Optimal periodic flexible policies for two-stage serial supply chains

Liu, Fang; Gavirneni, Srinagesh

2018

Liu, F., & Gavirneni, S. (2018). Optimal periodic flexible policies for two-stage serial supply chains. *European Journal of Operational Research*, 268(1), 148-160.
doi:10.1016/j.ejor.2018.01.004

<https://hdl.handle.net/10356/90102>

<https://doi.org/10.1016/j.ejor.2018.01.004>

© 2018 Elsevier B.V. All rights reserved. This paper was published in *European Journal of Operational Research* and is made available with permission of Elsevier B.V.

Downloaded on 26 Apr 2023 23:20:37 SGT

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Optimal periodic flexible policies for two-stage serial supply chains

Fang Liu^{a,*}, Srinagesh Gavirneni^b^a Nanyang Business School, Nanyang Technological University, Singapore^b Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca, New York 14853, USA

ARTICLE INFO

Article history:

Received 23 May 2017

Accepted 4 January 2018

Available online xxx

Keywords:

Inventory

Periodic flexibility

Dynamic inventory system

Calculus of variations

ABSTRACT

In a two-stage serial supply chain, a periodic flexible policy (PF policy) allows the retailer to receive fixed orders that may depend on demand history in one period of the ordering cycle and order freely in other periods. Existing literature has shown that certain PF policies can significantly reduce the inefficiency in a decentralized supply chain. However, these works have mostly defined and implemented ad-hoc periodic flexible policies and have not attempted to identify the optimal periodic flexible policies. In this paper, we characterize the structure of the optimal PF policy using calculus of variations. In particular, we show that under the optimal PF policy, the retailer receives shipments either according to a state dependent capacitated policy or a state dependent order up to policy. Furthermore, we can approximate the retailer's optimal restricted ordering function by a piecewise linear function and show numerically that this approximation is near optimal.

© 2018 Published by Elsevier B.V.

1. Introduction

Periodic Flexibility (PF), a strategy under which retailers' orders are intermittently constrained, is commonly found in the \$670 billion meat and food distribution industry, a major part of the U.S. economy. Griller's Pride, a meat and poultry supplier in Atlanta, uses a "Scheduled Day-of-the-Week Deliveries" strategy in which they deliver to each zip code in the Atlanta area on a specific day of the week. On a larger scale, the Houston-based company Sysco ships 21.5 million tons of produce, meats, prepared meals, and other food-related products annually. It offers 400,000 products to a third of American eating establishments including restaurants, cafeterias, and sports stadiums, which resulted in \$37 billion in sales in 2010. Many of these customers receive shipments at regular intervals with one or more deliveries per week. Given the regularity with which these deliveries occur, many suppliers facilitate the distribution process by providing the option to specify a standing order – a combination of products that are received on a consistent schedule. Winder Farms in Seattle, Washington allows their customers to define standing orders that can be delivered every week, every other week, or every month. Many organizations allow these standing orders to be modified frequently at no cost, although some organizations such as Victoria Organic Delivery do charge for changes. For organizations like VOD, customers must restrict the number of times they change these standing orders.

Standing orders may not be the same from one delivery to another. "Maybe Monday is a different quantity than a Wednesday or Friday" according to a small milk distributor in Hawaii. However, according to Zhu, Gavirneni, Kapuscinski (2010) and Chen and Gavirneni (2010), the optimal cycle length of the periodical flexible policies is most often two and never more than five. Thus, it is sufficient to study periodic flexible policies with a short cycle length.

We focus on the total supply chain cost as the primary measure of performance and, in doing so, develop strategies that will significantly reduce the cost of the supply chain. A strategy of centralized control will, of course, dominate any decentralized policy, but is usually not preferred in practice because many organizations do not feel inclined to give up complete control of their inventory management. A local restaurateur we interviewed mentioned that their day-to-day demand profile has so much qualitative information that it is almost impossible to accurately convey it to an outsider. He also mentioned that, while retailers would prefer to have complete flexibility in order to respond to demand changes, they also monitor costs very closely and are, as such, open to giving up some flexibility in order to reduce their cost burden. There is also evidence that suggests that Vendor Managed Inventory (VMI), an approach that many firms have adopted to operationalize centralized control, has resulted in many more problems than it has solved. A recent survey by Gatepoint Research illustrated that "Though there's an increased interest in inventory management programs, such as VMI, that aim to reduce costs and improve customer service levels, companies continue to face challenges in actually establishing these programs." They

* Corresponding author.

E-mail addresses: liu_fang@ntu.edu.sg (F. Liu), nagesh@cornell.edu (S. Gavirneni).

identified that “visibility into time-sensitive supply chain data” is a major challenge, with over half of the survey respondents struggling with this issue. Niranjana, Wagner, and Thakur-Weigold (2011) highlight the very mixed results from various VMI implementations and through the analysis of various situations to evaluate how specific features of the company, their products, and the supply chain played a vital role in determining VMI success. They developed a questionnaire that can aid companies in evaluating their readiness for VMI implementation. A survey of ten companies using this framework indicated that some of these firms are clearly not ready for VMI implementation and that it is best for them to continue to operate in a decentralized manner.

Another major recent change in supply chain management has been the realization that supply chains compete with each other while companies within each supply chain collaborate with each other. It was recently observed that “This increased collaboration – with technology as the catalyst – means that instead of companies competing against companies, by 2010 supply chains will be competing against supply chains.” Similarly noted is that “The classic model of company vs. company is starting to give way to a new model: supply chain vs. supply chain.” When a supply chain is able to offer a product at a lower cost than its competitors, it will be able to increase both market share and profits. However, there is reduced emphasis on how these profits are shared amongst the supply chain members. Our proposed strategies align with this management philosophy and are focused on reducing total supply chain cost and maximizing total supply chain efficiency.

Recognizing the importance and relevance of such policies, Zhu, Gavirneni, and Kapuscinski (2009) formally defined, modeled and analyzed PF policies with constant shipments in a serial supply chain. They showed numerically that the PF policies are able to improve the performance of decentralized supply chains on average by 11%, and alleviate about 43% of the gap between centralized and decentralized supply chains. Gavirneni (2006) showed that Hi-Lo pricing can be used to reduce the retailer's ordering flexibility and improve supply chain performance. Numerical results indicate on average a 5% improvement in supply chain performance. They also studied restricted order quantities that are in the form of piecewise linear functions of past demand. Although these papers rigorously and comprehensively establish the effectiveness of these policies and describe how to effectively choose the shipment quantity, they do not discuss whether the form of the fixed shipment quantity or simple functions of past demand itself minimizes the supply chain cost. In this paper, we consider a set of more general periodic flexible policies that allow the restricted order quantities to take any continuous functional forms of demand history and apply calculus of variations to identify the optimal periodic flexible policy that minimizes the supply chain cost among this general set of PF policies.

Before we go into detailed mathematical analysis, we clarify two issues regarding the generalizability and impact of history dependent PF policies. First, in an n -period cycle, a PF policy, denoted by $PF(n)$, lets the retailer receive a sequence of restricted shipments in the first $n - 1$ periods followed by an unrestricted order in the n th period. At one extreme, the $PF(1)$ policy, which does not restrict the retailer's ordering quantity in any period, is equivalent to decentralized control, and the retailer follows a base stock policy and orders the demand of his previous period. At the other extreme, the optimal $PF(\infty)$ policy, which restricts every retailer's order, achieves the performance of the centralized system. The supply chain may thus, shift from fully decentralized to centralized if we vary the cycle length of the PF policies from 1 to ∞ . Second, previous works show that PF policies in general reduce and increase the supplier's and retailer's cost, respectively. The supply chain's total cost is reduced by a PF policy only if the cost reduction at the supplier exceeds the cost increment at the retailer. For

any cycle length n , (because the decentralized control, the $PF(1)$ policy, is a feasible $PF(n)$ policy) the optimal $PF(n)$ policy always results in a lower supply chain cost, implying that the supplier's cost reduction is enough to cover the cost increment at the retailer. As a result, the optimal $PF(n)$ policy can be implemented under various forms of cost sharing contracts as long as the cost reduction is shared appropriately between supply chain members.

In Section 2, we first summarize the past results and demonstrate the effectiveness of the periodic flexible policy using an example. In Section 3, we formulate the $PF(2)$ problem. In Section 4, we solve the $PF(2)$ problem and show that only two possible optimal functional forms exist, and that the optimal restricted-ordering function can be determined by computing the supplier's and retailer's optimal ordering policies, which reduces to a conventional optimization problem. Then, we numerically compute the optimal $PF(2)$ policy and demonstrate its efficiency, and followed by summarizing the results on the functional forms of the optimal $PF(n)$ policies and providing numerical studies on the optimal $PF(3)$ policy. We demonstrate numerically that the optimal $PF(3)$ policy can achieve near optimal performance and its linear approximation result is within 1% of optimality gap. Finally, we allow the retailer's restricted ordering functions to take any value and characterize the structure of the optimal $PF(2)$ policy when the retailer is allowed to make free returns. In Section 5, we conclude the paper. Unless stated otherwise, all proofs are in the Appendix.

2. Model Setup

Consider a two-stage serial supply chain in which a single supplier (she) provides a single product to a retailer (he) who, in turn, faces independent and identically distributed stochastic end-customer demands. The sequence of events in each period is as follows: (a) The supplier completes production decided in step (e) of the previous period. (b) The retailer places an order with the supplier. (c) The supplier, if needed, outsources (at an additional cost as in Lee, So, & Tang, 2000) to satisfy shortages if needed. The outsourced units become available to the supplier immediately (see e.g., Gavirneni, Kapuscinski, and Tayur 1999, Lee et al. (2000) for similar assumptions on the immediate availability of the expedited items). (d) The retailer's complete order is filled using the supplier's inventory after outsourcing. (e) The supplier decides how much to produce. (f) End-customer demand is realized, and the retailer uses his on-hand inventory to satisfy the end-customer demand, while unsatisfied demand at the retailer is backlogged. All costs are calculated at the end of the period.

We summarize the commonly used notation in Table 1, omitting some of the less commonly used notations. We formulate the problem as follows: The end-customer demand in period t , denoted by $\tilde{d}_t$, is a continuous random variable with positive support on $[0, \infty)$ (or finite interval $[0, \bar{d}]$, $\bar{d} < \infty$) and has cdf $F(\cdot)$ and pdf $f(\cdot)$. Let $F^0(\cdot) = 1 - F(\cdot)$, and $(F^0)^{-1}(\cdot)$ be the inverse of $F^0(\cdot)$. The retailer's holding and penalty cost rates are h_r and p_r per unit, respectively. The supplier's holding and expediting cost rates are h_s and p_s per unit, respectively. Assume, without loss of generality, that $h_r < p_r$ and $h_s < p_s$, and there are neither fixed costs nor capacity limits at either supplier or the retailer. Furthermore, in every period we assume that the retailer and supplier make ordering decisions before the end-customer demand is realized. Other than that, there are no lead times, therefore, both the supplier and the retailer must contend with one period of uncertainty when they make their decisions. The zero-lead-time assumption is also made in Zhu, Gavirneni, Kapuscinski (2010) and Chen and Gavirneni (2010) mainly for the sake of analytical tractability. However numerical studies have shown, see section 6.2 of Chen and Gavirneni (2010) and section 5.3.4 of Zhu, Gavirneni, Kapuscinski (2010), that these policies are effective under non-zero leadtimes as well.

Table 1
Commonly used notation.

Notation	Definition
$\otimes$	convolution operator
$a \vee b, a \wedge b, (a - b)^+, (a - b)^-$	$\max\{a, b\}, \min\{a, b\}, \max\{a - b, 0\}, \max\{b - a, 0\}$
$\tilde{d}$	a random variable with the same distribution as the end-customer demand in period t
d	a realization of $\tilde{d}$
$\tilde{d}_t$	end-customer demand in period t [period $t - 1$]
d_t	a realization of $\tilde{d}_t$ [$\tilde{d}_{t-1}$]
$f(\cdot), F(\cdot)$	pdf of end-customer demand, cdf of end-customer demand
$F^0(\cdot), (F^0)^{-1}(\cdot)$	$1 - F(\cdot)$, inverse of $F^0(\cdot)$
$Q(d)$	the [optimal] retailer restricted-ordering function
$F_{d-Q}(\cdot), F_Q(\cdot)$	cdf of $d_{t-1} - Q(d_{t-1})$, cdf of $Q(d_{t-1})$
h_r	holding/penalty cost rate per unit at the retailer [supplier]
Q	$Q = \{Q(d) 0 \leq Q(d) \leq d, Q(d) \in C[0, \infty)\}$
S_s^c	the centralized optimal echelon produce up to level at the supplier
S_r^c	the centralized optimal order up to level at the retailer if the supplier has enough [does not have enough] on-hand inventory
S_s^d	optimal produce [order] up to level under fully decentralized system
S_r^f	[optimal] order up to level at the retailer in the free-ordering period
$S_s^f(d_{t-1})$	produce up to level at the supplier in the retailer free-ordering period given that the demand in the previous period is d_{t-1}
S_s^{R*}	[optimal] produce up to level at the supplier in the restricted-ordering period
S_t	the supplier's on-hand inventory after he receives the orders in period t
z_s^f	a constant such that $z_s^f = S_s^f(d_{t-1}) - (d_{t-1} - Q(d_{t-1}))$
z_s^{f*}	the optimal z_s^f at the supplier
$\tilde{z}_s^f$	a constant such that $\tilde{z}_s^f = \tilde{S}_s^f(d_{t-1}) - (d_{t-1} - Q(d_{t-1}))$
π	a feasible policy $\pi = (S_s^R, z_s^f, S_r^f)$
$\mathcal{A}$	$\mathcal{A} = \{\pi \pi = (S_s^R, z_s^f, S_r^f), 0 \leq S_s^R \leq z_s^f, S_r^f \leq S_s^f\}$
$\hat{Q}(d_{t-1}, S_s^R, z_s^f, S_r^f)$	the optimal ordering function for the unrestricted problem
$Q_s^*(d_{t-1}, S_s^R, z_s^f, S_r^f)$	the optimal retailer restricted-ordering function evaluated at $\pi \in \mathcal{A}$
$C(S_s^R, z_s^f, S_r^f Q(\cdot), d_{t-1})$	the supply chain cost in a cycle given d_{t-1} when the cycle starts from the restricted ordering period t
$C_e(S_s^R, z_s^f, S_r^f Q(\cdot))$	the expected supply chain cost in a cycle
$C_r(S_r^f Q(\cdot), d_{t-1})$	the retailer's cost in a cycle given d_{t-1} when the cycle starts from the restricted ordering period t
$C_s(S_s^R, z_s^f Q(\cdot), d_{t-1})$	the supplier's cost in a cycle given d_{t-1} and $Q(d_{t-1})$ when the cycle starts from the restricted ordering period t

183 We consider three types of policies: fully decentralized policy, cen-
184 tralized policy, and periodic flexible policy.

185 2.1. Fully decentralized policy

186 Under the fully decentralized policy, the retailer and supplier
187 minimize their total inventory cost independently, and neither has
188 restrictions on their order quantities. Because we assume no fixed
189 costs or capacity limits, both parties operate under a stationary,
190 newsvendor-type order up to policy. Let S_s^d and S_r^d be the opti-
191 mal order up to levels at the supplier and the retailer respectively.
192 These order up to levels can be computed using simple newsven-
193 dor formulae, being, $S_r^d = (F)^{-1}(\frac{p_r}{h_r + p_r})$ and $S_s^d = (F)^{-1}(\frac{p_s}{h_s + p_s})$. As a
194 result of a lack of coordination between supplier and retailer, the
195 supply chain cost under this policy is high.

196 2.2. Centralized policy

197 Under the centralized policy, a hypothetical single decision
198 maker controls the inventory replenishment processes at both
199 parties and minimizes the total supply chain cost. [Zhu et al. \(2009\)](#)
200 prove that the optimal policy can be characterized by three
201 parameters, S_s^c , S_r^c , and S_r^0 . Under the optimal policy, the supplier
202 produces up to the echelon produce up to level S_s^c every period.
203 If the supplier has enough on-hand inventory to bring the retailer
204 inventory level up to S_r^c , the retailer orders up to S_r^c . Otherwise, if
205 the echelon inventory level is higher than S_r^0 , the retailer orders
206 up to the echelon inventory level and the supplier fulfills the
207 retailer's order using all the on-hand inventory. If the echelon
208 inventory level is less than S_r^0 , the retailer orders up to S_r^0 and the

supplier fulfills the order first using his on-hand inventory, then
expediting until the retailer's inventory level reaches S_r^0 . Although
not available in closed form, these parameters satisfy $S_r^0 \leq S_r^c \leq S_s^c$
and can be efficiently computed using Infinitesimal Perturbation
Analysis (IPA).

214 2.3. Periodic flexible policy

215 Although periodic flexible policies can be of any arbitrary
216 length, we discuss in details a policy of length of two. Under the
217 two-period setting, the retailer follows a restricted order policy in
218 one period, followed by free-ordering in the next period, and this
219 pattern repeats itself throughout the time horizon. In the restricted
220 period, if the retailer has observed an end-customer demand of d ,
221 he orders the restricted quantity $Q(d)$. This period is referred to
222 as the *retailer restricted-ordering period* and $Q(d)$ as the *retailer re-*
223 *stricted ordering function*. $Q(d)$ is determined by the supplier with
224 the objective of minimizing the total supply chain cost, and is
225 known to both the supplier and the retailer. In the free-ordering
226 period, the retailer is free to order any non-negative quantity. We
227 refer to this period as the *retailer free-ordering period*. Clearly, when
228 $Q(d) = d$, the PF policy is equivalent to the fully decentralized pol-
229 icy. [Zhu et al. \(2009\)](#) use a constant function $Q(d) = K$. [Gavirneni \(2006\)](#)
230 uses the functional form $Q(d) = (d - \delta)^+$, with a δ as a
231 known constant. These functional forms are all special cases, and
232 the optimal functional form of $Q(d)$ is not known. In each of the
233 above examples, retailer's demand information is shared with the
234 supplier once the demand is realized so that the retailer cannot
235 game the system by misreporting his demand. Here, we also make

Fig. 1. Percentage improvement under $Q(d) = d \wedge a$ compared to the fully decentralized system.

the same assumption. We demonstrate the effectiveness of PF(2) policies in the following example.

Example 1. Let $d_t \sim \exp(1/100)$. Its density and cumulative distribution functions are

$$f(x) = \frac{1}{100} e^{-\frac{x}{100}}, \quad F(x) = 1 - e^{-\frac{x}{100}}.$$

We set the cost parameters as $h_r = 1$, $p_r = 9$, $p_s = 19$, and consider two cases, $h_s = 0.5$ and 1.5 . Fig. 1 shows the percentage change in cost at the retailer, supplier, and supply chain, respectively, under PF(2) with the restricted-ordering function taking the form $Q(d) = d \wedge a$. In both sub-figures, we let “ a ” change from 1 to 300 in increments of 1. The percentage change in cost for the supply chain and its members, respectively, is computed by comparing their expected costs under $Q(d) = d \wedge a$ to $Q(d) = d$.

In both sub-figures, the percentage cost reduction for the supplier is always negative, i.e., the supplier's expected cost is always higher under the decentralized control, whereas the percentage cost reduction for the retailer is always positive, i.e., the retailer's expected cost is always lesser under the decentralized control. The percentage cost reduction for the supply chain, however, can be either negative or positive. For example, in the first figure ($h_s = 0.5$), when $a < 70$, the percentage cost reduction for the supply chain is negative, implying that the supply chain's expected cost is higher under the periodic flexible policy; but when $a > 70$, the percentage cost reduction for the supply chain is positive, so the supply chain's expected cost is lower under the periodic flexible policy. In the second figure, when $h_s = 1.5$, the periodic flexible policies always result in lower supply chain costs. The increasing in supply chain cost may occur only under ad hoc PF policies. Thus, to ensure an improvement in the supply chain performance, optimal PF policy is necessary.

3. Analysis of the PF(2) policy

In this section, we formulate the PF(2) problem and solve for the optimal PF(2) policy.

3.1. Optimal ordering and production policies

Given the restricted ordering function $Q(d)$, the retailer chooses an ordering policy to minimize his expected long run average cost, and in response, the supplier chooses her production policy to minimize her long run average cost. This setting is applicable when neither the supplier nor the retailer has the power to interfere with the other party's local operational decisions, but they may jointly select the restricted order quantity function $Q(d)$ and divide

the cost savings through negotiation or contracting. The coordination of the supply chain is feasible so long as the total supply chain cost is reduced. Indeed, the coordination can be achieved even under very general information structures. One example is the relational contract studied in Plambeck and Taylor (2006) and Taylor and Plambeck (2007a, 2007b), and another example is the dynamic efficient mechanisms studied in Lewis et al. (2017), Athey and Segal (2013) and Bergemann and Välimäki (2010).

Consider the case that the supplier and the retailer operate under the PF(2) policy over an infinite horizon. To facilitate the characterization of the long-run average cost, we make the following assumption.

Assumption 1. $Q(d) \leq d$. The restricted order quantity is less than the last period's demand.

Note that Gavirneni (2006) model relies on this assumption. Zhu et al. (2009) model and the Chen and Gavirneni (2011) model are consistent with this assumption when the fixed shipment quantity is zero. When this assumption is violated (as is the case when the fixed shipment quantity is equal to the mean demand), regeneration of the supply chain is not guaranteed, which makes the formulation of the optimization problem challenging. When the model satisfies Assumption 1, we will show in Sections 3.1.1 and 3.1.2 that the system regenerates itself at the end of every free-ordering period, and therefore, the long-run average cost is equal to the expected cost within a cycle. In Section 4.4, we will discuss the implications of violating this assumption and prove that, under some conditions, the policies where fixed shipment quantity is equal to mean demand used by Zhu et al. (2009) and Chen and Gavirneni (2011) are indeed optimal.

We begin our analysis by assuming that period t is the retailer restricted-ordering period in which the retailer is required to order $Q(d_{t-1})$. Fig. 2 illustrates the sequence of events under a PF(2) policy in a two period cycle starting from the restricted period (hereafter referred to as a cycle). The supplier's actions are listed above the time line whereas the retailer's actions are listed below. For analytical convenience, we derive the retailer's long run average cost in periods of two from period t to $t + 1$, while for the supplier from period $t - 1$ to t . The different grouping of the periods will not affect the value of the long run average costs.

3.1.1. Retailer's optimal ordering policy and long run average cost

We first analyze the retailer's optimal ordering policy. Given the restricted ordering function $Q(d)$, the retailer chooses how much to order in the free-ordering period $t + 1$ to minimize his long run average cost. The retailer's inventory in the restricted period t before receiving order and free-ordering period $t + 1$ after receiving order

Cost is evaluated at the end of every period

Fig. 2. Sequence of events for a PF(2) policy.

are denoted as x_r^t and y_r^t , respectively. The retailer finds the ordering policy that minimizes his long run average cost by solving the following problem

$$\Upsilon_r = \min_{\pi_r \in \Pi_r} \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T C_r^t(x_r^t, y_r^t | Q, d_{t-1})}{T},$$

where

$$\begin{aligned} C_r^t(x_r^t, y_r^t | Q, d_{t-1}) = & h_r E_{\tilde{d}_t} [x_r^t + Q(d_{t-1}) - \tilde{d}_t]^+ \\ & + p_r E_{\tilde{d}_t} [x_r^t + Q(d_{t-1}) - \tilde{d}_t]^- \\ & + h_r E_{\tilde{d}_{t+1}} [y_r^t - \tilde{d}_{t+1}]^+ + p_r E_{\tilde{d}_{t+1}} [y_r^t - \tilde{d}_{t+1}]^- \end{aligned}$$

is the retailer's cost in a cycle given demand d_{t-1} , and Π_r is the set of retailer's ordering policies. The following lemma shows that the retailer's optimal policy is stationary under the long run average criteria.

Lemma 1. Under Assumption 1, in every free-ordering period, the retailer's optimal policy is a base stock policy, that is, $y_r^{t*} = s^0$, where s^0 is a constant defined by

$$\begin{aligned} s^0 = \arg \min_s \{ & h_r E_{\tilde{d}} [s - \tilde{d}]^+ + p_r E_{\tilde{d}} [s - \tilde{d}]^- \\ & + h_r E_{\tilde{d}, \tilde{d}} [s - \tilde{d} + Q(\tilde{d}) - \hat{d}]^+ + p_r E_{\tilde{d}, \tilde{d}} [s - \tilde{d} + Q(\tilde{d}) - \hat{d}]^- \}, \end{aligned}$$

and $\hat{d}$ is a random variable with a c.d.f. identical to the retailer's demand.

Lemma 1 implies that minimizing the retailer's long run average cost is equivalent to minimizing his expected cost within a cycle. Given $Q(\cdot)$, the retailer has only one decision to make in every cycle: the order quantity in the free-ordering period. The retailer's order up to level in the free-ordering period is denoted as S_r^F , and the retailer's resulting cost in a cycle is

$$\begin{aligned} C_r(S_r^F | Q(\cdot), d_{t-1}) = & h_r E_{\tilde{d}_t} [S_r^F - d_{t-1} + Q(d_{t-1}) - \tilde{d}_t]^+ \\ & + p_r E_{\tilde{d}_t} [S_r^F - d_{t-1} + Q(d_{t-1}) - \tilde{d}_t]^- \\ & + h_r E_{\tilde{d}_{t+1}} [S_r^F - \tilde{d}_{t+1}]^+ + p_r E_{\tilde{d}_{t+1}} [S_r^F - \tilde{d}_{t+1}]^-, \quad (1) \end{aligned}$$

and the retailer minimizes his expected cost in a cycle

$$(\mathcal{R}) \min_{S_r^F} E_{\tilde{d}_{t-1}} [C_r(S_r^F | Q(\cdot), \tilde{d}_{t-1})]$$

by selecting the optimal order up to level in the free-ordering period. For a certain $Q(\cdot)$, the following lemma provides a necessary condition for S_r^F to be optimal.

Lemma 2. The retailer's optimal base-stock level, denoted by S_r^{F*} , satisfies $S_r^{F*} \geq S_r^d$ and

$$S_r^{F*} = \inf \left\{ S \mid F^0(S) + (F \otimes F_{d-Q})^0(S) < \frac{2h_r}{h_r + p_r} \right\}. \quad (2)$$

Lemmas 1 and 2 also imply that it is sufficient to consider only stationary ordering policies for the retailer when finding the supplier's optimal production policy as well as the supply chain optimal restricted-ordering function. Under the optimal order up to policy, the retailer orders $d_{t-1} - Q(d_{t-1}) + d_t$ from the supplier in the free-ordering period $t + 1$. The supplier's demand in the free-ordering period $t + 1$ is, thus, $d_{t-1} - Q(d_{t-1}) + d_t$. This demand pattern implies that in the free-ordering period $t + 1$, the supplier is guaranteed to observe a demand of $d_{t-1} - Q(d_{t-1})$ on top of the random demand $\tilde{d}_t$. Thus, the supplier can pre-plan for this guaranteed demand in the free-ordering period $t + 1$ and hence reduce her cost. In the following section, we characterize the supplier's optimal production policy.

3.1.2. Supplier's optimal production policy and long run average cost

Given that the retailer follows a stationary order up to policy, the supplier chooses how much to produce in both the restricted-ordering period t and the free-ordering period $t + 1$ to minimize her long run average cost. Denote x_s^t as the supplier's inventory in the restricted period t and y_s^t as the supplier's inventory in the free-ordering period $t + 1$ after the production in the last period is completed. She solves the following problem to minimize her long run average cost.

$$\Upsilon_s = \min_{\pi_s \in \Pi_s} \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T C_s^t(x_s^t, y_s^t | Q, d_{t-1})}{T},$$

where

$$\begin{aligned} C_s^t(x_s^t, y_s^t | Q, d_{t-1}) = & h_s [x_s^t - Q(d_{t-1})]^+ + p_s [x_s^t - Q(d_{t-1})]^- \\ & + h_s E_{\tilde{d}_t} [y_s^t - d_{t-1} + Q(d_{t-1}) - \tilde{d}_t]^+ \\ & + p_s E_{\tilde{d}_t} [y_s^t - d_{t-1} + Q(d_{t-1}) - \tilde{d}_t]^- \end{aligned}$$

is the supplier's cost in a cycle given d_{t-1} and Π_s is the set of supplier's production policies. Let $x^0 = (F^0)_{Q^{-1}}^{-1}(\frac{h_s}{h_s + p_s})$. The following

lemma characterizes the supplier's optimal production policy under the long run average criterion.

Lemma 3. The supplier's optimal production policy is a modified stationary policy. In the retailer's free-ordering period $t-1$, the supplier's optimal policy is a modified myopic policy, i.e., $x_s^{t*} = \max\{x^0, z^0 - d_{t-2}\}$. In the retailer's restricted-ordering period t , the supplier's optimal policy is to produce up to a constant z^0 after adjusting for the guaranteed demand, i.e., $y_s^{t*} + d_{t-1} - Q(d_{t-1}) = z^0$, and

$$z^0 = \arg \min_z \left\{ h_s E_{\tilde{d}_t} [z - \tilde{d}_t]^+ + p_s E_{\tilde{d}_t} [z - \tilde{d}_t]^- + h_s E_{\tilde{d}_{t-2}, \tilde{d}_{t-1}} [\max\{x^0, z - \tilde{d}_{t-2}\} - Q(\tilde{d}_{t-1})]^+ + p_s E_{\tilde{d}_{t-2}, \tilde{d}_{t-1}} [\max\{x^0, z - \tilde{d}_{t-2}\} - Q(\tilde{d}_{t-1})]^- \right\}.$$

Lemma 2 suggests that the dynamics of the supplier's on-hand inventory regenerates itself in every cycle. As a result, we only need to analyze the supplier's production cost within a cycle. Let S_s^R be the supplier's inventory level after production in the retailer restricted-ordering period, and let $S_s^F(d_{t-1})$ be the supplier's inventory level after production in the retailer free-ordering period t given that in the previous period $t-1$ the realized demand is d_{t-1} . Because the supplier is providing 100% service level to the retailer, no backlog or lost sales occur at the supplier. The inventory levels after production (i.e., S_s^R and $S_s^F(d_{t-1})$) must then be non-negative. Note that in the free-ordering period t , the supplier is guaranteed to see a demand of at least $d_{t-1} - Q(d_{t-1})$. Thus, at the end of the free-ordering period t , this amount will be for sure deducted from the supplier's on-hand inventory. It is, as such, optimal for the supplier to produce $d_{t-1} - Q(d_{t-1})$ plus a constant amount in the same period. Let z_s^F be the remainder of the supplier's on-hand inventory at the end of the free-ordering period t satisfying $z_s^F = S_s^F(d_{t-1}) - (d_{t-1} - Q(d_{t-1}))$. From **Lemma 3**, we know that the optimal z_s^F is independent of $\tilde{d}_{t-1}$. Later, we will use z_s^F instead of $S_s^F(d_{t-1})$ to characterize the supplier's optimal policy and the optimal PF(2) policy. Given d_{t-2} and d_{t-1} , the supplier's cost in a cycle can be written as a function of S_s^R and z_s^F as follows:

$$C_s(S_s^R, z_s^F | Q(\cdot), d_{t-2}, d_{t-1}) = h_s [S_s^R \vee (z_s^F - d_{t-2}) - Q(d_{t-1})]^+ + p_s [S_s^R \vee (z_s^F - d_{t-2}) - Q(d_{t-1})]^- + h_s E_{\tilde{d}_t} [z_s^F - \tilde{d}_t]^+ + p_s E_{\tilde{d}_t} [z_s^F - \tilde{d}_t]^- \quad (3)$$

For a certain $Q(\cdot)$, the supplier solves the following problem to minimize her expected cost within a cycle:

$$(S) \min_{S_s^R \geq 0, z_s^F} E_{\tilde{d}_{t-2}, \tilde{d}_{t-1}} [C_s(S_s^R, z_s^F | Q(\cdot), \tilde{d}_{t-2}, \tilde{d}_{t-1})].$$

Let $(F^0)^{-1}(x) = \infty$ when $x \leq 0$, and $(F^0)^{-1}(x) = 0$ when $x \geq 1$. The following lemma characterizes the supplier's optimal production policy.

Lemma 4. The supplier's optimal base-stock level in the retailer-restricted ordering periods satisfies $0 \leq S_s^{R*} \leq S_s^d$ and

$$S_s^{R*} = \inf \left\{ S | F_Q^0(S) < \frac{h_s}{h_s + p_s} \right\}. \quad (4)$$

Under optimality, we must have $z_s^{F*} \geq S_s^{R*}$.

3.1.3. The supply chain cost and PF(2) problem

Sections 3.1.1 and 3.1.2 imply that minimizing the long run average supply chain cost is equivalent to minimizing the cost in the two-period cycle starting from the retailer restricted-ordering period. The expected supply chain cost in a cycle is given in the following equation:

lowing equation:

$$C_e(S_s^R, z_s^F, S_r^F | Q(\cdot)) = E_{\tilde{d}_{t-2}, \tilde{d}_{t-1}} [C(S_s^R, z_s^F, S_r^F | Q(\cdot), \tilde{d}_{t-2}, \tilde{d}_{t-1})] \\ = E_{\tilde{d}_{t-1}} [C_r(S_r^F | Q(\cdot), \tilde{d}_{t-1})] \\ + E_{\tilde{d}_{t-2}, \tilde{d}_{t-1}} [C_s(S_s^R, z_s^F | Q(\cdot), \tilde{d}_{t-2}, \tilde{d}_{t-1})], \quad (5)$$

where $C(S_s^R, z_s^F, S_r^F | Q(\cdot), d_{t-2}, d_{t-1}) = C_r(S_r^F | Q(\cdot), d_{t-1}) + C_s(S_s^R, z_s^F | Q(\cdot), d_{t-2}, d_{t-1})$ is the expected supply chain cost in the cycle given d_{t-2} and d_{t-1} .

The PF(2) problem is to find a retailer restricted-ordering function that minimizes the supply chain cost given that the supplier and retailer select ordering policies that minimize their own costs. Because the supply chain cost is separable in both parties' decisions due to expediting, the PF(2) problem can be formulated as follows:

$$(\mathcal{O}) \min_{Q(\cdot) \in \mathcal{Q}} \left\{ \min_{S_r^F} E_{\tilde{d}_{t-1}} [C_r(S_r^F | Q(\cdot), \tilde{d}_{t-1})] + \min_{S_s^R \geq 0, z_s^F} E_{\tilde{d}_{t-2}, \tilde{d}_{t-1}} [C_s(S_s^R, z_s^F | Q(\cdot), \tilde{d}_{t-2}, \tilde{d}_{t-1})] \right\}.$$

Unlike the conventional calculus of variations problem, the PF(2) problem also minimizes the supplier's and retailer's decisions, which prevents us from solving the PF(2) problem directly from the above formulation because we cannot completely characterize the minimizers $(S_s^{R*}, z_s^{F*}, S_r^{F*})$ in terms of $Q(d)$. To overcome this challenge, note that given any $Q(d)$, the retailer and supplier's optimal order and produce-up-to levels also minimize the expected supply chain cost, particularly for the optimal $Q(d)$. As a result, **Lemmas 2** and **4** also hold for optimal $Q(d)$. We, therefore, first reformulate the PF(2) problem and apply calculus of variations to find the optimal structure of the restricted-ordering function. Then, we further characterize the optimal structure by imposing the necessary conditions in **Lemmas 2** and **4**. We are, to our knowledge, the first to combine calculus of variations with conventional optimization in order to solve inventory problems.

3.2. Equivalent formulation of PF(2) problem

Because problem $(\mathcal{O})$ consists of two minimizations, we can reformulate it by switching their ordering:

$$(\mathcal{P}) \min_{(S_s^R, z_s^F, S_r^F) \in \mathcal{A}} \left\{ \min_{Q(\cdot) \in \mathcal{Q}} C_e(S_s^R, z_s^F, S_r^F | Q(\cdot)) \right\}.$$

$(\mathcal{P})$ gives the same optimal base-stock levels and expected cost as problem $(\mathcal{O})$.

Remark. Problem $(\mathcal{P})$ does not have the first derivative term of the variation functional. As a result, the functional derivative becomes analogous to the first order condition under conventional optimizations, which significantly reduces the complexity of the problem. In general, policies can be viewed as a function of the states in most dynamic programming problems do not account for the cost of the speed of change of the state, and, thus, can be formulated as simple calculus of variations problems. Moreover, calculus of variations can provide a better characterization of the optimal policy structure as a function of the system states, which is especially beneficial under non-stationary settings where, in general, dynamic programming can only provide a partial characterization of the optimal policy structure.

4. Finding the optimal PF(2) policy

4.1. Structure of the optimal $Q(d)$

For any given $\pi = (S_s^R, z_s^F, S_r^F) \in \mathcal{A}$, we minimize the supply chain's expected cost within a cycle from period t to period $t+1$

over the retailer's restricted-ordering function $Q(d_{t-1})$, i.e.,

$$\min_{Q \in \mathcal{Q}} \int_0^\infty E_{\tilde{d}_{t-2}} [C(S_s^R, z_s^F, S_r^F | Q(\cdot), \tilde{d}_{t-2}, d_{t-1})] f(d_{t-1}) dd_{t-1}. \quad (6)$$

Consider any perturbation $Q(d) + \epsilon \eta(d)$ of the retailer's restricted-ordering function, where ϵ is a small number and $\eta(d)$ is a function on $[0, \infty)$ that has continuous first derivatives and vanishes at the end points, i.e., $\eta(0) = \eta(\infty) = 0$. If $Q(\cdot)$ is the minimizer of the supply chain's cost, then for any η , $\int_0^\infty E_{\tilde{d}_{t-2}} [C(S_s^R, z_s^F, S_r^F | Q(\cdot) + \epsilon \eta(\cdot), \tilde{d}_{t-2}, d_{t-1})] f(d_{t-1}) dd_{t-1}$ is minimized at $\epsilon = 0$, and the first order condition holds. Because $E_{\tilde{d}_{t-2}} [C(S_s^R, z_s^F, S_r^F | Q(\cdot), \tilde{d}_{t-2}, d_{t-1})]$ has a continuous second derivative with respect to $Q(\cdot)$, the Euler–Lagrange equation remains ground for the retailer's optimal restricted-ordering function (see Gelfand & Fomin, 2000, Section 5). We can solve for the structure of the optimal $Q(d)$ by first characterizing it without Assumption 1.

Compared to the traditional optimization procedure where the first order properties of $Q(d)$ such as monotonicity are characterized by showing convexity and modularity properties of the objective function, calculation of variations allows us to characterize higher order properties of $Q(d)$ in d . In particular, in the following analysis, we demonstrate that the optimal $Q(d)$ is a piecewise non-linear function of d with each piece possesses a different functional form, which can be characterized in explicit forms.

A relaxed problem. For any $\pi \in \mathcal{A}$ and d_{t-1} , let $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ be the solution to the calculus of variations in (6) when $Q(\cdot)$ is allowed to violate Assumption 1. Then, $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ must satisfy the unconstrained Euler–Lagrange equation. Note that because C does not contain the term $\partial Q / \partial d$, the Euler–Lagrange equation for the unconstrained problem is equivalent to

$$\frac{\partial}{\partial Q} E_{\tilde{d}_{t-2}} [C(S_s^R, z_s^F, S_r^F | Q(\cdot), \tilde{d}_{t-2}, d_{t-1})] |_{Q(d_{t-1}, S_s^R, z_s^F, S_r^F)} = \hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$$

= 0. Let

$$\Gamma_1(Q | d_{t-1}) = -h_s \text{Prob}(z_s^F - \tilde{d}_t \geq Q) + p_s \text{Prob}(z_s^F - \tilde{d}_t < Q)$$

$$+ h_r \text{Prob}(S_r^F - d_{t-1} + Q \geq \tilde{d}_t)$$

$$- p_r \text{Prob}(S_r^F - d_{t-1} + Q < \tilde{d}_t),$$

$$\Gamma_2(Q | d_{t-1}) = -h_s + h_r \text{Prob}(S_r^F - d_{t-1} + Q \geq \tilde{d}_t)$$

$$- p_r \text{Prob}(S_r^F - d_{t-1} + Q < \tilde{d}_t).$$

$\partial C / \partial Q$ is written in terms of Γ_1 and Γ_2 as follows:

$$\begin{aligned} \frac{\partial}{\partial Q} E_{\tilde{d}_{t-2}} [C(S_s^R, z_s^F, S_r^F | Q(\cdot), \tilde{d}_{t-2}, d_{t-1})] |_{Q(d_{t-1}, S_s^R, z_s^F, S_r^F)} &= \hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) \\ &= \begin{cases} \Gamma_1(\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) | d_{t-1}) & \hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) > S_s^R \\ \Gamma_2(\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) | d_{t-1}) & \hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) \leq S_s^R \end{cases} \end{aligned}$$

Observe that Γ_1 and Γ_2 are increasing functions of Q , and $\Gamma_1(Q) \geq \Gamma_2(Q)$, for all Q and d_{t-1} . Reorganizing the Euler–Lagrange equation in terms of the Γ functions, we get:

$$\Gamma_1(\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) | d_{t-1}) = 0 \text{ iff } \Gamma_1(S_s^R | d_{t-1}) < 0$$

$$\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) = S_s^R \text{ iff } \Gamma_1(S_s^R | d_{t-1}) \geq 0 \text{ and}$$

$$\Gamma_2(S_s^R | d_{t-1}) \leq 0$$

$$\Gamma_2(\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) | d_{t-1}) = 0 \text{ iff } \Gamma_2(S_s^R | d_{t-1}) > 0, \quad (7)$$

and further find the explicit form for $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$:

$$\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) > S_s^R \text{ iff } \Gamma_1(S_s^R | d_{t-1}) < 0$$

$$\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) = S_s^R \text{ iff } \Gamma_1(S_s^R | d_{t-1}) \geq 0 \text{ and } \Gamma_2(S_s^R | d_{t-1}) \leq 0$$

$$\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) < S_s^R \text{ iff } \Gamma_2(S_s^R | d_{t-1}) > 0.$$

Denote $Q_\pi^*(d_{t-1}, S_s^R, z_s^F, S_r^F)$ as the optimal retailer restricted-ordering function evaluated at $\pi \in \mathcal{A}$ that satisfies Assumption 1.

The results that characterize $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ and $Q_\pi^*(d_{t-1}, S_s^R, z_s^F, S_r^F)$ are as follows:

Proposition 1. (Optimal Restricted-Ordering Function for the Relaxed and Original Problems) The optimal order quantity for the relaxed problem satisfies:

$$\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) = \begin{cases} d_{t-1} + (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) - S_r^F & d_{t-1} \leq -(F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) + S_r^F + S_s^R \\ S_s^R & S_r^F + S_s^R - (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) < d_{t-1} \\ & \leq S_r^F + S_s^R - \Delta(z_s^F - S_s^R) \\ (\Gamma_1)^{-1}(0 | d_{t-1}) & S_r^F + S_s^R - \Delta(z_s^F - S_s^R) < d_{t-1} \\ & \leq z_s^F + S_r^F - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) \\ d_{t-1} + (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) - S_r^F & z_s^F + S_r^F - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) < d_{t-1} \end{cases} \quad (8)$$

where

$$\Delta(z_s^F - S_s^R) = (F^0)^{-1} \left(\frac{h_r - h_s + (h_s + p_s)F^0(z_s^F - S_s^R)}{h_r + p_r} \right).$$

When $h_s \geq h_r$, we have $\frac{h_r - h_s}{h_r + p_r} \leq 0$, so the first term in $\hat{Q}$ disappears, and $\hat{Q}$ becomes

$$\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) = \begin{cases} S_s^R & S_r^F + S_s^R - (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) < d_{t-1} \\ & \leq S_r^F + S_s^R - \Delta(z_s^F - S_s^R) \\ (\Gamma_1)^{-1}(0 | d_{t-1}) & S_r^F + S_s^R - \Delta(z_s^F - S_s^R) < d_{t-1} \\ & \leq z_s^F + S_r^F - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) \\ d_{t-1} + (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) - S_r^F & z_s^F + S_r^F - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) < d_{t-1} \end{cases}$$

When $p_s \geq p_r$, we have $\frac{h_r + p_s}{h_r + p_r} \geq 1$, so the last term in $\hat{Q}$ disappears, and $\hat{Q}$ becomes

$$\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) = \begin{cases} d_{t-1} + (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) - S_r^F & d_{t-1} \leq -(F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) + S_r^F + S_s^R \\ S_s^R & S_r^F + S_s^R - (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) < d_{t-1} \\ & \leq S_r^F + S_s^R - \Delta(z_s^F - S_s^R) \\ (\Gamma_1)^{-1}(0 | d_{t-1}) & S_r^F + S_s^R - \Delta(z_s^F - S_s^R) < d_{t-1} \end{cases}$$

The optimal order quantity is $Q_\pi^*(d_{t-1}, S_s^R, z_s^F, S_r^F) = (\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F) \wedge d_{t-1})^+$.

Proposition 1 characterizes the optimal restricted-ordering function for the relaxed problem under the cases (1) $h_s \geq h_r$ and $p_s \geq p_r$, (2) $h_s < h_r$ and $p_s \geq p_r$, (3) $h_s \geq h_r$ and $p_s < p_r$, and (4) $h_s < h_r$ and $p_s < p_r$. Under all four cases, the optimal restricted-ordering function is found by restricting the optimal restricted-ordering function for the relaxed problem to take non-negative values that are no more than d_{t-1} . As a result, the optimal restricted-ordering function may have four possible structures, depending on the relationships between p_s and p_r and between h_s and h_r . Fig. 3 provides a schematic representation of $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ with all four possible $\hat{Q}(\cdot)$ s. If $h_r < h_s$, then $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ is always at or above S_s^R ; if $p_s > p_r$, then $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ is always at or below z_s^F . Taking the same parameter settings as in Example 1 ($h_s = 0.5$), Fig. 4 illustrates the actual $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ given that $S_s^R = 300$, $z_s^F = 400$, and $S_r^F = 250$. As shown in Fig. 4, it is evident that $\hat{Q}(d_{t-1}, S_s^R, z_s^F, S_r^F)$ is almost piecewise linear in d_{t-1} . The following theorem further characterizes the structure of $Q^*(d)$.

Fig. 3. A schematic representation of the optimal retailer restricted-ordering function ($\hat{Q}(\cdot)$) for the relaxed problem given values of S_s^R , z_s^F , S_r^F . $A = -(F^0)^{-1}(\frac{h_r - h_s}{h_r + p_r}) + S_r^F + S_s^R$, $B = S_r^F + S_s^R - \Delta(z_s^F - S_s^R)$, $C = z_s^F + S_r^F - (F^0)^{-1}(\frac{h_r + p_r}{h_r + p_s})$. 1 stands for S_s^R , and 2 stands for $z_s^F - (F^0)^{-1}(\frac{h_r + p_r}{h_r + p_s})$.

Fig. 4. The optimal retailer restricted-ordering function ($\hat{Q}$) for Example 1, when $S_s^R = 300$, $z_s^F = 400$, and $S_r^F = 250$.

3. When $p_s \geq p_r$, the optimal restricted order quantity function satisfies:

$$Q^*(d_{t-1}) = \begin{cases} d_{t-1} \wedge S_s^{R*} & d_{t-1} \in [0, S_r^{F*} + S_s^{R*} - \Delta(z_s^{F*} - S_s^{R*})) \\ \hat{Q}(d_{t-1}, S_s^{R*}, z_s^{F*}, S_r^{F*}) & d_{t-1} \in [S_r^{F*} + S_s^{R*} - \Delta(z_s^{F*} - S_s^{R*}), S_r^{F*} + z_s^{F*} - (F^0)^{-1}(\frac{h_r + p_r}{h_s + p_s})) \\ z_s^{F*} - (F^0)^{-1}(\frac{h_s + p_r}{h_s + p_s}) & d_{t-1} \in [S_r^{F*} + z_s^{F*} - (F^0)^{-1}(\frac{h_s + p_r}{h_s + p_s}), \infty) \end{cases}$$

Theorem 1 implies that the retailer's optimal PF(2) policy in the restricted period is either a state dependent capacitated ordering policy if $p_s \geq p_r$, or a state dependent order-up-to policy if $p_s < p_r$. When $p_s \geq p_r$, the optimal ordering function $Q^*(d)$ has a finite upper bound, and as $d_{t-1} \rightarrow \infty$, the least upper bound of $Q^*(d)$ increases from S_s^{R*} to $z_s^{F*} - (F^0)^{-1}(\frac{h_r + p_r}{h_s + p_s})$. When $p_s < p_r$, the retailer follows a modified order-up-to policy in the restricted periods. As $d_{t-1} \rightarrow \infty$, $Q^*(d_{t-1}) \rightarrow \infty$ and the modified order-up-to level $S_r^{F*} - d_{t-1} + Q(d_{t-1})$ decreases from S_r^{F*} to $(F^0)^{-1}(\frac{h_r + p_s}{h_r + p_r})$. If $p_s < p_r$, $S_s^{R*} = 0$ and $z_s^{F*} = 0$, then $Q^*(d_{t-1})$ further degenerates to $Q(d) = (d - a)^+$, the functional form used in Gavirneni (2006) where a is a non-negative constant. The optimal parameters can be found by searching through all $\pi \in \mathcal{A}$. We propose an exhaustive enumeration search to find the optimal PF(2) policy illustrated in the following example.

Remark. Finding the optimal parameters $\pi^* = (S_s^{R*}, z_s^{F*}, S_r^{F*})$ can be computationally challenging because the supply chain cost under $Q^*(d)$ is non-convex in $\pi = (S_s^R, z_s^F, S_r^F)$. However, when the random demand has a finite density function, it is evident that the supply chain cost is a Lipschitz function of $\pi = (S_s^R, z_s^F, S_r^F)$, and standard approaches such as cutting angle method can be applied to solve for the optimal π (see Chapter 7 of Rubinov & Q.Yang (2003)).

Fig. 5. The optimal retailer restricted-ordering function.

4.2. Numerical examples for PF(2) policy

In this section and Section 4.3, we focus on numerical examples assuming $h_s \geq h_r$ because it is the setting in which the periodic flexible policies are more commonly effective. To illustrate, Fig. 4 (right) in Zhu, Gavirneni, Kapuscinski (2010) has shown that whenever the retailer's share of the decentralized supply chain cost is less than 70%, these periodic flexible policies are effective. Thus, what really matter is not a comparison of the relative holding costs, but rather a combination of both the holding and penalty costs at the retailer and supplier.

Example 2. We assume the same parametric settings for the retailer ($h_r = 1$, $p_r = 9$) as in Example 1, and demonstrate the optimal order quantity for two types of suppliers, in which the first type has cost parameters $h_s = 1.5$, $p_s = 19$, and the second type has $h_s = 1.5$, and $p_s = 12$. Fig. 5 illustrates the optimal retailer restricted-ordering function for both cases. When $p_s = 19$, the optimal parameters are $S_s^{R*} = 202$, $z_s^{F*} = 263$, and $S_r^{F*} = 244$. The optimal retailer restricted-ordering function is $Q^*(d) = d \wedge 202$, and it has a single turning point in the form of $Q(d) = d \wedge a$, where the constant $a = S_s^{R*}$. Under $Q^*(d)$, the retailer's ordering policy in the restricted period is a capacitated ordering policy, where the order cannot exceed the supplier's base-stock level, which will prevent any expediting costs at the supplier in the restricted period that would have occurred due to large demand realizations. In the free-ordering period, the retailer's optimal order-up-to level is 244, which is higher than that of the decentralized case (230). The extra units held in the flexible period can compensate for the inventory loss, when the demand in the restricted-ordering period is over 202. The supplier's optimal base-stock level in the free-ordering period is 263, almost equal to that under the decentralized setting (262), implying the operations in the free-ordering period have little effect on the supplier's decision in the restricted period. The system cost under $Q^*(d)$ is 1071.7, which is a 4.56% improvement from $Q(d) = d \wedge S_s^d$, and a 13.93% improvement from $Q(d) = d$.

When $p_s = 12$, the optimal parameters are $S_s^{R*} = 128$, $z_s^{F*} = 188$, and $S_r^{F*} = 265$. The optimal retailer restricted-ordering function is no longer a simple function with one turning point; rather, it has two flat regions and can be characterized as follows:

$$Q^*(d) = \begin{cases} d \wedge 128 & d \in [0, 356) \\ d + 34.6 - 100 \ln \left(-0.5 + \sqrt{0.25 + 540e^{-\frac{453-d}{100}}} \right) & d \in [356, 428) \\ 162.86 & d \in [428, \infty) \end{cases}$$

The first turning point is at $d = S_s^{R*} = 128$, and the subsequent turning points are at $d = 356$, and $d = 428$. Similarly to the first case in the restricted period, when $d < 128$, the retailer will order a quantity equal to the demand. When $d \in [128, 356)$, the retailer will order 128 to prevent the supplier from incurring back-order costs. However, as demand increases, the retailer's inventory in the free-ordering period decreases and the retailer's potential loss at the retailer due to low inventory increases. When $d = 356$ the retailer still orders 128, if his inventory in the following period is only 37 units, then a huge shortage cost may occur. When $d \in [356, 428)$, the retailer orders beyond the supplier's base-stock level to reduce his potential loss in the free-ordering period. When $d \geq 428$, because of the large p_s , the loss at the supplier exceeds the potential loss at the retailer, and hence the retailer's order is restricted to 162.86. The supplier's base-stock level in the restricted period is less than that when $p_s = 19$ because the retailer orders less on average in the restricted period, while in the free-ordering period, the retailer's base-stock level is higher than that when $p_s = 19$ because the demand variance is larger in the free-ordering period. The system cost under $Q^*(d)$ is 950.92.

Interestingly, the optimal restricted function $Q^*(d)$ is piecewise linear when $p_s = 19$, and almost piecewise linear when $p_s = 12$ (see Fig. 4). This functional structure raises the possibility that we can approximate $Q^*(d)$ using piecewise linear functions. More specifically, approximating $Q^*(d)$ with the piecewise linear function:

$$Q^{app}(d) = \begin{cases} d \wedge 128 & d \in [0, 356) \\ 0.48d + 43.38 & d \in [356, 428), \\ 162.86 & d \in [428, \infty) \end{cases}$$

the cost difference is 6.8, which is less than 1%. Thus, $Q^{app}(d)$ is near optimal.

We generate Fig. 4 by searching for only the integer values of the base-stock levels over $\mathcal{A}$. We then evaluate the cost function by taking the average of the objective values for over 50000 pairs of randomly selected demand realizations. The time to generate $Q^*(d)$ when $p_s = 19$ is 264 seconds, and when $p_s = 12$ is 1464 seconds.

Because computing $Q^*(d)$ is computationally prohibitive, we approximate $Q^*(d)$ with $Q(d) = d \wedge a^*$, where a^* is optimal among all forms of $Q(d) = d \wedge a$. Table 1 shows the efficiency of this approximation under different cost parameters. C_{a^*} is the expected supply chain cost under $Q(d) = d \wedge a^*$, C_e^* is the optimal expected supply chain cost under the optimal restricted-ordering function, and C^c is the supply chain cost under the centralized policy. We fix $h_r = 1$, and $p_r = 9$, and vary h_s and p_s . First, among all cases we studied, the average error from using $Q(d) = d \wedge a^*$ compared to $Q^*(d)$ is around 0.9%, implying that the approximation is very effective. Specifically, the efficiency of the approximation increases as p_s increases, and when $p_s = 19$, the optimal order function is in the form of $Q(d) = d \wedge a$, and hence the efficiency reaches 100%. Another observation is that z_s^{F*} increases in p_s .

Second, for our numerical examples, the optimal PF(2) policy has an average loss of 15.45% compared to the centralized policy. Specifically, when h_s and p_s are both large or both small, the percentage loss is large, with values above 15%. However, when either h_s or p_s is large while the other one is small, the percentage loss is much lower, with values around 10%, indicating that the PF(2) policy is more effective under these settings. However, even when the PF(2) policy is the most efficient, a 10% gap from the centralized supply chain cost still exists. In the following example, we explore the percentage improvement of the same approximation compared to the decentralized policy when both p_r and p_s are large.

Fig. 6. Percentage improvement versus the decentralized policy as a function of α .

Example 3. Assume the same parameters as in Example 1, that is, $h_r = 1$, $p_r = 9$, and $p_s = (1 + \alpha)p_r$. Fig. 6 shows the percentage cost improvement versus the decentralized policy under the order function $Q(d) = d \wedge S_s^{R*}$, where α is selected based on an increment of 0.2 and S_s^R and z_s^F take integer values. When $h_s = 1$, the percentage improvement is increasing in α . As $\alpha \rightarrow \infty$, the percentage improvement has an upper bound and converges to 16%. However, when $h_s > h_r$, the percentage improvement is no longer monotonous in α , and the PF(2) policy under $Q(d) = d \wedge S_s^{R*}$ achieves its maximum efficiency at a finite α . Specifically, when $h_s = 3$, the maximum percentage improvement under $Q(d) = d \wedge S_s^{R*}$ is achieved at $\alpha = 3(3.2)$ with a value of 27.51%.

4.3. The optimal PF(n) Policy and the numerical examples for PF(3) policy

When $n > 2$, the retailer orders according to the restricted-ordering functions in the first $n - 1$ periods but is free to order any quantity in the n th period. We can obtain the optimal PF(n) policy by applying calculus of variations. Under optimality, the retailer orders according to the same functional form across all the restricted periods and the optimal order quantity only depends on the preceding period's demand and the sum of the previous periods' demands in a cycle. Moreover, the optimal order quantity expressed as a function of the preceding period's demand has a functional structure similar to the optimal PF(2) policy that across all the restricted-ordering periods, the retailer either follows a state dependent capacitated policy or follows a state dependent order-up-to policy, and the order-up-to levels can be found by applying conventional optimization methods. (A characterization of the optimal PF(n) policy is provided in Appendix B.)

Near optimal PF(3) policies. We demonstrate through numerical studies that the optimal PF(3) policies on average have an inefficiency of less than 1% by approximating the optimal PF(3) policy with near optimal policies and using them as upper bounds to study the system inefficiency. Below, we provide two different approximations of the optimal PF(3) policies.

Example 4. We approximate the optimal retailer restricted-ordering function in the second restricted period $\tilde{Q}^*(d_t, d_{t-1})$ by $\tilde{Q}^*(d_t)$, which is a function of only period t 's demand d_t . Assume the same parametric settings for the retailer ($h_r = 1$, $p_r = 9$) as in Example 1 and $h_s = 1.5$ and $p_s = 19$, we illustrate the approximation of the retailer restricted-ordering function in the first restricted period $Q_t^*(d_{t-1})$ and in the second restricted period $\tilde{Q}^*(d_t)$.

Fig. 7. $Q_t^*(d_{t-1})$ and $\tilde{Q}^*(d_t)$.

Fig. 7 displays the optimal $Q_t^*(d_{t-1})$ and $\tilde{Q}^*(d_t)$. Under the optimal policy, the retailer's new ordering functions in both restricted periods are limited to $z_s^F - (F^0)^{-1}(\frac{h_s + p_r}{h_s + p_s})$. In both restricted periods, the retailer must follow a capacitated ordering policy where the order cannot exceed the supplier's base-stock level. This policy will prevent any expediting costs that would have otherwise been caused by large demand realizations from being incurred by the supplier in the restricted period.

In the free-ordering period, the retailer's optimal order up to level is 286, which is higher than 230 in the decentralized case. The extra units held in the flexible period can compensate for the reduction in inventory when the demand is over 108 in the first restricted-ordering period and 94 in the second restricted-ordering period. The supplier's optimal base-stock level in the free-ordering period is 194, which is lower than that under the decentralized setting (262) because in the restricted periods, the retailer orders less than the actual demand and a small z_s^F can lower the supplier's inventory in the restricted periods and, thus, reduce cost. The system cost under $Q_t^*(d_{t-1})$ and $\tilde{Q}^*(d_t)$ is less than 1% higher than that of the centralized system. If we further approximate the functional form of $Q_t^*(d_{t-1})$ and $\tilde{Q}^*(d_t)$ with $d_{t-1} \wedge a$ and $d_t \wedge b$, respectively, then the optimal parameters of the approximated restricted-ordering functions can be computed using classical optimization methods.

Example 5. In Table 2, we compare the performance of the optimal PF(2) policy and a near optimal PF(3) policy to the centralized case under the same setting as Example 2. First, we compare the inefficiency of $Q(d) = k^*$ with that of $Q^*(d)$ under PF(2), where k^* is the optimal fixed order quantity that minimizes the supply chain cost. The numbers in the third column illustrate an inefficiency of 35.66% on average under $Q(d) = k^*$ compared to the centralized system, which are on average 20% less efficient than $Q^*(d)$, indicating that $Q^*(d)$ can significantly improve the supply chain performance. Second, we compare the cost Table 3 improvement under the PF(3) and PF(2) policies. The rightmost column shows the percentage reduction on inefficiency from a specific type of PF(3) policy where $Q_t(d_{t-1}) = d_{t-1} \wedge a^*$, $Q_{t+1}(d_t, d_{t-1}) = d_{t-1} - Q_t(d_{t-1}) + d_t \wedge a^*$, and a^* is optimized over all nonnegative real numbers. This column computed from a specific type of PF(3) policy represents an upper bound on the reduction in efficiency for the optimal PF(3) policies. Interestingly, the average of the last column is 3.45%, which is significantly less than that under the optimal PF(2) policies, and diminishes when either h_s or p_s is large. Thus, the optimal PF policies can significantly improve the supply chain performance with more restriction periods.

Table 2
Efficiency of $Q(d) = d \wedge a^*$ and $Q^*(d)$.

h_s	p_s	S_s^R	z_s^F	S_r^F	a^*	C_e^*	C_{a^*}	C^*	$\frac{C_{a^*} - C_e^*}{C_e^*} \times 100\%$	$\frac{C^* - C_e^*}{C_e^*} \times 100\%$
1.01	9.0001	144	211	245	156	816.39	826.31	611.99	1.20	33.40
	12	146	219	251	165	854.48	864.89	758.93	1.20	12.59
	15	145	232	253	165	887.56	894.18	795.30	0.74	11.60
	19	221	265	238	180	927.91	927.91	843.79	0.00	9.97
1.5	9.0001	117	173	260	120	898.00	906.80	788.52	0.97	13.88
	12	128	188	265	129	950.90	960.20	824.89	0.97	15.28
	15	147	205	262	153	994.94	1004.30	861.26	0.93	15.52
	19	202	263	244	202	1071.70	1071.70	909.75	0.00	17.80
3	9.0001	81	117	270	78	1046.00	1056.60	946.70	1.00	10.49
	12	79	129	275	87	1137.00	1146.20	1007.23	0.81	12.88
	15	83	143	274	84	1214.50	1224.20	1063.18	0.79	14.23
	19	100	161	270	102	1308.40	1308.40	1111.67	0.00	17.70

Table 3

Percentage reduction in efficiency from PF(2) with $Q(d) = k^*$ and PF(3) with $Q_1(d_t) = d_t \wedge a^*$, $Q_2(d_t, d_{t+1}) = d_t - Q_1(d_t) + d_{t+1} \wedge a^*$ compared to centralized system.

h_s	p_s	PF2: $Q(d) = k^*$ %	PF3: $Q_1(d_t) = d_t \wedge a^*$, $Q_2(d_t, d_{t+1}) = d_t - Q_1(d_t) + d_{t+1} \wedge a^*$ %
1.01	9.0001	50.01	24.73
	12	27.70	4.03
	15	26.90	2.29
	19	22.81	0.00
1.5	9.0001	29.03	4.56
	12	33.36	3.44
	15	36.54	1.86
	19	35.82	0.00
3	9.0001	35.80	0.53
	12	38.85	0.00
	15	43.89	0.00
	19	47.24	0.00

The figures $h_s = 1.01$, $h_s = 1.5$ and $h_s = 3.0$ in Fig. 8 imply that (1) when h_s increases from 1.01 to 3.0, the difference between the percentage reduction in efficiency under the optimal fixed quantity and the optimal PF (both the optimal PF(2) and the PF(3) policies) increases. As h_s increases, the supply chain benefits more from applying optimal PF policies than the ad hoc fixed quantity policies. (2) For a fixed h_s , as p_s increases, the difference between the percentage reduction in efficiency under optimal PF(2) and PF(3) policies increases. Therefore, as p_s increases, the supply chain benefits more from applying optimal PF policies with a larger n . (3) When $h_s = 1.01$ and $p_s = 9.0001$, all three policies incur a reduction in efficiency of over 20%. Thus, when h_s and p_s are both small, the fixed quantity policy is inefficient, and the supply chain should adopt a PF policy with a larger n . Finally, the figure displaying the average percentage reduction in efficiency in Fig. 8 implies that, except when h_s and p_s are small, the performance of PF policies (optimal PF(2), PF(3) and $Q^* = k$) is relatively stable on the change of the cost coefficients.

The two examples above demonstrate that $Q(d) = d \wedge a^*$ is an effective approximation of Q^* for both PF(2) and PF(3) systems.

4.4. PF policy with free returns: what happens if $Q(d) > d$?

The previous results are all based on the assumption that $Q(d) \leq d$. From the characterization of $\hat{Q}(\cdot)$, we know that this assumption automatically holds if $S_r^{F*} = (F^0)^{-1}(\frac{h_r - h_s}{h_r + p_r}) < \infty$. However, when $S_r^{F*} < (F^0)^{-1}(\frac{h_r - h_s}{h_r + p_r})$, $\hat{Q}(\cdot)$ may exceed d and we can lower the system's cost by violating $Q(d) \leq d$.

Note that when $Q(d) > d$ is allowed, the system is no longer regenerative, because the on-hand inventory in some periods may

exceed the optimal base-stock levels. In this situation, we consider a free-returns problem in which the retailer is allowed to return, without any additional costs, any amount in the free-ordering periods, and the supplier is allowed to dispose of extra inventories, again, without any additional costs, in both the free-ordering and the restricted periods. Chen and Gavirneni (2010) show that the probability of these inventory transactions occurring is low, and even when they do occur, the magnitude of the transactions is also minute. In the presence of these options, it is optimal for both the supplier and the retailer to follow base-stock policies. Let $\tilde{S}_s^R$ be the order-up-to level at the supplier in the restricted-ordering periods; let $\tilde{S}_s^F$ be the order-up-to level at the supplier in the free-ordering periods; and let $\tilde{z}_s^F$ be the remaining inventory after adjusting to the known demand and or returns, i.e., $d_{t-1} - Q(d_{t-1})$, which can be computed as $\tilde{S}_s^F - (d_{t-1} - Q(d_{t-1}))$. The supply chain regenerates itself every two periods we refer to the optimization problem under this setting as the “free-returns problem.” The supply chain cost under the free-returns problem is a lower bound on that when these returns are not allowed. Similarly to Section 4.1, we formulate the free-returns problem as follows:

$$(\mathcal{P}') \min_{\pi \in \mathcal{A}} \left\{ \min_{\tilde{Q}(\cdot) \in \tilde{\mathcal{Q}}} E_{\tilde{d}_{t-1}} [\tilde{C}(\tilde{S}_s^R, \tilde{z}_s^F, \tilde{S}_r^F | \tilde{Q}(\cdot), \tilde{d}_{t-1})] \right\},$$

where $\tilde{\mathcal{Q}} = \{\tilde{Q}(d) | 0 \leq \tilde{Q}(d), \tilde{Q}(d) \in C[0, \infty)\}$, $\mathcal{A} = \{(\tilde{S}_s^R, \tilde{z}_s^F, \tilde{S}_r^F) | \tilde{S}_s^R \geq 0, \tilde{z}_s^F \geq 0, \tilde{S}_r^F \geq 0\}$, and

$$\begin{aligned} \tilde{C}(\tilde{S}_s^R, \tilde{z}_s^F, \tilde{S}_r^F | \tilde{Q}(\cdot), d_{t-1}) &= \tilde{C}_s(\tilde{S}_s^R, \tilde{z}_s^F | \tilde{Q}(\cdot), d_{t-1}) + \tilde{C}_r(\tilde{S}_r^F | \tilde{Q}(\cdot), d_{t-1}), \\ \tilde{C}_s(\tilde{S}_s^R, \tilde{z}_s^F | \tilde{Q}(\cdot), d_{t-1}) &= h_s [\tilde{S}_s^R - \tilde{Q}(d_{t-1})]^+ + p_s [\tilde{S}_s^R - \tilde{Q}(d_{t-1})]^- \\ &\quad + h_s E_{\tilde{d}_t} [\tilde{z}_s^F - \tilde{d}_t]^+ + p_s E_{\tilde{d}_t} [\tilde{z}_s^F - \tilde{d}_t]^- , \\ \tilde{C}_r(\tilde{S}_r^F | \tilde{Q}(\cdot), d_{t-1}) &= h_r E_{\tilde{d}_t} [\tilde{S}_r^F - d_{t-1} + \tilde{Q}(d_{t-1}) - \tilde{d}_t]^+ \\ &\quad + p_r E_{\tilde{d}_t} [\tilde{S}_r^F - d_{t-1} + \tilde{Q}(d_{t-1}) - \tilde{d}_t]^- \\ &\quad + h_r E_{\tilde{d}_{t+1}} [\tilde{S}_r^F - \tilde{d}_t]^+ + p_r E_{\tilde{d}_{t+1}} [\tilde{S}_r^F - \tilde{d}_t]^- . \end{aligned}$$

Similarly to Lemmas 2–4, we have the following characterization of the optimal parameters in the free-returns problem.

Lemma 5. The optimal $(\tilde{S}_s^{R*}, \tilde{z}_s^{F*}, \tilde{S}_r^{F*})$ for problem $(\mathcal{P}')$ is as follows:

1. The retailer's optimal base-stock level, $\tilde{S}_r^{F*}$ is $\tilde{S}_r^{F*} = \inf\{S | F^0(S) + (F \otimes F_{d-\tilde{Q}})^0(S) < \frac{2h_r}{h_r + p_r}\}$, where $F_{d-\tilde{Q}}$ is the cdf of $d_{t-1} - \tilde{Q}(d_{t-1})$.
2. The supplier's optimal base-stock level in the retailer restricted-ordering periods is $\tilde{S}_s^{R*} = \inf\{S | F_{\tilde{Q}}^0(S) < \frac{h_s}{h_s + p_s}\}$.
3. $\tilde{z}_s^{F*} = S_s^d$.

Applying calculus of variations, we have results similar to those in Theorem 1, which characterize the structure of the optimal

Fig. 8. Percentage reduction in efficiency compared to the centralized solution under different PF policies.

restricted-ordering function in the free returns problem. Let the optimal restricted-ordering function for (P') be $Q_R^*(d_{t-1})$.

Theorem 2. $Q_R^*(d_{t-1}) = [\tilde{Q}(d_{t-1})]^+$, where

$$\tilde{Q}(d_{t-1}) = \begin{cases} (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right) & d_{t-1} \in [0, \tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right)] \\ -\tilde{S}_r^{F*} + d_{t-1} & \\ \tilde{S}_s^{R*} & d_{t-1} \in [\tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right), \tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right)] \\ \tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) & \\ (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right) & d_{t-1} \in (\tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right), \infty) \\ -\tilde{S}_r^{F*} + d_{t-1} & \end{cases}$$

When $h_s \geq h_r$, the first term disappears, and when $p_s \geq p_r$, the last term disappears.

Theorem 2 shows that the optimal restricted-ordering function in the free-returns problem is a piecewise linear function. For example, when $d_{t-1} \in [0, \tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right)]$, the restricted-ordering function is increasing in d_{t-1} at a constant rate 1, and the retailer follows an order up to policy where the order up to level is $(F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right)$. When $d_{t-1} \in [\tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r - h_s}{h_r + p_r} \right), \tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right)]$, the optimal restricted-ordering function is a constant $\tilde{S}_s^{R*}$. This constant order quantity reduces the supplier's cost by reducing the uncertainty of her demand. When $d_{t-1} \in (\tilde{S}_s^{R*} + \tilde{S}_r^{F*} - (F^0)^{-1} \left(\frac{h_r + p_s}{h_r + p_r} \right), \infty)$, the optimal restricted-ordering function is again increasing at a constant rate 1, and the retailer's order up to level is $\tilde{S}_s^{R*}$. Note that when $h_s \geq h_r$, and $p_s \geq p_r$, the optimal restricted-ordering function degenerates to a constant $Q_R^*(d_{t-1}) = \tilde{S}_s^{R*}$ that is identical to the ordering function used in Chen and Gavirneni (2010), and Zhu et al. (2009).

5. Conclusion

This paper is a first attempt to characterize optimal periodic flexible policies for two-stage supply chains. We demonstrate a general way to find the optimal PF policies using calculus of variations, in particular, characterizing various structural properties of the two-period optimal periodic flexible policy and using them to develop a procedure for computing the optimal policy. Recognizing that determining optimal policies is computationally expensive, we identified and evaluated the effectiveness of a heuristic rule. Although we were able to make some progress on analyzing periodic flexible policies, many avenues for future research remain, the most important being the identification of the optimal periodic flexible policies for more complex supply chains.

Uncited references

Fudenberg and Tirole (1991); Gavirneni (2002); Glenn, Kilmer, and Stevens (2002); Kuribko, Lewis, Liu, and Song (2017).

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.ejor.2018.01.004](https://doi.org/10.1016/j.ejor.2018.01.004).

References

- Athey, S., & Segal, I. (2013). An efficient dynamic mechanism. *Econometrica*, 81, 2463–2485.
- Bergemann, D., & Välimäki, J. (2010). The dynamic pivot mechanism. *Econometrica*, 78, 771–789.
- Chen, L., & Gavirneni, S. (2010). Using scheduled ordering to improve the performance of distribution supply chains. *Management Science*, 56(9), 1615–1632.
- Fudenberg, D., & Tirole, J. (1991). *Game theory*. MIT Press.
- Gavirneni, S. (2002). Information flows in capacitated supply chains with fixed ordering costs. *Management Science*, 48(5), 644–651.
- Gavirneni, S. (2006). Price fluctuations, information sharing, and supply chain performance. *European Journal of Operational Research*, 174(3), 1651–1663.

- Gelfand, I. M., & Fomin, S. V. (2000). *Calculus of variations*. New York, USA: Dover Publications.
- Glenn, S. J., Kilmer, R. L., & Stevens, T. J. (2002). Supply chain management & production agriculture: a florida dairy case. *Choices*, 17(4), 2425.
- Kuribko, N., Lewis, T. R., Liu, F., & Song, J.-S. (2017). *Long term partnership for achieving efficient capacity allocation*. Working paper
- Lee, H. L., So, K., & Tang, C. (2000). The value of information sharing in a two-level supply chain. *Management Science*, 46(5), 626–643.
- Niranjan, T. T., Wagner, S. M., & Thakur-Weigold, B. (2011). Are you ready for VMI? *Industrial Engineer*, 43(2), 39–44.
- Plambeck, E. L., & Taylor, T. A. (2006). Partnership in a dynamic production system with unobservable actions and noncontractible output. *Management Science*, 52(10), 1509–1527.
- Rubinov, A. M., & Q. Yang, X. (2003). *Lagrange-type functions in constrained non-convex optimization*. Boston, USA: Kluwer Academic Publishers.
- Taylor, T. A., & Plambeck, E. L. (2007a). Simple relational contracts to motivate capacity investment: Price only vs. price and quantity. *Manufacturing & Service Operations Management*, 9(1), 94–113.
- Taylor, T. A., & Plambeck, E. L. (2007b). Supply chain relationships and contracts: The impact of repeated interaction on capacity investment and procurement. *Management science*, 53(10), 1577–1593.
- Zhu, W., Gavirneni, S., & Kapuscinski, R. (2009). Periodic flexibility, information sharing, and supply chain performance. *IIE Transactions*, 42(3), 173–187.