t \in [\kappa\tau - i\tau, T - i\tau]$ *and* $0 \leq i \leq \kappa$, *and* $\tilde{p}^*_t = \tilde{p}^*_{\kappa,2i-1}$ *for* $t \in [T - i\tau, (\kappa+1)\tau - i\tau]$ *and* $1 \leq i \leq \kappa$, *where* $(\tilde{p}^*_{\kappa,0}, \ldots, \tilde{p}^*_{\kappa,2\kappa})$ *is the optimal solution for Problem* (14).

To explain the structure of the optimal policy of DPR, Proposition 6(i) shows that the seller uses only NMBG if the stock level is sufficiently high (i.e., $c \geq \bar{c}$). As the initial stock level decreases (i.e., $\underline{c} \leq c < \bar{c}$), the seller uses both sales modes (i.e., MBG and NMBG) with the same perceived price, as indicated in Proposition 6(ii). Moreover, the duration of MBG (i.e., $T - \eta^*$) decreases with the initial stock. If the stock becomes smaller (i.e., $c < \underline{c}$), the seller offers MBGs before $T - \tau$ so that returned products can be resold. To understand the sales process, Figure 3 illustrates the price processes and stock levels for different initial stocks and return times. If the initial stock is not very scarce, the optimal solutions of rDPR-0 and DPR coincide (i.e., $\kappa = 0$), which corresponds to Proposition 6(iii) a. As shown in Figure 3a and c, the seller first uses price $\tilde{p}^*_{0,1}$ during the MBG phase $[0, T - \tau]$ and then uses price $\tilde{p}^*_{0,0}$ during the NMBG phase $[T - \tau, T]$. Nevertheless, as returns come after time τ, the sales rate (the purchase rate minus the return rate) is not constant during the MBG (resp. NMBG) phase in Figure 3a (resp. Figure 3c). If the initial stock is very scarce, the optimal solution of DPR coincides with that of rDPR-1 rather than rDPR-0, that is, $\kappa = 1$. The solution for a short (resp. long) return time is characterized in Proposition 6(iii)a (resp. Proposition 6(iii)b) and illustrated by Figure 3b (resp. Figure 3d). For each scenario, the seller's optimal pricing policies consist of three pieces. Furthermore, the seller clears the initial inventory by the time $\max(T - \tau, \tau)$ (see Figure 3b and d), and then sells returned products by keeping zero inventory. Next we further characterize some structural properties of the optimal policy.

PROPOSITION 7. *For* $c < \underline{c}$,

(i) *if* $\kappa < K$, $\tilde{p}^*_{\kappa,i}$ *decreases in* i; *if* $\kappa = K$, *both* $\tilde{p}^*_{\kappa,2i}$ *and* $\tilde{p}^*_{\kappa,2i-1}$ *decrease in* i;
(ii) *the sales rate is decreasing in time during the whole selling season.*

Proposition 7(i) implies that the optimal perceived price $\tilde{p}^*_t$ is essentially increasing in time t, although the notion is in a periodic manner[5] for $\kappa = K$. This is consistent with the observation from the pricing policies in Figure 3. Proposition 7(ii) states that the sales rate decreases as time passes, which is clear from process of the stock level in Figure 3. Intuitively, due to the MBG policy, the seller is willing to sell more at a lower price (implying a higher sales rate) at the beginning as he has more opportunity to sell the returned items.

Figure 3 The Stock Level and Perceived Price for Different Return Times and Initial Inventories. We Set Arrival Rate $\lambda = 0.1$, Fit Probability $\beta = 0.6$, Hassle Cost $h = 3$, Sales Period $T = 500$, and Willingness-to-Pay Having a Gamma Distribution with Shape Parameter 15 and Scale Parameter 1 [Color figure can be viewed at wileyonlinelibrary.com]

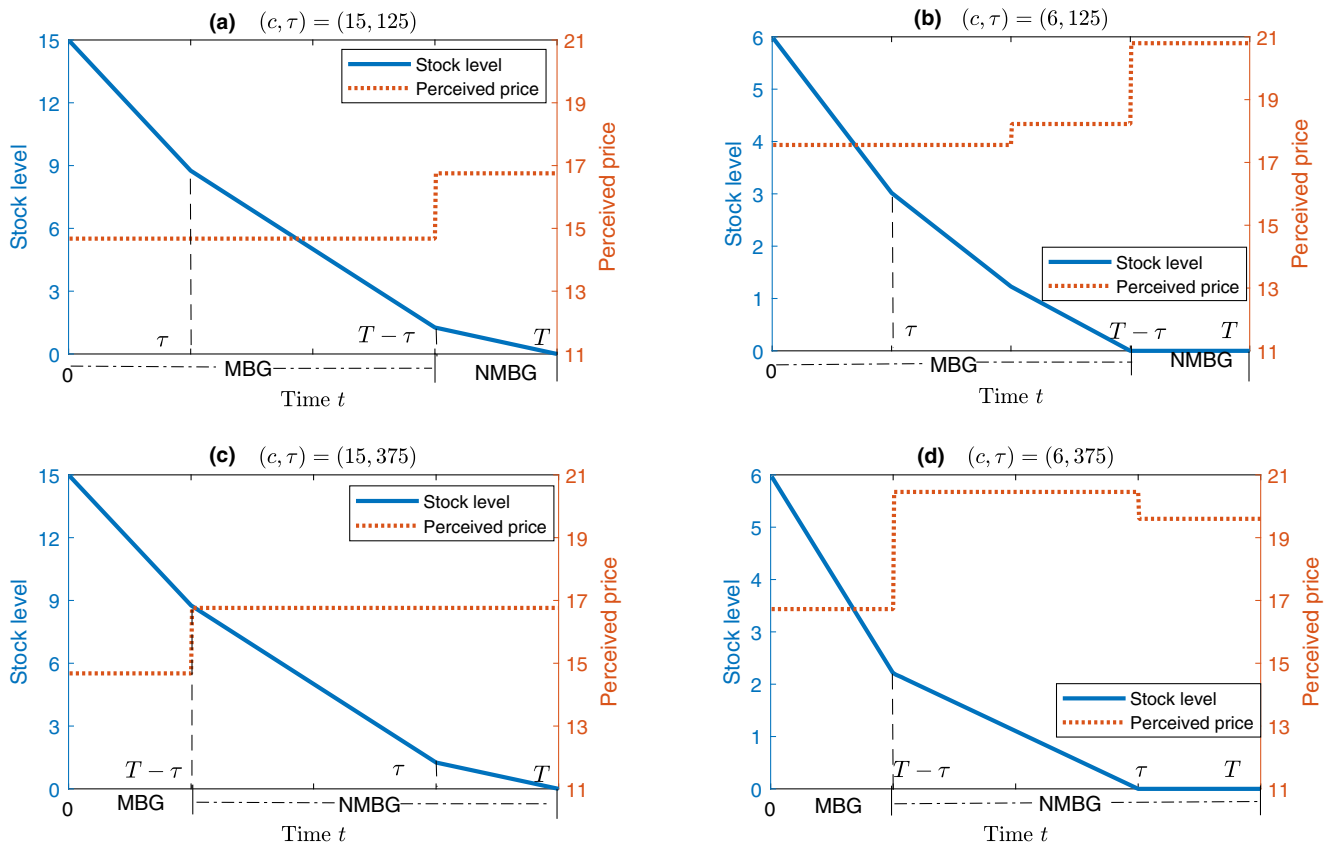

The optimal policy of DPR provides several insights. First, similar to the instant-return model in section 3, the lower the initial stock level, the longer the seller uses the MBG policy. Second, dynamic pricing is particularly valuable when free returns are allowed. As shown in Proposition 6(iii), it is optimal for the seller not to set a uniform price even during the MBG phase. This is a key difference from the dynamic pricing literature, in which the fluid model usually yields a constant price. The difference arises from the non-redundant inventory constraints for all $t \in [0, T]$. Third, the seller may keep zero inventory approaching the end of the season. The revenue is mainly generated from reselling returned items.

To conclude this section, we emphasize the applicability of Proposition 7. It seems that the increasing price trajectory is not consistent with the observation from the fashion retail industry, which motivates our research. This is because in practice, the consumer valuations are typically decreasing over the selling season, whereas we consider stationary consumer valuations. The decreasing valuation results in decreasing prices, which would counter the price increase caused by product returns analyzed in this study. Nevertheless, we hope this study has provided comprehensive analysis from the perspective of product returns, upon which the analysis of non-stationary valuations may be built in future research. Moreover, although Proposition 7 is derived for fixed return times, the key insight can be extended to industries with flexible return times. For instance, it provides another explanation for the common practice of raising prices over time in hotel bookings.

4.2. Heuristic and Bounds

The deterministic optimal solution specified by Proposition 6 can be used as a simple heuristic policy for the stochastic system. For instance, for $c < \underline{c}$ and $\kappa < K$, the seller offers MBGs with perceived price $\tilde{p}^*_{\kappa,\kappa+1}$ during $[0, T - (\kappa + 1)\tau]$, adjusts the perceived price to $\tilde{p}^*_{\kappa,i}$ to sell any remaining stock (including the returned items) during $[T - (i + 1)\tau, T - i\tau]$ for $i = \kappa, \ldots, 1$, and uses NMBG with $\tilde{p}^*_{\kappa,0}$ to sell all available items during $[T - \tau, T]$. This is referred to as the DPR heuristic. It is similar to the heuristics in previous research (e.g., Gallego and Van Ryzin 1994, Hu et al. 2017), but has distinctive

characteristics. Unlike the heuristic in Gallego and Van Ryzin (1994), the DPR heuristic needs not be stationary because there may be price changes during the sales horizon. It is also different from the switch heuristic in Hu et al. (2017), as the sales may depend on the underlying return process, which is a thinned Poisson process.

Let $V_H(T, c)$ denote the expected revenue of the stochastic system if the retailer uses the DPR heuristic with initial inventory c and sales season length T. We have

PROPOSITION 8. $V_H(T, c) \leq V_0(c, \mathbf{0}) \leq \bar{V}(T, c)$.

Recall that $V_0(c, \mathbf{0})$ is the optimal expected revenue for the stochastic system. As a sub-optimal policy, the heuristic generates less revenue than the optimal policy. In addition, Proposition 8 further shows that the performance of the stochastic system is dominated by its deterministic counterpart. This is commonly observed in the dynamic pricing literature. As the scale of the system becomes large, one may expect the standard result that the DPR heuristic for the stochastic system is asymptotically optimal. However, the asymptotic optimality in this problem presents unique technical challenges. Since the seller relies on the returned products to sell, the system resembles a queueing system. As a result, the analysis becomes intractable and we evaluate the seller's expected revenue under the heuristic by simulation.

4.3. Performance of the DPR Heuristic

In this section, we conduct numerical studies to show the performance of the DPR heuristic. In the numerical study, the willingness-to-pay V follows a Gamma distribution with mean $\mu = 15$ and standard deviation $\sigma = 5$. We fix the arrival rate $\lambda = 0.5$, time periods $T = 100$, fit probability $\beta = 0.7$, and hassle cost $h = 3$, and let the initial stock c vary from 1 to 40. To characterize the asymptotic performance of the DPR heuristic, we consider a sequence of problems with sales horizons $T^k = kT$, initial capacities $c^k = kc$, and return times $\tau^k = k\tau$, indexed by k. Since the optimal value of the fluid model provides an upper bound, we consider the relative performance $V_H(kT, kc)/\bar{V}(kT, kc)$, which provides a lower bound for the performance of the DPR heuristic in the stochastic system. We show how the relative performance varies with respect to the return time τ, the initial stock c, and the scaling factor k.

To evaluate $V_H(kT, kc)$, we simulate 1000 sample paths and apply the DPR heuristic. The realized revenues are averaged and the results are summarized in Table 2. We make the following observations. First, given return time τ and initial stock c, the relative performance increases as the scale of the system k increases. This implies that the DPR heuristic is asymptotically optimal if the scale of the system is sufficiently large, which is consistent with Gallego and Van Ryzin (1994) and Hu et al. (2017). Second, the relative performance approaches 100% as capacity c increases, which is similar to the standard findings in the revenue management literature. Third, for small stock levels, the relative performance first decreases and then increases as return time τ increases. If the return time is relatively long (resp. short), the retailer would implement NMBG (resp. MBG) for most of the sales season. This result indicates that the heuristic is more effective for these regimes.

Table 2 Relative Performance of the Heuristic to the Fluid Model (%)

		Initial stock (c)							
k	τ/T	5	10	15	20	25	30	35	40
1	0	82.0	87.6	90.0	91.3	93.6	93.3	93.4	94.4
	0.25	80.3	86.5	89.7	91.2	92.6	93.2	93.2	94.5
	0.50	77.6	85.1	89.0	90.6	92.1	93.3	93.4	94.6
	0.75	78.3	84.8	88.4	90.6	91.7	92.3	93.2	94.4
	1.00	82.9	87.4	90.0	91.4	91.9	92.5	93.1	94.7
10	0	94.4	95.9	96.8	97.2	97.9	97.9	97.9	98.8
	0.25	93.1	95.9	96.9	97.5	97.7	97.9	97.9	98.8
	0.50	92.5	95.8	96.7	97.3	97.6	97.8	97.9	98.7
	0.75	92.8	95.7	96.7	97.2	97.5	97.8	97.9	98.7
	1.00	94.3	95.9	96.7	97.3	97.5	97.6	97.9	98.8
100	0	98.2	98.8	99.0	99.1	99.4	99.3	99.3	99.9
	0.25	97.9	98.8	99.0	99.2	99.3	99.3	99.3	99.9
	0.50	97.8	98.8	99.0	99.2	99.3	99.3	99.3	99.9
	0.75	98.0	98.8	99.0	99.1	99.3	99.3	99.3	99.9
	1.00	98.2	98.8	98.9	99.1	99.2	99.3	99.3	99.9

5. Value of MBG

In this section, we identify the conditions under which implementing MBG can significantly increase the revenue. To that aim, we compare the performance when the retailer is flexible in using MBG and that without MBG. In particular, to evaluate the effects of the stock level, fit probability, and hassle cost, we focus on the instant-return model and measure the value of MBG by the percentage improvement $(V_F - V_N)/V_N$, where recall that V_F and V_N are the expected revenues for the stochastic flexible model and the NMBG-only model, respectively. To assess the impact of return time, we use the fluid model and measure the value of MBG by the percentage improvement $(\bar{V}(T, c) - \bar{V}_N(T, c))/\bar{V}_N(T, c)$. The baseline parameter setting is the same as that in section 4.3.

Effect of Stock Level. Figure 4 displays the expected revenues under MBG and NMBG, and the percentage improvement for different stock

Figure 4 Effect of Stock on Percentage Improvement of Using Money-Back Guarantees [Color figure can be viewed at wileyonlinelibrary.com]

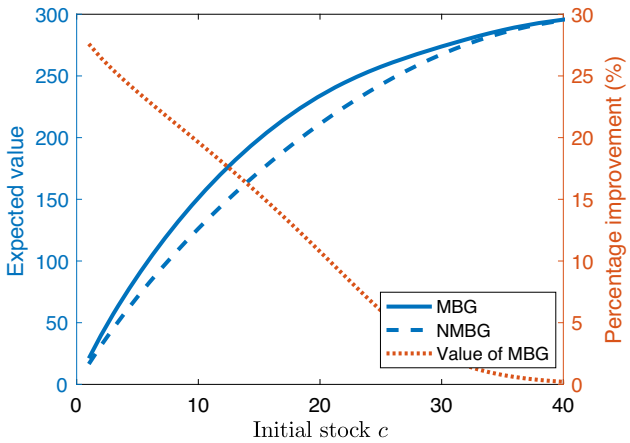

levels. As the stock level increases, the expected revenues for both MBG and NMBG increase, but the value of MBG decreases. This is consistent with Lemma 1: a small marginal value of inventory disincentivizes the retailer from using MBG. Therefore, MBG is less valuable to the retailer if his stock is relatively high.

Effects of Fit Probability and Hassle Cost. Table 3 presents the percentage improvement of using MBG for different values of stock, fit probability, and hassle cost. Interestingly, the value of MBG first increases and then decreases in the fit probability when the stock level is high. This is because the fit probability has two effects: when it is relatively small, the buyer returns a large portion of items. Hence the stock level becomes relatively high to the seller, and the value of MBG decreases. When the fit probability is relatively high, the product is more likely to be a good match and hence the seller does not need MBG to resell the product. Higher hassle costs decrease the value of MBG. This is easy to understand because the hassle cost increases the buyer's return cost and makes MBG less valuable. Overall, the percentage improvement can be rather significant, achieving as high as 54.1%.

Table 3 Percentage Improvement (%) for Different Parameters

	Fit probability β					Hassle cost h				
c	0.5	0.6	0.7	0.8	0.9	1	2	3	4	5
1	54.1	40.0	27.6	17.0	7.8	32.7	30.1	27.6	25.1	22.5
5	42.4	33.2	23.7	14.9	7.0	29.6	26.7	23.7	20.8	17.9
10	30.3	26.1	19.6	12.8	6.1	26.2	22.9	19.6	16.4	13.2
20	9.1	11.3	10.8	8.1	4.3	18.0	14.3	10.8	7.4	4.5
30	1.3	1.9	2.3	2.5	1.7	7.2	4.4	2.3	1.1	0.5
40	0.1	0.2	0.2	0.2	0.1	0.8	0.4	0.2	0.1	0.0

Figure 5 Effect of Return Time (τ) on Percentage Improvement of Using Money-Back Guarantees [Color figure can be viewed at wileyonlinelibrary.com]

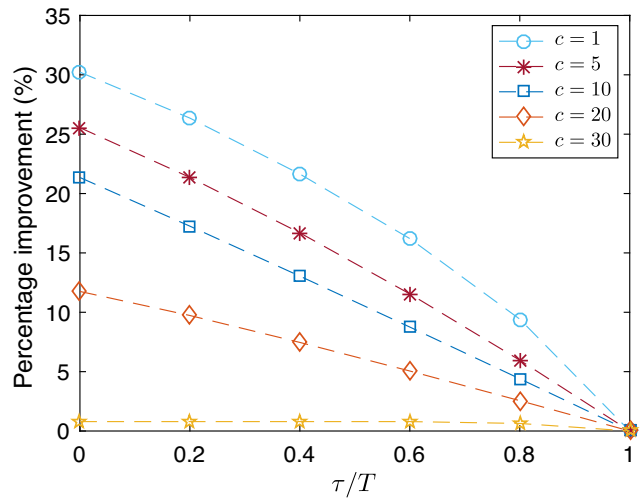

Effect of Return Time. Figure 5 shows the percentage improvement of using MBG for different return times. As per this figure, the value of MBG decreases in the return time. Due to the limited length of the sales season, a longer return time decreases the resale value of the returned products, thus reducing the value of MBG. Moreover, the value of MBG decreases in the stock level, which is consistent with Figure 4.

6. Extensions

In this section, we consider two extensions to the model in order to incorporate a few realistic features: first, we consider customers returning products after a random period of time; second, the resale of products may have a processing cost.

6.1. Random Return Times

Another approach to modeling the return time is to use a random variable that is independent for each customer. If the random variable has a geometric distribution, then the memoryless property keeps the dynamic programming formulation tractable. This is the approach taken by, for example, Aydin et al. (2017) and Hu et al. (2019). In practice, the firm can choose between the two approaches (exponential or fixed return times) based on which better reflects the behavioral pattern of consumer returns displayed in the data.

Next we consider this alternative formulation for our problem. Suppose each customer independently chooses to keep the product, return it due to misfitting, or delay her decision with probabilities q_k, q_r, and q_d in each period until the end of the selling horizon. To link to the base model, we calibrate the

probabilities so that, *on average*, a customer is satisfied with the product with probability β and if not, returns it after τ periods. Equivalently, we have $\tau = 1/(1-q_d)$ and $\beta = q_k + q_k q_d + q_k q_d^2 + \cdots$. As a result, similar to section 3, the expected revenue obtained from a consumer is βp. Note that for given T, β has only finite terms, and the approximation is not exact. The error is typically negligible for the range of sales horizon in practice. Once we calibrate q_d and q_k in this fashion, let $q_r = 1 - q_d - q_k$.

We keep track of the number of tentative buyers who may return the product in each period, denoted as state y. As it proceeds from t to period $t+1$, a random number $\mathbf{M}_k(y)$ among the y customers would keep the product. Similarly, $\mathbf{M}_r(y)$ ($\mathbf{M}_d(y)$) would return the product (postpone the decision). We have $\mathbf{M}_k(y) + \mathbf{M}_r(y) + \mathbf{M}_d(y) = y$, and they have a multinomial distribution with parameters q_k, q_r, and q_d. As for the dynamic programming, let n and y be the (remaining) on-hand inventory and tentative buyers in period t, respectively. Then the Bellman equation is given by

$$\begin{aligned} V_t(n,y) = \max\Big\{ & \max_p \{\lambda \bar{F}(p/\beta)[p + \mathbb{E}V_{t+1}(n-1+\mathbf{M}_r(y),\ \mathbf{M}_d(y))] \\ & + [1-\lambda\bar{F}(p/\beta)]\mathbb{E}V_{t+1}(n+\mathbf{M}_r(y),\ \mathbf{M}_d(y))\}, \\ & \max_p \{\lambda \bar{F}(p+(1-\beta)h/\beta)[\beta p + \mathbb{E}V_{t+1}(n-1+\mathbf{M}_r(y),\ 1+\mathbf{M}_d(y))] \\ & + [1-\lambda\bar{F}(p+(1-\beta)h/\beta)]\mathbb{E}V_{t+1}(n+\mathbf{M}_r(y),\ \mathbf{M}_d(y))\}\Big\} \end{aligned} \tag{15}$$

for $n \geq 1$ and $V_t(0,\ y) = \mathbb{E}V_{t+1}(\mathbf{M}_r(y),\ \mathbf{M}_d(y))$ with boundary conditions $V_{T+1}(n,\ y) = 0$ for any $(n,\ y)$, where the expectation is taken with respect to the multinomial distribution with parameters $(q_k,\ q_r,\ q_d,\ y)$. In this formulation, if there is no purchase at t, then the state of the system becomes $(n + \mathbf{M}_r(y),\ \mathbf{M}_d(y))$ at $t+1$. On the other hand, if there is a purchase at t, the state of the system becomes $(n-1+\mathbf{M}_r(y),\ 1+\mathbf{M}_d(y))$ or $(n-1+\mathbf{M}_r(y),\ \mathbf{M}_d(y))$, depending on whether MBG is offered or not. After some algebraic manipulation, we have

$$\begin{aligned} V_t(n,y) = {} & \mathbb{E}V_{t+1}(n+\mathbf{M}_r(y),\ \mathbf{M}_d(y)) \\ & + \lambda \max\Big\{ \max_{\tilde{p}} \bar{F}(\tilde{p})\big[\beta\tilde{p} - \Delta_{t+1}^N(n,\ y)\big], \\ & \max_{\tilde{p}} \bar{F}(\tilde{p})\big[\beta\tilde{p} - (1-\beta)h - \Delta_{t+1}^M(n,\ y)\big]\Big\} \end{aligned}$$

where $\Delta_{t+1}^N(n,\ y) = \mathbb{E}V_{t+1}(n+\mathbf{M}_r(y),\ \mathbf{M}_d(y)) - \mathbb{E}V_{t+1}(n-1+\mathbf{M}_r(y),\ \mathbf{M}_d(y))$ and $\Delta_{t+1}^M(n,\ y) = \mathbb{E}V_{t+1}(n+\mathbf{M}_r(y),\ \mathbf{M}_d(y)) - \mathbb{E}V_{t+1}(n-1+\mathbf{M}_r(y),\ 1+\mathbf{M}_d(y))$ are the opportunity costs for selling one unit of inventory with NMBG and MBG policies, respectively. Accordingly, the retailer's optimal policy is given as follows.

Lemma 6. *If $\Delta_{t+1}^N(n,\ y) > (1-\beta)h + \Delta_{t+1}^M(n,\ y)$, it is optimal to offer MBGs with the perceived price $\tilde{p}^*((1-\beta)h + \Delta_{t+1}^M(n,\ y))$; otherwise, it is optimal to use NMBG with the perceived price $\tilde{p}^*(\Delta_{t+1}^N(n,\ y))$.*

Lemma 6 extends Lemma 1 to the case of random return times. Notice that the dynamic model in (15) also serves as an approximate dynamic model (ADM) for our main model in (5). As shown in Aydin et al. (2017), one can use its optimal policy as a heuristic solution for the main model. That is, we solve the dynamic program in (15) to obtain the optimal policy. Then, we use the derived policy as our MBG and pricing policy for the model in (5). It is interesting to compare this ADM heuristic with the DPR heuristic in section 4.2. To that aim, we numerically study the relative performance of the two heuristics to the fluid model. The parameter settings are the same as section 4.3 except that $T = 50$, such that the calculation time for the ADM heuristic is not extremely long. Table 4 shows the relative performance and the corresponding calculation time (in the bracket) for

Table 4 Relative Performance of Heuristics to the Fluid Model (%) and the Computation Time (sec)

		Scaling factor $k = 1$		Scaling factor $k = 5$	
τ	c	ADM	DPR	ADM	DPR
5	5	88.5 (0.7)	82.0 (0.1)	96.3 (385.1)	91.8 (0.1)
	10	90.4 (4.1)	87.6 (0.0)	96.8 (4932.7)	94.8 (0.1)
	15	94.1 (14.6)	90.6 (0.0)	98.1 (23365.9)	95.6 (0.0)
10	5	87.8 (0.6)	81.4 (0.1)	95.8 (380.8)	91.4 (0.1)
	10	90.4 (3.9)	87.7 (0.1)	96.2 (4899.2)	94.8 (0.1)
	15	93.2 (14.2)	90.4 (0.0)	98.0 (24053.4)	95.7 (0.0)
15	5	86.8 (0.6)	80.3 (0.1)	95.0 (381.9)	91.4 (0.2)
	10	90.0 (4.1)	87.5 (0.2)	96.1 (4997.2)	94.7 (0.1)
	15	93.1 (14.2)	90.5 (0.0)	97.6 (25292.7)	95.9 (0.0)

obtaining the heuristic solution. It indicates that the effects of the inventory level and the scale of the system on the two heuristics are similar. Moreover, it seems that the ADM heuristic always outperforms the DPR heuristic. As in Aydin et al. (2017), this is because the ADM heuristic makes use of the dynamic information of the system. To capture this dynamic information, the calculation time for the ADM heuristic increases dramatically in the stock level and the number of selling periods. Nevertheless, the calculation time for the DPR heuristic is independent of the scale of the system and is almost negligible. More importantly, the performance gap between the two heuristics shrinks as the stock level and the scale of the system increase.

6.2. Processing Cost

We consider a fixed cost g if the seller processes a returned item for resale. We first consider the instant-return model and then investigate positive return times.

6.2.1. Instant Returns. The model setup and assumptions are identical to those detailed in section 3. Let $V_c(t,\ n)$ be the optimal expected revenue-to-go in period t with n units of remaining inventory. The optimization problem follows from the Bellman equation:

$$V_c(t,\ n) = \max\Big\{\max_p\{\lambda\bar{F}(p/\beta)[p + V_c(t+1,\ n-1)] + [1-\lambda\bar{F}(p/\beta)]V_c(t+1,\ n)\},\ V_c^m(t,\ n)\Big\} \tag{16}$$

where

$$V_c^m(t,n) = \max_p\big\{\lambda(1-\beta)\bar{F}(p+(1-\beta)h/\beta)\max\{V_c(t+1,n)-g,\ V_c(t+1,n-1)\} + \lambda\beta\bar{F}(p+(1-\beta)h/\beta)[p+V_c(t+1,n-1)] + [1-\lambda\bar{F}(p+(1-\beta)h/\beta)]V_c(t+1,n)\big\}$$

with boundary conditions $V_c(t,\ 0) = 0$ for $t = 1, \ldots, T$ and $V_c(T+1,\ n) = 0$ for all $n \le c$. In the outer bracket, the optimal value function takes the maximum of (i) the optimal revenue without MBG, which is the same as that in the base model, and (ii) the optimal revenue with MBG, $V_c^m(t,\ n)$. In (ii), the first term represents a customer purchasing but returning the product, and then the retailer considering whether to restock the returned inventory. The second term represents a customer purchasing and keeping the product, and the third term represents no product being sold in that period.

For any $(t,\ n)$, let $\Delta V_c(t,\ n) = V_c(t,\ n) - V_c(t,\ n-1)$ be the marginal value of inventory. Rearranging Equation (16), we obtain

$$V_c(t,\ n) = \lambda\max\{R_N(\Delta V_c(t+1,\ n)),\ R_M^c(\Delta V_c(t+1,\ n))\} + V_c(t+1,\ n),$$

where $R_N(\Delta) = \max_p \bar{F}(p/\beta)(p-\Delta)$ (see section 3.2) and $R_M^c(\Delta) \triangleq \max_p \bar{F}(p\ +(1-\beta)h/\beta)[\beta(p-\Delta)-(1-\beta)\min(g,\ \Delta)]$. Here, $R_M^c(\Delta)$ is the retailer's expected additional gain by offering MBG when the opportunity cost for selling the unit is Δ and the processing cost is g. After transforming to the perceived price $\tilde{p}$, we have $R_N(\Delta) = \max_{\tilde{p}} \bar{F}(\tilde{p})(\beta\tilde{p}-\Delta)$ and $R_M^c(\Delta) = \max_{\tilde{p}} \bar{F}(\tilde{p})[\beta\tilde{p}-(1-\beta)(h+\min(g,\ \Delta)) - \beta\Delta]$. The structure of the optimal pricing policy for Problem (16) is largely parallel to the base model without costs.

Proposition 9.

(i) *For any $(t,\ n)$, if $\Delta V_c(t+1,n) > h+g$, it is optimal to offer an MBG with the perceived price $\tilde{p}^*(\beta\Delta V_c(t+1,\ n)+(1-\beta)(h+g))$; otherwise, it is optimal to use NMBG with the perceived price $\tilde{p}^*(\Delta V_c(t+1,\ n))$.*

(ii) *Given n, it is optimal for the retailer to offer an MBG if and only if $t < t_n^c$, where $t_n^c \triangleq \min\{t : \Delta V_c(t+1,\ n) \le h\}$. Moreover, t_n^c is decreasing in n.*

Proposition 9(i) shows that the retailer's optimal policy has the same structure as the base model in Lemma 1, whereas replacing the cost h with its counterpart $h+g$. Intuitively, the retailer allows the return of a product only when he plans to restock this return, under which case the customer incurs hassle cost h and the retailer incurs processing cost g. As a result, for a given inventory level, the return policy is also analogous to that for the base case in Proposition 1.

6.2.2. Positive Return Times. We formulate the dynamic programming similar to that in section 4. Given the time and on-hand inventory $(t,\ n)$ and the MBG-sales history $\mathbf{s}^t$, let $V_t^c(n,\ \mathbf{s}^t)$ be the optimal expected revenue. By the Bellman equation, we have

$$V_t^c(n,\ \mathbf{s}^t) = \max\Big\{\max_p\{\lambda\bar{F}(p/\beta)[p + U(t+1, n, \mathbf{s}^t, 1, 0)] + [1 - \lambda\bar{F}(p/\beta)]U(t+1, n, \mathbf{s}^t, 0, 0)\},$$
$$\max_p\{\lambda\bar{F}(p + (1-\beta)h/\beta)[\beta p + U(t+1, n, \mathbf{s}^t, 1, 1)]$$
$$+[1 - \lambda\bar{F}(p + (1-\beta)h/\beta)]U(t+1, n, \mathbf{s}^t, 0, 1)\}\Big\} \tag{17}$$

for $n \geq 1$ and

$$V_t^c(0, \mathbf{s}^t) = s_\tau^t(1-\beta)\max\{V_{t+1}^c(1, \Phi(\mathbf{s}^t, 0)) - g, V_{t+1}^c(0, \Phi(\mathbf{s}^t, 0))\} + (1 - s_\tau^t(1-\beta))V_{t+1}^c(0, \Phi(\mathbf{s}^t, 0)),$$

with terminal conditions $V_{T+1}^c(n, \mathbf{s}^{T+1}) = 0$ for any $(n, \mathbf{s}^{T+1})$, where $U(t+1, n, \mathbf{s}^t, z, m) = s_\tau^t(1-\beta)\max\{V_{t+1}^c(n-z+1, \Phi(\mathbf{s}^t, z\cdot m)) - g, V_{t+1}^c(n-z, \Phi(\mathbf{s}^t, z\cdot m))\} + [1 - s_\tau^t(1-\beta)]V_{t+1}^c(n-z, \Phi(\mathbf{s}^t, z\cdot m))$ is the expected revenue regrading to the event $\mathbf{B}(s_\tau^t(1-\beta))$ with the sale and return policy information (z, m) at time t. Here $z \in \{0, 1\}$ represents the sale information such that $z = 1$ represents there is a sale at time t and otherwise $z = 0$, and $m \in \{0, 1\}$ represents the associated return policy such that $m = 1$ represents MBG and $z = 0$ means NMBG. For the term $U(t+1, n, \mathbf{s}^t, z, m)$, if there is a return, the retailer decides whether to restock the returned inventory by choosing the one leading to a higher revenue—that is, $\max\{V_{t+1}^c(n-z+1, \Phi(\mathbf{s}^t, z\cdot m)) - g, V_{t+1}^c(n-z, \Phi(\mathbf{s}^t, z\cdot m))\}$. As in section 4, Problem (17) is analytically intractable because of the curse of dimensionality, we then study the fluid approximation to derive heuristics for the problem.

The model setup and assumptions are identical to those detailed in section 4.1, but there is a processing cost if the retailer restocks the returned inventory. As the DPR problem in (6), the seller's fluid model with return time τ and processing cost g can be formulated as

$$\text{(DPC)}\quad \bar{V}_c(T, c) = \max_{m_t \in \{-1;0;1\},\ \tilde{p}_t}\Big[\int_0^T \lambda I_{\{m_t=1\}}\{\bar{F}(\tilde{p}_t)[\beta\tilde{p}_t - (1-\beta)h] - \bar{F}(\tilde{p}_t)(1-\beta)g\}$$
$$+\lambda I_{\{m_t=0\}}\beta\tilde{p}_t\bar{F}(\tilde{p}_t) + \lambda I_{\{m_t=-1\}}\bar{F}(\tilde{p}_t)[\beta\tilde{p}_t - (1-\beta)h]dt\Big]$$
$$\text{subject to}\int_0^s \lambda\bar{F}(\tilde{p}_t)dt - \int_0^{(s-\tau)^+}\lambda I_{\{m_t=1\}}(1-\beta)\bar{F}(\tilde{p}_t)dt \leq c, \qquad \forall 0 \leq s \leq T, \tag{18}$$

where $I_{\{\cdot\}}$ is the indicator function. The ternary decision m_t captures the return policy and associated restocking policy (if any returned inventory) such that $m_t = 1$ represents MBG and restocking the returned inventory, $m_t = 0$ represents NMBG, and $m_t = -1$ represents MBG but not restocking the returned inventory. If the retailer uses MBG and restocks the partially returned inventory ($m_t = 1$) after τ, he incurs a unit processing cost g. In Problem (18), "DPC" stands for the (D)eterministic formulation with (P)ositive return time and processing (C)ost.

For the DPC problem in (18), it is easy to see that the MBG but no-restocking policy ($m_t = -1$) is dominated by the NMBG policy, since NMBG generates more revenue than the MBG but no-restocking policy if they use the same price. Intuitively, the retailer using MBGs aims to resell the returned inventory to obtain more revenue, but the no-restocking policy precludes this possibility. Consequently, the retailer never uses the MBG but no-restocking policy and hence the DPC problem becomes the DPR problem in (6), where the cost now includes both the hassle cost and the processing cost—that is, $h + g$.

Proposition 10. *The DPC problem* (18) *is equivalent to the DPR problem in* (6) *by replacing h with $h + g$.*

7. Conclusions

In this study, we consider a firm's joint MBG and pricing policy when consumers can return their unsatisfactory products. If the misfit products are returned instantly, we find the optimal pricing and MBG strategy is a threshold policy in which MBG is offered if the remaining sales season is sufficiently long. We further add a new monotonicity result to the revenue management literature by showing that the marginal value of inventory is increasing in the fit probability.

To address the problem with positive return times, we focus on the underlying fluid model (DPR). The firm needs to decide the joint MBG and pricing policy for the whole sales season, accounting for both returns and purchases. Thus, the fluid model is an inventory control problem involving two decision variables and an infinite number of constraints. We show that the optimal policy has a simple structure: first MBGs are offered and then no MBGs are allowed. This reduces the DPR problem to deciding how long to offer MBGs and the pricing policy. We develop an iterative approach that is guaranteed to solve the fluid model in finite iterations. In particular, we consider a series of associated relaxed problems that have a nested structure. We use the Lagrangian approach to explicitly solve these relaxed problems and show that the optimal solution for DPR coincides with one of the relaxations. Our numerical analysis implies that the heuristic derived from the fluid model is asymptotically optimal when the scale of the system is large enough.

Our analytical and computational results generate many managerial insights. First, the firm may sell all its products in the middle of the season and then rely solely on returned items to earn profits in the rest of the season. Second, the less inventory the firm has, the longer it should implement the MBG policy. Third, MBGs increase revenue significantly when the inventory is low, when the product does not fit the consumers very well, or when the return time is relatively short.

Our modeling framework and analysis lead to several important extensions for future development. First, we assume that the return time is exogenous; one could endogenize the return (i.e., cancelation) option by studying the "cancel by certain date" problem for the hotel industry. Second, we test the asymptotic property for the DPR heuristic through numerical simulation; it is important to analytically characterize the underlying system. Third, the fit probability is the same for all customers and independent of the price. It would be interesting to study the impact of heterogeneous and price-dependent fit probability. Finally, we assume there is a single channel for the firm to sell the product; it is important to consider the revenue management problem with product return for omnichannel retailing.

Acknowledgments

The authors thank the department editor (Dan Zhang), the senior editor, and the two anonymous reviewers for their valuable comments that have significantly improved the study. Yan Liu is partially supported by the National Natural Science Foundation of China [Grants 71801209, 71921001, and 71731010]. The research of Ningyuan Chen is partially supported by the NSERC Discovery Grants Program RGPIN-2020-04038.

Notes

[1]https://www.statista.com/forecasts/758683/revenue-of-the-apparel-market-worldwide-by-country (accessed date August 18, 2021).

[2]The implicit assumption is $p \geq h$, as the refund must exceed the cost of returning. As this condition does hold for the optimal price with MBG, we do not explicitly include it for ease of presentation.

[3]Consider an example where $V(T, 1) > h > 0$. Hence $V(T-1, 1)$ is strictly decreasing in h. Moreover, as $\Delta V(T, 2) = 0$, $V(T-1, 2)$ is independent of h. Thus, $\Delta V(T-1, 2) = V(T-1, 2) - V(T-1, 1)$ is strictly increasing in h.

[4]For $k = K$, in addition to the inventory constraints at $T - i\tau$ for $i = 0, \ldots, K-1$, we further impose inventory constraints at $i\tau$ for $i = 1, \ldots, K$. As discussed above, if the optimal policy for rDPR-$(K-1)$ is not feasible, that is, it leads to a higher return rate than the purchase rate during $[\tau, T-(K-1)\tau)$, then we add an additional constraint at τ to guarantee non-negative inventory during $[\tau, T-(K-1)\tau]$. The additional constraint can be shown to be binding for the optimal solution of rDPR-K (otherwise, the optimal policy for rDPR-$(K-1)$ would be feasible for DPR), which implies that the firm depletes the inventory at τ. The products sold during $[0, \tau]$ are returned during $[\tau, 2\tau]$, requiring an additional inventory constraint at 2τ. Recursively, rDPR-K requires inventory constraints at $i\tau$ for all $i = 1, \ldots, K$.

[5]It is easily verified that $\tilde{p}^*_{\kappa,2i+1} > \tilde{p}^*_{\kappa,2i}$ for all $0 \leq i \leq \kappa - 1$, hence $\tilde{p}^*_t$ is decreasing within period $[T-(i+1)\tau, T-i\tau]$.

References

Akcay, Y., T. Boyaci, D. Zhang. 2013. Selling with money-back guarantees: The impact on prices, quantities, and retail profitability. *Prod. Oper. Manag.* **22**(4): 777–791.

Altug, M. S. 2012. Optimal dynamic return management of fixed inventories. *J. Revenue Pricing Manag.* **11**(6): 569–595.

Aydin, G., S. Ziya. 2008. Pricing promotional products under upselling. *Manuf. Serv. Oper. Manag.* **10**(3): 360–376.

Aydin, N., Ş. İ. Birbil, H. Topaloğlu. 2017. Delayed purchase options in single-leg revenue management. *Transp. Sci.* **51**(4): 1031–1045.

Bell, D. R., S. Gallino, A. Moreno. 2018. Offline showrooms in omnichannel retail: Demand and operational benefits. *Management Sci.* **64**(4): 1629–1651.

Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manuf. Serv. Oper. Manag.* **5**(3): 203–229.

Bitran, G. R., S. V. Mondschein. 1997. Periodic pricing of seasonal products in retailing. *Management Sci.* **43**(1): 64–79.

Cohen, M. C., R. Lobel, G. Perakis. 2018. Dynamic pricing through data sampling. *Prod. Oper. Manag.* **27**(6): 1074–1088.

Dai, J., A. J. Kleywegt, Y. Xiao. 2019. Network revenue management with cancellations and no-shows. *Prod. Oper. Manag.* **28**(2): 292–318.

Davis, S., E. Gerstner, M. Hagerty. 1995. Money back guarantees in retailing: Matching products to consumer tastes. *J. Retail.* **71**(1): 7–22.

Debo, L. G., L. B. Toktay, L. N. Van Wassenhove. 2005. Market segmentation and product technology selection for remanufacturable products. *Management Sci.* **51**(8): 1193–1205.

Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overviews, current practices, and future directions. *Management Sci.* **49**(10): 1287–1309.

Ferrer, G., J. M. Swaminathan. 2006. Managing new and remanufactured products. *Management Sci.* **52**(1): 15–26.

Gallego, G., G. Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8): 999–1020.

Gallego, G., M. Z. Li, Y. Liu. 2020. Dynamic nonlinear pricing of inventories over finite sales horizons. *Oper. Res.* **68**(3): 655–670.

Gallino, S., A. Moreno. 2018. The value of fit information in online retail: Evidence from a randomized field experiment. *Manuf. Serv. Oper. Manag.* **20**(4): 767–787.

Gao, F., X. Su. 2016. Online and offline information for omnichannel retailing. *Manuf. Serv. Oper. Manag.* **19**(1): 84–98.

Hu, S., X. Hu, Q. Ye. 2017. Optimal rebate strategies under dynamic pricing. *Oper. Res.* **65**(6): 1429–1731.

Hu, X., Z. Wan, N. N. Murthy. 2019. Dynamic pricing of limited inventories with product returns. *Manuf. Serv. Oper. Manag.* **21**(3): 501–518.

Kuo, C.-W., H.-S. Ahn, G. Aydin. 2011. Dynamic pricing of limited inventories when customers negotiate. *Oper. Res.* **59**(4): 882–897.

Li, M., Y. Liu. 2021. Beneficial product returns in supply chains. *Prod. Oper. Manag.* Forthcoming.

Mann, D. P., J. P. Wissink. 1988. Money-back contracts with double moral hazard. *RAND J. Econ.* **19**(2): 285–292.

McWilliams, B. 2012. Money-back guarantees: Helping the low-quality retailer. *Management Sci.* **58**(8): 1521–1524.

Moorthy, S., K. Srinivasan. 1995. Signaling quality with a money-back guarantee: The role of transaction costs. *Market. Sci.* **14**(4): 442–466.

Mostard, J., R. Teunter. 2006. The newsboy problem with resalable returns: A single period model and case study. *Eur. J. Oper. Res.* **169**(1): 81–96.

Orendorff, A. 2019. The Plague of Ecommerce Return Rates and How to Maintain Profitability. Available at https://www.shopify.com/enterprise/ecommerce-returns (accessed date August 23, 2021).

Shen, Z.-J. M., X. Su. 2009. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Prod. Oper. Manag.* **16**(6): 713–728.

Shulman, J. D., A. T. Coughlan, R. C. Savaskan. 2009. Optimal restocking fees and information provision in an integrated demand-supply model of product returns. *Manuf. Serv. Oper. Manag.* **11**(4): 577–594.

Song, B., M. Z. Li. 2018. Dynamic pricing with service unbundling. *Prod. Oper. Manag.* **27**(7): 1334–1354.

Su, X. 2009. Consumer return policies and supply chain performance. *Manuf. Serv. Oper. Manag.* **11**(4): 595–612.

Subramanian, J., S. J. Stidham, C. J. Lautenbacher. 1999. Airline yield management with overbooking, cancellations, and no-shows. *Transp. Sci.* **33**(2): 147–167.

Sun, M. R., K. Tyagi. 2020. Product fit uncertainty and information provision in a distribution channel. *Prod. Oper. Manag.* **29**(10): 2381–2402.

Xiao, Y., S. X. Zhou. 2020. Trade-in for cash or for upgrade? Dynamic pricing with customer choice. *Prod. Oper. Manag.* **29**(4): 856–881.

Xie, J., E. Gerstner. 2007. Service escape: Profiting from customer cancellations. *Market. Sci.* **26**(1): 18–30.

Yan, X., X. Chao, Y. Lu, S. X. Zhou. 2017. Optimal policies for selling new and remanufactured products. *Prod. Oper. Manag.* **26**(9): 1746–1759.

Zhang, D., W. L. Cooper. 2005. Revenue management for parallel flights with customer-choice behavior. *Oper. Res.* **53**(3): 415–431.

Ziya, S., H. Ayhan, R. D. Foley. 2004. Relationships among three assumptions in revenue management. *Oper. Res.* **52**(5): 804–809.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix A: Notations and Proofs.