European Journal of Operational Research 228 (2013) 484–493

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

Managing inventories with one-way substitution: A newsvendor analysis

Yannick Deflem, Inneke Van Nieuwenhuyse *

Research Center for Operations Management, Department of Decision Sciences and Information Management, KU Leuven, Naamsestraat 69, 3000 Leuven, Belgium

ARTICLE INFO

Article history:
Received 3 April 2012
Accepted 3 February 2013
Available online 16 February 2013

Keywords:
Inventory
One-way substitution
Newsvendor

ABSTRACT

This paper presents an insightful approach to analyze two-item periodic inventory systems with one-way substitution. The objective is to minimize the expected total cost per period, which consists of expected purchasing costs, expected inventory holding costs, expected shortage costs, and expected adjustment costs. This approach helps derive the optimality conditions in both single-period and infinite horizon settings and yields useful insights into the impact of substitution on the service level, the optimality of a borderline case in which the order-up-to level of the inflexible item is reduced to zero, and the pivotal role of the purchasing cost.

1. Introduction

In many supply chains, mismatches between supply and demand can be (at least partially) mitigated by keeping inventories at different levels of the supply chain (e.g., raw materials, components, semifinished products, end items). The task of inventory management is to balance the benefits of inventory (i.e., reducing lost sales or limiting backorders) and the associated cost (which is typically reflected in the inventory holding cost).

One way to reduce the cost associated with inventory is to pool the demands for multiple items on the same (flexible) inventory item: provided that demands are not perfectly positively correlated, this process allows for a reduction in the required amount of safety stock and, thus, a reduction in inventory holding cost. This is referred to as "risk pooling" or "statistical economies of scale" (Van Mieghem, 2008). However, flexibility tends to come at a cost, which can be boiled down to a *product cost premium* (when the flexible item is inherently more expensive to manufacture or purchase) or an additional *adjustment cost* (when the item needs to undergo additional processing or transportation to make it "fit for use" when demand arises).

This observation has spurred research on so-called substitution (or "tailored-pooling") systems, in which flexible (and, thus, more expensive) stock is used as a substitute only when the regular (cheaper) item is out of stock. Tailored pooling can be obtained in various ways, such as through the use of *manufacturer-driven one-way substitution*[1] (e.g., Rutten and Bertrand, 1998; Bassok et al., 1999; Rao et al., 2004), *lateral transshipments* (e.g., Robinson, 1990; Herer et al., 2006), and *tailored postponement* (Tibben-Lembke and Bassok, 2005). It offers a compromise between a setting with shared inventory (when demand for a particular product type is always rerouted to the stock of the flexible product, and no product-specific stock is held) and separate inventories (when only product-specific stock is held, and demand can never be rerouted to stock of a different item).

In general, determining the optimal inventory control parameters in systems with substitution is complex: demands are only "partially pooled" on the inventory of the flexible item, and the amount of demand that can be rerouted to the flexible item depends on the order policies of both the dedicated product and the substitute. In this paper, we analyze a two-item inventory system with one-way substitution, assuming that both items are managed according to a *base stock* policy with periodic review. (As Bassok et al., 1999 show, this policy is optimal in multiproduct inventory problems with one-way substitution and zero setup costs.) The objective is to minimize the expected total cost per period, which consists of expected purchasing costs, expected inventory holding costs, expected shortage costs, and expected adjustment costs. The proposed framework enables us to (1) prove the cost conditions for which one-way substitution outperforms separate inventories, (2) present optimal first-order conditions for the respective order-up-to levels, and (3) discuss optimality conditions for a "borderline case" in which the order-up-to level of the inflexible item is set to zero.

The remainder of this paper proceeds as follows: in Section 2, we discuss the relevant literature. Section 3 further details the research problem and introduces notation. We explain the single-period newsvendor model in Section 4 and extend the findings to a setting with infinite time horizon in Section 5. In Section 6, we discuss the optimality of the borderline case. Finally, in Section 7, we summarize the main conclusions.

* Corresponding author. Tel.: +32 16 326743; fax: +32 16 326624.
E-mail addresses: yannick.deflem@econ.kuleuven.be (Y. Deflem), inneke.vannieuwenhuyse@econ.kuleuven.be (I. Van Nieuwenhuyse).

[1] Manufacturer-driven one-way substitution can be split further into component substitution (in an assemble-to-order setting, see e.g., Hale et al., 2001) or product substitution (at end-item level, e.g., Bassok et al., 1999).

0377-2217/$ - see front matter
http://dx.doi.org/10.1016/j.ejor.2013.02.003

2. Literature overview

The newsvendor approach is a popular method for analyzing periodic inventory systems with shared inventories (see, e.g., Hillier, 1999, 2000). In contrast, most studies considering periodic inventory systems with tailored pooling have used simulation optimization to determine the optimal order-up-to levels through either a gradient-based search or an exhaustive search (see, e.g., Robinson, 1990; Herer et al., 2006; Gong and Yücesan, 2012, for research on transshipment problems; see Khouja et al., 1996; Bassok et al., 1999; Rao et al., 2004, for related work on product substitution systems; and see Tibben-Lembke and Bassok, 2005, for research on delayed customization).

Studies that use a newsvendor approach to examine inventory systems with substitution are scarce. Hale et al. (2001) consider a setting with two end products, each comprised of two components, one of which can be downward substituted (i.e., component substitution). They succeed in deriving optimal first-order conditions, along with bounds on the order-up-to levels, but restrict attention to the single-period setting. Hillier (2002) considers the multiperiod "commonality-as-backup" setting, in which a common product acts as a backup for all unique products (this coincides with product substitution). However, it is impossible to find a closed-form expression for the optimal stocking level; only for the special case of two products with homogeneous costs and uniform demand distributions can the expected cost per period be derived (though not the critical fractile).

In addition, a separate stream of literature has emerged on continuous-review systems with tailored pooling (see, e.g., Bayindir et al., 2005; Liu and Lee, 2007, for product substitution; see Axsäter, 2003; Olsson, 2010; Van Wijk et al., 2012, for transshipment systems). The usual assumption in this research stream is that demands are uncorrelated, unit sized, and derived from Poisson processes, which enables analysis through continuous-time Markov chains.

In this paper, we focus on periodic (single-period and infinite horizon) inventory settings with substitution. We extend the work of Hale et al. (2001) and Hillier (2002) by (1) proving the cost conditions for which one-way substitution outperforms separate inventories, (2) presenting optimal first-order conditions for the respective order-up-to levels, and (3) discussing optimality conditions for a "borderline case" in which the order-up-to level of the inflexible item is set to zero (this coincides with shared inventories in the single-period case, though not in the infinite horizon case, as discussed subsequently). Our goal is to obtain analytical insights into the optimality conditions, and so we focus on a setting with two product types, one of which can act as a substitute for the other (i.e., product substitution). Our approach is inspired by Van Mieghem's (1998) work on optimal investment decisions in flexible capacity; note, however, that in Van Mieghem's study, the flexible resource acts purely as a backup for the dedicated resources and has no own demand to fulfill. This clearly differs from the inventory substitution setting. Throughout the analysis, we assume zero replenishment lead times (as is common in the literature; see, e.g., Khouja et al., 1996; Bassok et al., 1999).

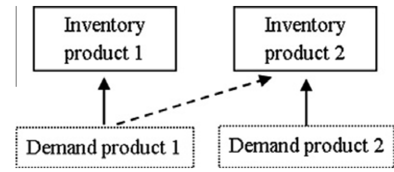

Fig. 1. A two-product inventory system with one-way substitution.

3. Problem description

We consider a setting with two product types (product 1 and product 2) as in Fig. 1. Demand d_i for a specific product type i in an arbitrary time period is preferably satisfied by means of its corresponding *product-specific* or *dedicated* inventory, as indicated by the solid arrows in Fig. 1. Only when the dedicated inventory for product 1 is out of stock can demand be satisfied by the substitute item (item 2); see the dashed arrow in Fig. 1. As such, part of the demand for item 1 can be "rerouted" to the stock of item 2; each unit of rerouted demand incurs a unit adjustment cost a.

Both inventories are managed according to a periodic base stock inventory policy. Fig. 2 shows the sequence of activities over a given period:

- At the start of every period, the decision maker places an order such that the inventory position is raised to the order-up-to level S_i (for $i = 1, 2$) (Chopra and Meindl, 2007). In a single-period setting, it is common to assume that the starting inventory position is zero (e.g., Heyman and Sobel, 1990; Khouja et al., 1996; Hale et al., 2001), whereas in the infinite horizon case, the inventory position at the start of a period equals the inventory position at the end of the previous period. Because the replenishment lead time is assumed to be zero, orders are received immediately; consequently, the net inventory immediately rises to S_i after an order for item i has been placed. The unit purchasing cost is represented by c_i for $i = 1, 2$.
- At the end of every period, the decision maker allocates the observed demand to the different inventories, constrained by earlier inventory investments. Any leftover inventory of product i at the end of the period incurs a unit holding cost h_i. Demand for product i that cannot be satisfied at the end of a period is penalized at a unit shortage cost p_i and is backordered to the next period (in the infinite horizon model) or lost (in the single-period model).

This problem fits the broader framework of a two-stage stochastic program. In the first stage, before demand is known, the optimal order-up-to levels are determined. In the second stage, after demand has been observed, the allocation decision is made.

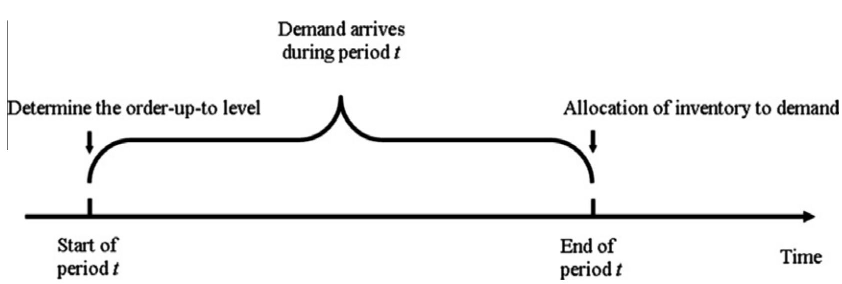

Fig. 2. Sequence of activities in an arbitrary period.