

INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Multiechelon Production/Inventory Systems: Optimal Policies, Heuristics, and Algorithms

Geert-Jan van Houtum

To cite this entry: Geert-Jan van Houtum. Multiechelon Production/Inventory Systems: Optimal Policies, Heuristics, and Algorithms. *In* INFORMS TutORials in Operations Research. Published online: 14 Oct 2014; 163-199.
<https://doi.org/10.1287/educ.1063.0026>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2006, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Multiechelon Production/Inventory Systems: Optimal Policies, Heuristics, and Algorithms

Geert-Jan van Houtum

Department of Technology Management, Technische Universiteit Eindhoven, P.O. Box 513, 5600 MB, Eindhoven, The Netherlands, g.j.v.houtum@tm.tue.nl

Abstract The theory on multiechelon production/inventory systems is a core theory within supply chain management. It provides useful insights for design of supply chains and may be used for tactical and operational planning decisions. The multiechelon theory started with the seminal paper of Clark and Scarf in 1960. In this tutorial, we describe for which systems optimal policies are known, which key features are needed for these optimal policy structures, and we discuss heuristics for systems of which the optimal policy structure is not known. We describe the complete analysis for the most basic multiechelon production/inventory system: The serial, two-echelon production/inventory system with linear inventory holding and backordering costs. We show that base-stock policies are optimal, derive a decomposition result for the determination of optimal base-stock levels, present newsboy equations for the optimal base-stock levels, and discuss computational procedures. Next, we describe a variety of systems for which generalized classes of base-stock policies have been derived to be optimal. This includes assembly systems and systems with fixed batch sizes, fixed replenishment intervals, generalized demand processes, and a service-level constraint instead of back-ordering costs. Finally, we discuss approaches that have been taken for distribution systems and systems with a general structure.

Keywords production/inventory; multiechelon; stochastic demand; stochastic dynamic programming; base-stock policies; newsboy equations

1. Introduction

Supply chain management is a broad area that covers strategic, tactical, and operational management decisions. The objective of a supply chain is to deliver products of the right quality, at the right time, in the right amount, and, preferably, with low costs. Two primary sources of costs in supply chains are *capacity costs* and *material costs*. Typically, capacity decisions are made for a longer term than material decisions; thus, capacity decisions are often made first, and material decisions follow. Material decisions may also be made sequentially, according to a *hierarchical approach* with two decision levels.

(i) A *first level* decides on such things as the form of batching, the batch sizes and replenishment intervals, and the (planned) lead times, where a *multi-item, multiechelon view* is taken. Via these decisions, one can accommodate setups, capacity constraints, capacity partitioning, and shipment consolidation. These decisions may be reviewed annually, for example;

(ii) A *second level* decides on reorder and base-stock levels, adapted on a daily, weekly, or monthly basis (e.g., when procedures like exponential smoothing are used for demand forecasting). Here, the batching rule is taken as given, and a *single-item, multiechelon view* can be incorporated.

The essential feature of this approach is that batching decisions are separated from safety stock decisions, as advocated by Graves [39]. For the second-level material decisions, excellent support may be provided by *multiechelon production/inventory models*. In addition,

the multiechelon models give insights into the effect of lead times, batch sizes, and demand uncertainty on total costs. They, thus, may also support first-level material decisions, capacity decisions, and design decisions (see also de Kok and Graves [17], Tayur et al. [59]).

The theory of multiechelon production/inventory decisions is the topic of this chapter. This theory was started by Clark and Scarf [14] in 1960. In their paper, a basic model for a supply chain consisting of multiple stages with a serial structure is considered. The stages are numbered $1, \dots, N$. Stage N orders at an external supplier, stage $N - 1$ orders at stage N , stage $N - 2$ orders at stage $N - 1$, and so on. Finally, at the most downstream stage, stage 1, external demand occurs. A stage may represent a production node, in which case input material is transformed into another product, or a transportation node, in which case a product is moved from one location to another. At the end of each stage, products can be kept on stock in a stockpoint, where they stay until they are demanded by either the next stage or the external customers. Time consists of periods of equal length, which may be days, weeks, or months, and the time horizon is infinite. Each stage is allowed to order at the beginning of each period. One can never order more than the amount available at the supplying stage, and the ordered amount by a stage n is assumed to arrive at the stockpoint at the end of stage n after a deterministic lead time. For the demand, a stationary, stochastic demand process is assumed. Costs consist of (linear) inventory-holding costs, which models the costs of working capital in the supply chain, and linear penalty costs for backordering, which constitute the counterpart for the inventory-holding costs. Clark and Scarf proved that so-called base-stock policies based on echelon inventory positions are optimal, and they showed that the optimal base-stock levels are obtained by the minimization of one-dimensional convex cost functions (this is known as the decomposition result). We refer to their model as the Clark-Scarf model.

Since 1960, much research has been executed to extend the work of Clark and Scarf. Extensions that have been considered are systems with a pure assembly/convergent structure, fixed batch sizes or fixed replenishment intervals, a service-level constraint, and advance demand information. Also, alternative approaches were developed to derive the main results for the Clark-Scarf model, which has contributed to a better understanding of which features are key to obtain the optimality of base-stock policies.

The objective of this tutorial is to expose for which systems optimal policies are known, which key features are needed to be able to derive the structure of optimal policies, and to discuss heuristics for systems of which the optimal policy structure is not known. We will start with a complete analysis of the most basic system: The two-echelon, serial system. From there on, we describe many extensions that have been made. For these extensions, generalized forms of base-stock policies have been shown to be optimal. This includes assembly/convergent systems. For distribution/divergent systems, base-stock policies are optimal under the so-called balance assumption, but they are not optimal without that assumption.

Systems with a general structure (i.e., with a mixed convergent-divergent structure) are most difficult. For those systems, concepts have been developed based on base-stock policies, and those concepts can be related to insights for basic systems (see §5.3). In the past few years, these concepts have been successfully applied in practice. In de Kok et al. [18], Graves and Willems [40], and Lin et al. [46], applications in large-scale projects at IBM, Eastman Kodak, and Philips Electronics have been reported. There are also several applications in smaller projects, and, currently, there is also commercial software available that is based on multiechelon theory. Generally, multiechelon theory is increasingly incorporated into the practice of supply chain management.

The foreknowledge that we assume is basic probability theory, basic inventory theory (e.g., Axsäter [3], Zipkin [71]), and stochastic dynamic programming (e.g., Porteus [49], Puterman [50]). This tutorial is intended to be accessible for anyone with that foreknowledge. It may also serve as a starting point for a Ph.D. course on multiechelon production/inventory systems, and for starting researchers in this research area.

The organization is as follows. In §2, we give a complete treatment of a two-echelon, serial system, and we denote the key features that lead to the optimality of base-stock policies, the decomposition result, and newsboy equations for optimal base-stock levels. Next, in §3, we describe the generalized results for multiechelon, serial systems, and we discuss exact and approximate procedures for the computation of an optimal policy and the corresponding optimal costs. In §4, we describe a variety of model variants and extended models for which pure or generalized forms of base-stock policies are optimal. This includes assembly/convergent systems and systems with a service-level constraint, fixed batch sizes, and fixed replenishment intervals. Then, in §5, we discuss systems with a distribution/divergent structure and systems with a mixed convergent-divergent structure. After that, in §6, we classify multiechelon systems as nice and complicated systems, and we conclude.

2. Analysis of the Two-Echelon, Serial System

In this section, we give a complete analysis of the two-echelon, serial system. In §2.1, we describe the model. Next, in §2.2, we derive the optimality of base-stock policies under general convex echelon cost functions, and we show that the optimal base-stock levels follow from the minimization of convex, one-dimensional functions (this is known as the decomposition result). Subsequently, in §2.3, for the common case with linear inventory holding and penalty costs, we derive simpler expressions in terms of so-called shortfalls and backlogs for these convex, one-dimensional functions. These alternative expressions facilitate computational procedures, and we use them to derive newsboy equations for the optimal base-stock levels.

2.1. Model


Consider a supply chain consisting of two stages, in which a single product is produced to stock. The upstream stage is called stage 2 and the downstream stage is called stage 1. Both stage 1 and stage 2 consists of a production step, a transportation step, or a network of such steps, with a stockpoint at the end of the stage. The stockpoint at the end of stage $n = 1, 2$ is called stockpoint n . For simplicity, we say that stage 2 is fed with raw materials, that an intermediate product is obtained from stage 2 and stored in stockpoint 2, and that a final product is obtained from stage 1 and stored in stockpoint 1. External demand occurs for the final product, i.e., at stockpoint 1.

Time is divided into periods of equal length. W.l.o.g., the length of each period, is assumed to be equal to 1. The time horizon that we consider is infinitely long. The periods are numbered $0, 1, \dots$, and denoted by the index t ($t \in \mathbb{N}_0 := \{0\} \cup \mathbb{N}$).

Both stages or stockpoints are allowed to place orders at the beginning of each period. An amount ordered by stage 2 at the beginning of a period t arrives at stockpoint 2 after a deterministic lead time $l_2 \in \mathbb{N}$. We assume that sufficient raw material is always available, and, thus, orders by stockpoint 2 are never delayed. An amount ordered by stage 1 at the beginning of a period t arrives at stockpoint 1 after a deterministic lead time $l_1 \in \mathbb{N}_0$ ($l_1 = 0$ is allowed), provided that there is sufficient stock at stockpoint 2 available at the beginning of period t . If the available stock is smaller than the ordered amount, then the available amount is sent into stage 1 and becomes available after l_1 periods, while the rest is delivered as soon as possible.

The demands in different periods are independent and identically distributed on $[0, \infty)$. The cumulative demand over periods $t_1, \dots, t_2$, $0 \leq t_1 \leq t_2$, is denoted by D_{t_1, t_2} . F is the generic distribution function for the demand $D_{t, t}$ in an arbitrary period $t \in \mathbb{N}_0$. The mean demand per period is $\mu > 0$. We implicitly assume that we have a continuous product and that order sizes and inventory levels are real-valued variables. The demand distribution function, however, is not necessarily continuous. There may be positive probability masses at specific points. In the case of a discrete product, it is more natural to limit order sizes and inventory levels to integer values. That case is discussed in §4.2.

FIGURE 1. The serial, two-echelon production/inventory system.



A picture of the serial, two-echelon system is given in Figure 1. We have the following events in each period.

- (i) at each stage, an order is placed;
- (ii) arrival of orders;
- (iii) demand occurs; and
- (iv) one-period costs are assessed (these costs are specified below).

The first two events take place at the beginning of the period, and the order of these two events may be interchanged, except for the most downstream stage when its lead time equals 0. The last event occurs at the end of a period. The third event, the demand, may occur anywhere in between.

2.1.1. Echelon Stocks and Costs Attached to Echelons. The analysis of multiechelon systems is generally based on the concepts *echelon stock* and *echelon inventory position*, as introduced by Clark [13] in 1958 (see also Zipkin [71], pp. 120–124). Below, we describe these concepts and define costs attached to echelons.

In general, the echelon stock (or *echelon inventory level*) of a given stockpoint denotes all physical stock at that stockpoint plus all materials in transit to or on hand at any stockpoint downstream minus eventual backlogs at the most downstream stockpoints. The chain under consideration is called the *echelon*. An echelon stock may be negative, indicating that the total backlog at the most downstream stockpoints is larger than the total physical stock in that echelon. Echelons are numbered according to the highest stockpoint in that echelon. In our two-echelon system, we have two echelons:


- echelon 1, consisting of stockpoint 1; and
- echelon 2, consisting of stockpoint 2, stockpoint 1, and the pipeline in between.

The echelon stock of echelon 1 is also called *echelon stock 1*, and is the same as the installation stock of stockpoint 1. The echelon stock of echelon 2 is also called *echelon stock 2*.

The *echelon inventory position* of a stockpoint is defined as its echelon stock plus all materials that are in transit to the stockpoint. We assume that a stockpoint never orders more than what is available at the next upstream stockpoint. In our two-echelon system, this implies that stockpoint 1 never orders more than what is available at stockpoint 2. As we study the optimal behavior of the system under centralized control, this assumption can be made w.l.o.g.; instead of creating a backlog position at stockpoint 2, stockpoint 1 will attempt to order that difference at the next period. Under this assumption, the echelon inventory position is also equal to the echelon stock plus all materials on order. The echelon inventory position of echelon n is also called *echelon inventory position n* , $n = 1, 2$. The echelon stocks and echelon inventory positions are visualized in Figure 2.

We now define our costs, which are assessed at the *end* of each period, based on the echelon stocks. For $n = 1, 2$, we pay costs $c_n(x_n)$, where x_n denotes echelon stock n at the *end* of a period. Notice that, by the above definitions, it holds that $x_2 \geq x_1$. The function $c_n(x_n)$ denotes the *costs attached to echelon n* . We assume that the cost functions $c_n(x_n)$, $n = 1, 2$, are *convex*. In addition, to avoid mathematical complexities, we assume that these cost functions are such that it is suboptimal to let the backlog grow to infinity. That one-period costs can be expressed as the sum of separable, convex functions based on echelon stocks is a crucial assumption. This was already pointed out by Clark and Scarf [14] (Assumption 3, pp. 478–479).

FIGURE 2. The concepts echelon stock and echelon inventory position.



A *special cost structure* is obtained when we have *linear inventory-holding and penalty costs*. That structure is often assumed and is as follows. A cost of $h_2 \geq 0$ is charged for each unit that is on stock in stockpoint 2 at the end of a period and for each unit in the pipeline from stockpoint 2 to stockpoint 1. A cost of $h_1 + h_2 \geq 0$ is charged for each unit that is on stock in stockpoint 1 at the end of a period. The inventory-holding cost parameters represent interest and storage costs. We assume that the additional inventory-holding cost at stage 1 is nonnegative, i.e., $h_1 \geq 0$. A penalty cost p is charged per unit of backordered demand at stockpoint 1 at the end of a period. This represents inconvenience for delayed fulfillment of demand and constitutes the counterpart for the inventory-holding costs. We assume that $p > 0$.

Let x_n , $n = 1, 2$, be echelon stock n at the end of a period. Then, the total inventory holding and backordering costs at the end of a period are equal to

$$h_2(x_2 - x_1) + (h_1 + h_2)x_1^+ + px_1^-,$$

where $x^+ = \max\{0, x\}$ and $x^- = \max\{0, -x\} = -\min\{0, x\}$ for any $x \in \mathbb{R}$. These costs may be rewritten as

$$\begin{aligned} & h_2(x_2 - x_1) + (h_1 + h_2)x_1^+ + px_1^- \\ &= h_2(x_2 - x_1) + (h_1 + h_2)x_1 + (p + h_1 + h_2)x_1^- \\ &= h_2x_2 + h_1x_1 + (p + h_1 + h_2)x_1^- \\ &= c_2(x_2) + c_1(x_1), \end{aligned}$$

with

$$c_1(x_1) = h_1x_1 + (p + h_1 + h_2)x_1^-, \quad (1)$$

$$c_2(x_2) = h_2x_2. \quad (2)$$

This shows that the case with linear inventory holding and penalty costs fits under the general cost structure. In this special case, $c_2(x_2)$ is linear and $c_1(x_1)$ is a convex function consisting of two linear segments. In the analysis below (in §2.2), we assume the general cost structure. After that, we derive additional results that hold under linear inventory holding and penalty costs (in §2.3).

2.1.2. Objective. Let Π denote the set of all possible ordering policies, and let $G(\pi)$ denote the average costs of ordering policy π for all $\pi \in \Pi$. We want to solve the following minimization problem to optimality.

$$\begin{aligned} (\text{P}): \quad & \min G(\pi) \\ \text{s.t.} \quad & \pi \in \Pi. \end{aligned}$$

So, the objective is to find an ordering policy under which the average costs per period are minimized.

2.2. Analysis

In this subsection, we derive the optimality of base-stock policies and the decomposition result. These results are due to Clark and Scarf [14], who derived these results via a stochastic dynamic program in a finite-horizon setting. Federgruen and Zipkin [29] extended these results to the infinite-horizon case. Alternative, easier proofs were developed by Langenhoff and Zijm [45] and by Chen and Zheng [12] (see also Chen [10]). We follow the approach of Chen and Zheng, where we add an explicit definition of a relaxed single-cycle problem (cf. van Houtum et al. [66] for a generalized system; Chen and Zheng have an implicit definition). We distinguish three steps:

1. definition of cycles and cycle costs;
2. solution of a relaxed single-cycle problem; and
3. solution of the infinite-horizon problem (P).

These steps are described in §§2.2.1–2.2.3. The introduction of the relaxed single-cycle problem and the property that the solution of the single-cycle problem also solves the infinite-horizon problem (P) are key in the line of proof. Interestingly, the relaxed single-cycle problem is a stochastic dynamic programming problem with a finite number of stages (two stages in this case). Thus, the solution of problem (P), which is a stochastic dynamic programming problem with an infinite horizon, follows in fact from a finite-horizon stochastic programming problem.

2.2.1. Step 1: Definition of Cycles and Cycle Costs. We consider the connection between order decisions at the two stages, and we describe which costs they affect.

For each $n = 1, 2$ and $t \in \mathbb{N}_0$, let $IL_{t,n}$ and $IP_{t,n}$ denote echelon stock n (= echelon inventory level n) and echelon inventory position n at the beginning of period t (just before the demand occurs), and let $C_{t,n}$ denote the costs attached to echelon n at the end of period t .

We now consider the following two connected decisions, starting with an order placed by stage 2 at the beginning of a period $t_0 \in \mathbb{N}_0$:

- *Decision 2:* Decision 2 concerns the decision at the beginning of period t_0 with respect to the order placed by stage 2. Suppose that this order is such that $IP_{t_0,2}$ becomes equal to some level z_2 . First of all, this decision directly affects the echelon 2 costs at the end of period $t_0 + l_2$. The expected value of these costs equals

$$\mathbb{E}\{C_{t_0+l_2,2} | IP_{t_0,2} = z_2\} = \mathbb{E}\{c_2(z_2 - D_{t_0,t_0+l_2})\}. \quad (3)$$

Second, by this decision, echelon stock 2 at the beginning of period $t_0 + l_2$ becomes equal to $IL_{t_0+l_2,2} = z_2 - D_{t_0,t_0+l_2-1}$, and this directly limits the level to which one can increase the echelon inventory position $IP_{t_0+l_2,1}$ of echelon 1 at the beginning of period $t_0 + l_2$. This is the second decision to consider.


- *Decision 1:* Decision 1 concerns the order placed by stage 1 at the beginning of period $t_0 + l_2$. Suppose that by this order, $IP_{t_0+l_2,1}$ becomes equal to some level z_1 . This decision directly affects the echelon 1 costs at the end of period $t_0 + l_2 + l_1$. The expected value of these costs equals

$$\mathbb{E}\{C_{t_0+l_2+l_1,1} | IP_{t_0+l_2,1} = z_1\} = \mathbb{E}\{c_1(z_1 - D_{t_0+l_2,t_0+l_2+l_1})\}. \quad (4)$$

Figure 3 visualizes the way in which the above decisions affect each other, and which costs are determined by them.

In the description above, we have explicitly described for decision 1 how the level z_1 to which $IP_{t_0+l_2,1}$ is increased is bounded from above. We will need this in the analysis below. Obviously, for both decisions 2 and 1, it also holds that the levels z_2 and z_1 to which $IP_{t_0,2}$ and $IP_{t_0+l_2,1}$ are increased, are bounded from below (by the level that one already has for

FIGURE 3. The consequences of the decisions 1 and 2.



its echelon inventory position just before the new order is placed). In the analysis below, this is taken into account too. But, this bounding from below will appear to be less important.

The decisions 2 and 1 start with decision 2 taken in period t_0 . These decisions constitute a *cycle*, and the corresponding expected costs are equal to

$$C_{t_0} := C_{t_0+l_2,2} + C_{t_0+l_2+l_1,1}.$$

These costs are defined for each period $t_0 \in \mathbb{N}_0$, and we call them the *total costs attached to cycle t_0* . For each positive recurrent policy $\pi \in \Pi$, the average costs are equal to the average value of the costs C_{t_0} over all cycles t_0 .

$$\begin{aligned} G(\pi) &= \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left\{ \sum_{t=0}^{T-1} (C_{t,2} + C_{t,1}) \right\} \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left\{ \sum_{t=0}^{T-1} C_t + \sum_{t=0}^{l_2-1} C_{t,2} + \sum_{t=0}^{l_2+l_1-1} C_{t,1} - \sum_{t=T}^{T+l_2-1} C_{t,2} - \sum_{t=T}^{T+l_2+l_1-1} C_{t,1} \right\} \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} C_t \end{aligned} \quad (5)$$

2.2.2. Step 2: Solution of a Relaxed Single-Cycle Problem. Consider a cycle t_0 . We now consider how the decisions 1 and 2 can be taken such that the expected total costs attached to cycle t_0 ($= \mathbb{E} C_{t_0}$) are minimized. Decision n , $n = 1, 2$, is described by the level z_n , to which echelon inventory position n is increased at the beginning of period t_0 and $t_0 + l_2$, respectively. The choice for the level z_1 is limited from above by what is available at stage 2. Further, the choice for the level z_n , $n = 2, 1$, is bounded from below by the value of echelon inventory position n just before the order is placed. By neglecting the bounding from below, we obtain the following *relaxed problem*:

$$\begin{aligned} (\text{RP}(t_0)) : \quad & \text{Min} \quad \mathbb{E} C_{t_0} = \mathbb{E} C_{t_0+l_2,2} + \mathbb{E} C_{t_0+l_2+l_1,1} \\ & \text{s.t.} \quad \mathbb{E} C_{t_0+l_2,2} = \mathbb{E} \{ c_2(z_2 - D_{t_0,t_0+l_2}) \}, \\ & \quad \mathbb{E} C_{t_0+l_2+l_1,1} = \mathbb{E} \{ c_1(z_1 - D_{t_0+l_2,t_0+l_2+l_1}) \}, \\ & \quad z_1 \leq IL_{t_0+l_2,2}, \\ & \quad IL_{t_0+l_2,2} = z_2 - D_{t_0,t_0+l_2-1}. \end{aligned}$$

Problem $(\text{RP}(t_0))$ is a two-stage stochastic dynamic programming problem. Decision 2 is described by z_2 and is not limited at all; we, thus, may connect this decision to a dummy starting state. The resulting direct expected costs are equal to $\mathbb{E} \{ c_2(z_2 - D_{t_0,t_0+l_2}) \}$. Decision 1 is described by z_1 , and, via the constraint $z_1 \leq IL_{t_0+l_2,2}$, its decision space depends on the echelon stock 2 at the beginning of period $t_0 + l_2$, i.e., on $IL_{t_0+l_2,2}$. Hence, we use

$IL_{t_0+l_2,2}$ to describe the state of the system when decision 1 is taken. This state depends on decision 2 via the relation $IL_{t_0+l_2,2} = z_2 - D_{t_0,t_0+l_2-1}$. Decision 1 results in direct expected costs $E\{c_1(z_1 - D_{t_0+l_2,t_0+l_2+l_1})\}$.

For problem $(RP(t_0))$, we first determine what is optimal for decision 1, and after that we consider decision 2.

Let the function $G_1(y_1)$ be defined by

$$G_1(y_1) := E\{c_1(y_1 - D_{t_0+l_2,t_0+l_2+l_1})\}, \quad y_1 \in \mathbb{R}. \quad (6)$$

This function denotes the expected costs attached to echelon 1 at the end of a period $t_0 + l_1 + l_2$ if echelon inventory position 1 at the beginning of period $t_0 + l_2$ (i.e., l_1 periods earlier) has been increased up to level y_1 .

Lemma 1 (On the Optimal Choice for z_1). *It holds that*

- (i) $G_1(y_1)$ is convex as a function of y_1 , $y_1 \in \mathbb{R}$.
- (ii) Let $S_1 \in (\mathbb{R} \cup \{\infty\})$ be chosen such that

$$S_1 := \arg \min_{y_1 \in \mathbb{R}} G_1(y_1).$$

Then, for the problem $(RP(t_0))$, it is optimal to choose the level z_1 equal to S_1 , or as high as possible if this level cannot be reached.

Proof. The formula for $G_1(y_1)$ may be rewritten as

$$G_1(y_1) = \int_0^\infty c_1(y_1 - x) dF_{l_1+1}(x),$$

where F_{l_1+1} is the $(l_1 + 1)$ -fold convolution of F . Let $y_1^1, y_1^2 \in \mathbb{R}$, and $\alpha \in [0, 1]$, then, by the convexity of $c_1(\cdot)$,

$$\begin{aligned} G_1(\alpha y_1^1 + (1 - \alpha)y_1^2) &= \int_0^\infty c_1(\alpha(y_1^1 - x) + (1 - \alpha)(y_1^2 - x)) dF_{l_1+1}(x) \\ &\leq \int_0^\infty [\alpha c_1(y_1^1 - x) + (1 - \alpha)c_1(y_1^2 - x)] dF_{l_1+1}(x) \\ &= \alpha G_1(y_1^1) + (1 - \alpha)G_1(y_1^2), \end{aligned}$$

and, thus, $G_1(y_1)$ is convex. This proves Part (i).

Next, S_1 is defined as the point where $G_1(y_1)$ is minimized. If there are multiple points where $G_1(y_1)$ is minimized, then S_1 may be taken equal to any of these points. We can now show how decision 1, i.e., the choice for z_1 , may be optimized for problem $(RP(t_0))$. This decision is taken at the beginning of period $t_0 + l_2$, and the choice for z_1 is bounded from above by $IL_{t_0+l_2,2}$. This decision only affects the costs $EC_{t_0+l_2+l_1,1}$, which, by (6), are equal to $G_1(z_1)$. As the function G_1 is convex, these costs are minimized by choosing z_1 equal to $z_1 = S_1$ if $IL_{t_0+l_2,2} \geq S_1$, and equal to $z_1 = IL_{t_0+l_2,2}$ if $IL_{t_0+l_2,2} < S_1$. This completes the proof of Part (ii). $\square$

By Lemma 1, for decision 1, it is optimal to apply *base-stock policy* S_1 (i.e., a base-stock policy with base-stock level S_1). Let $G_2(y_1, y_2)$ be defined as the expected cycle costs when a base-stock policy with level $y_2 \in \mathbb{R}$ is applied for decision 2 and a base-stock policy $y_1 \in \mathbb{R}$ for decision 1 (notice that we allow that $y_2 < y_1$ and y_1 and y_2 may also be negative). Then, $z_2 = y_2$, as the external supplier can always deliver, and for z_1 , we find

$$z_1 = \min\{IL_{t_0+l_2,2}, y_1\} = \min\{y_2 - D_{t_0,t_0+l_2-1}, y_1\}.$$

Hence,

$$G_2(y_1, y_2) = E\{c_2(y_2 - D_{t_0,t_0+l_2}) + c_1(\min\{y_2 - D_{t_0,t_0+l_2-1}, y_1\} - D_{t_0+l_2,t_0+l_2+l_1})\}, \quad y_1, y_2 \in \mathbb{R}. \quad (7)$$

Lemma 2 (On the Optimal Choice for z_2). *It holds that*

- (i) $G_2(S_1, y_2)$ is convex as a function of y_2 , $y_2 \in \mathbb{R}$.
- (ii) Let $S_2 \in \mathbb{R} \cup \{\infty\}$ be chosen such that

$$S_2 := \arg \min_{y_2 \in \mathbb{R}} G_2(S_1, y_2).$$

Then, for problem $(RP(t_0))$, it is optimal to choose the level z_2 equal to S_2 .

Proof. Let F_{l_2} be the l_2 -fold convolution of F . The formula for $G_2(S_1, y_2)$ may be rewritten as

$$\begin{aligned} G_2(S_1, y_2) &= \mathbb{E}\{c_2(y_2 - D_{t_0, t_0+l_2})\} + \int_0^\infty \mathbb{E}\{c_1(\min\{y_2 - x, S_1\} - D_{t_0+l_2, t_0+l_2+l_1})\} dF_{l_2}(x) \\ &= \mathbb{E}\{c_2(y_2 - D_{t_0, t_0+l_2})\} + G_1(S_1) + \int_0^\infty [G_1(\min\{y_2 - x, S_1\}) - G_1(S_1)] dF_{l_2}(x) \\ &= \mathbb{E}\{c_2(y_2 - D_{t_0, t_0+l_2})\} + G_1(S_1) + \int_0^\infty \tilde{G}_1(y_2 - x) dF_{l_2}(x), \end{aligned} \quad (8)$$

where

$$\tilde{G}_1(y) = G_1(\min\{y, S_1\}) - G_1(S_1) = \begin{cases} G_1(y) - G_1(S_1) & \text{if } y < S_1, \\ 0 & \text{if } y \geq S_1. \end{cases}$$

Because $G_1(\cdot)$ is convex, with a minimum in S_1 , also the function $\tilde{G}_1(y)$ is convex. Hence, along the same lines as for Part (i) of Lemma 1, the first and third term in (8) may be shown to be convex. This implies that $G_2(S_1, y_2)$ is convex as a function of y_2 , which completes the proof of Part (i).

Next, S_2 is defined as the point that minimizes $G_2(S_1, y_2)$ as a function of y_2 . If there are multiple points where $G_2(S_1, y_2)$ is minimized, then S_2 may be taken equal to any of these points. We can now show how decision 2, i.e., the choice for z_2 , may be optimized for problem $(RP(t_0))$. This decision is taken at the beginning of period t_0 . This decision affects the costs $\mathbb{E}C_{t_0+l_2, 2}$ and $\mathbb{E}C_{t_0+l_2+l_1, 1}$. Whatever choice is made for z_2 , it is optimal to take decision 1 according to a base-stock policy with base-stock level S_1 (by Part (ii) of Lemma 1). Hence, by (7),

$$\mathbb{E}C_{t_0+l_2, 2} + \mathbb{E}C_{t_0+l_2+l_1, 1} = G_2(S_1, y_2).$$


These costs are minimized by choosing z_2 equal to $z_2 = S_2$. This completes the proof of Part (ii). $\square$

By Lemmas 1 and 2, for decisions 2 and 1 of problem $(RP(t_0))$, it is optimal to apply a base-stock policy with base-stock level S_2 and S_1 , respectively. The corresponding optimal costs are equal to $G_2(S_1, S_2)$. Because this problem was obtained by neglecting the bounding from below when placing orders, the optimal costs $G_2(S_1, S_2)$ constitute a lower bound for the optimal costs of the original problem (P).

2.2.3. Step 3: Solution of the Infinite-Horizon Problem (P). The functions $G_1(y_1)$ and $G_2(y_1, y_2)$ as defined above, have alternative interpretations; $G_1(y_1)$ represents the average costs of a base-stock policy y_1 in a specific single-echelon system, called *subsystem 1*, while $G_2(y_1, y_2)$ represents the average costs of a base-stock policy (y_1, y_2) for the full two-echelon system as depicted in Figure 1. This is shown first, and after that, we solve the original problem (P).

Subsystem 1 is defined as the downstream part of the full two-echelon system. It consists of stage 1 only, and it has stockpoint 2 as external supplier with infinite supply. The lead time for this subsystem is l_1 , the demands are the same as in the full system, and the costs consist of the costs attached to echelon 1; see Figure 4. Under a base-stock policy y_1

FIGURE 4. Subsystem 1.



($y_1 \in \mathbb{R}$), at the beginning of each period, nothing is ordered if the current inventory position is already at level y_1 or higher, and the inventory position is increased up to level y_1 if the current inventory position is lower than y_1 . That the inventory position before ordering is above the base-stock level y_1 may only happen in a limited number of periods. Hence, in steady state, the inventory position is always increased up to level y_1 , and, therefore, the average costs are equal to $G_1(y_1) = \mathbf{E}\{c_1(y_1 - D_{t_0+l_2, t_0+l_2+l_1})\}$.

Consider now a base-stock policy (y_1, y_2) , $y_1, y_2 \in \mathbb{R}$, for the full two-echelon system (we allow that $y_2 < y_1$). Under this policy, at the beginning of each period, stage 2 orders nothing if the current echelon inventory position 2 is already at level y_2 or higher, and its echelon inventory position is increased up to level y_2 if the current position is lower than y_2 . That echelon inventory position 2 before ordering is above the base-stock level y_2 may only happen in a limited number of periods. Hence, in steady state, echelon inventory position 2 is always increased up to level y_2 . Similarly, at the beginning of each period, stage 1 orders nothing if the current echelon inventory position 1 is already at level y_1 or higher, and, one aims to increase up to level y_1 if the current position is lower than y_1 . In the latter case, it may not be possible to increase up to y_1 because there is not sufficient material available in stockpoint 2. That echelon inventory position 1 before ordering is above the base-stock level y_1 may only happen in a limited number of periods. Hence, in steady state, we obtain that echelon inventory position 1 is increased up to level y_1 if echelon stock 2 is at least y_1 at that moment, and up to echelon stock 2 otherwise. Hence, in steady state, we obtain per cycle that the ordering behavior is precisely as depicted in Figure 3 in which base-stock policies with levels y_2 and y_1 are applied for decisions 2 and 1, respectively. Hence, the average costs of a base-stock policy (y_1, y_2) are given by the function $G_2(y_1, y_2)$.

Suppose now that base-stock policy (S_1, S_2) is used for the original problem (P). Then average costs $G_2(S_1, S_2)$ are obtained, and these costs are, thus, equal to the lower bound. This implies that base-stock policy (S_1, S_2) is optimal for the original problem (P). In fact, if base-stock policy (S_1, S_2) is used in all periods, then the lower bounds that were relaxed in problem (RP(t_0)) are only binding during a transient period (when the echelon inventory positions may be above S_1 and S_2 , and nothing should be ordered). In the long run, these lower bounds are not binding and, thus, the optimal solutions of the relaxed and unrelaxed problem are identical.

Theorem 1. *Base-stock policy (S_1, S_2) , with the S_i as defined in Lemmas 1 and 2, is optimal for problem (P).*

This theorem shows that the class of base-stock policies is optimal, and that the optimal base-stock levels can be obtained sequentially by the minimization of one-dimensional functions. The latter result is known as the *decomposition result*.

Notice that it may happen that $S_2 < S_1$. As stated above, for base-stock policies (y_1, y_2) in general, we allow that $y_2 < y_1$, i.e., that the base-stock level for echelon inventory position 1 is larger than the base-stock level for echelon inventory position 2. Nevertheless, in practice, it is more natural to use and communicate a base-stock policy (y_1, y_2) with $y_2 \geq y_1$. The following lemma shows that any base-stock policy (y_1, y_2) with $y_2 < y_1$ can be translated into a base-stock policy $(\tilde{y}_1, y_2)$ with $y_2 \geq \tilde{y}_1$ and equal average costs.

Lemma 3. *Let $y_1, y_2 \in \mathbb{R}$, and define $\tilde{y}_1 := \min\{y_1, y_2\}$. Then $G_2(\tilde{y}_1, y_2) = G_2(y_1, y_2)$.*

Proof. Let $y_1, y_2 \in \mathbb{R}$, and define $\tilde{y}_1 := \min\{y_1, y_2\}$. That $G_2(\tilde{y}_1, y_2) = G_2(y_1, y_2)$ is trivial in case $y_2 \geq y_1$, because then $\tilde{y}_1 = y_1$. In case $y_2 < y_1$, at the beginning of each period, stage 1 is confronted with a shortage of material at stockpoint 2, and all available material at stockpoint 2 will be forwarded into stage 2. This implies that stockpoint 2 is a stockless stockpoint. This will still be so if base-stock level y_1 is decreased to $\tilde{y}_1 = y_2$. Hence, under base-stock policy $(\tilde{y}_1, y_2)$, the orders are identical to the orders generated under base-stock policy (y_1, y_2) (at least in the long run; in the first periods of the horizon, there may be differences). Thus, both policies have the same average costs. An alternative, technical proof is obtained by (7): If $y_2 < y_1$, then

$$\begin{aligned} G_2(y_1, y_2) &= \mathbb{E}\{c_2(y_2 - D_{t_0, t_0+l_2}) + c_1(y_2 - D_{t_0, t_0+l_2-1} - D_{t_0+l_2, t_0+l_2+l_1})\} \\ &= G_2(y_2, y_2) = G_2(\tilde{y}_1, y_2). \quad \square \end{aligned}$$

This completes the whole analysis for the two-echelon serial system. All results are easily extended to serial systems with more than two stages. Proofs go by induction, where the induction step is identical to what we derived for stage 2 in this two-echelon system.

Remark 1 (Induced Penalty Cost Function). Equation (8) for $G_2(S_1, y_2)$ consists of three terms. The first term denotes the costs attached to echelon 2. The second term, $G_1(S_1)$, denotes the minimal costs for subsystem 1. The third term denotes the additional costs when echelon stock 2 is insufficient to increase echelon inventory position 1 to its optimal value S_1 . We defined S_2 as the point where $G_2(S_1, y_2)$ is minimized. Obviously, one finds the same optimal base-stock level by the minimization of the echelon 2 costs (the first term) plus the third term. This is how Clark and Scarf proceeded, and they interpreted the third term as an *induced penalty cost function*.

2.3. Linear Inventory Holding and Penalty Costs

In this subsection, we assume that the echelon cost functions $c_n(\cdot)$, $n = 1, 2$, are given by (1)–(2), i.e., we consider the special, but common, cost structure consisting of linear inventory holding and penalty costs. We derive interesting, additional results. First, in §2.3.1, we derive an alternative formula in terms of expected shortfalls and backlogs for the average costs of a base-stock policy. That formula facilitates computational procedures, and we exploit that formula to get the partial derivative to the base-stock level of echelon 2. For the average costs in subsystem 1, we obtain also a partial derivative, and the combination of both partial derivatives leads to newsboy equations for the optimal base-stock levels; see §2.3.2.

2.3.1. Alternative Cost Formulas for Base-Stock Policies. Assume the echelon costs functions as given by (1)–(2) and consider a base-stock policy (y_1, y_2) , $y_1, y_2 \in \mathbb{R}$. The average costs $G_2(y_1, y_2)$ may be obtained by a single-cycle analysis; see Figure 3. The costs consist of the terms $C_{t_0+l_2, 2}$ and $C_{t_0+l_2+l_1, 1}$. The expected value of the costs $C_{t_0+l_2, 2}$ equals

$$\mathbb{E}C_{t_0+l_2, 2} = \mathbb{E}\{c_2(y_2 - D_{t_0, t_0+l_2})\} = \mathbb{E}\{h_2(y_2 - D_{t_0, t_0+l_2})\} = h_2(y_2 - (l_2 + 1)\mu).$$

Next, we study $\mathbb{E}C_{t_0+l_2+l_1, 1}$. The level z_1 denotes the actual level to which $IP_{t_0+l_2, 1}$ is increased. The difference with the desired level y_1 is called the *shortfall*, which can also be seen as a “backlog” at stockpoint 2 (it would be the backlog at stockpoint 2 if stage 1 would order such that $IP_{t_0+l_2, 1}$ is increased up to y_1 , without taking into account how much is available at stockpoint 2). We denote this shortfall by B_1 . This shortfall is equal to

$$\begin{aligned} B_1 &= y_1 - z_1 \\ &= y_1 - \min\{y_2 - D_{t_0, t_0+l_2-1}, y_1\} \\ &= y_1 + \max\{-y_2 + D_{t_0, t_0+l_2-1}, -y_1\} \\ &= \max\{0, y_1 - y_2 + D_{t_0, t_0+l_2-1}\} \\ &= (D_{t_0, t_0+l_2-1} - (y_2 - y_1))^+ \end{aligned} \tag{9}$$

(notice that by definition this shortfall is positive if $y_1 > y_2$). Now, define B_0 as the backlog at stockpoint 1 at the end of period $t_0 + l_2 + l_1$. Given that $IP_{t_0+l_2,1}$ is increased up to $z_1 = y_1 - B_1$, B_0 becomes equal to

$$\begin{aligned} B_0 &= (z_1 - D_{t_0+l_2, t_0+l_2+l_1})^- \\ &= (D_{t_0+l_2, t_0+l_2+l_1} - z_1)^+ \\ &= (D_{t_0+l_2, t_0+l_2+l_1} - (y_1 - B_1))^+ \\ &= (B_1 + D_{t_0+l_2, t_0+l_2+l_1} - y_1)^+. \end{aligned} \quad (10)$$

Then, for the costs attached to echelon 1 at the end of period $t_0 + l_2 + l_1$, we obtain

$$\begin{aligned} EC_{t_0+l_2+l_1,1} &= E\{c_1(z_1 - D_{t_0+l_2, t_0+l_2+l_1})\} \\ &= E\{h_1(z_1 - D_{t_0+l_2, t_0+l_2+l_1}) + (p + h_1 + h_2)(z_1 - D_{t_0+l_2, t_0+l_2+l_1})^-\} \\ &= h_1(y_1 - EB_1 - (l_1 + 1)\mu) + (p + h_1 + h_2)EB_0. \end{aligned}$$

As a result, we find the following theorem. (The formula in this theorem stems from van Houtum and Zijm [62], where an equivalent formula has been derived, but with EB_1 and EB_0 expressed in integral form.)

Theorem 2. *Let the echelon cost functions $c_n(\cdot)$ be given by (1)–(2). Then, the average costs of a base-stock policy (y_1, y_2) , with $y_1, y_2 \in \mathbb{R}$, are equal to*

$$G_2(y_1, y_2) = h_2(y_2 - (l_2 + 1)\mu) + h_1(y_1 - EB_1 - (l_1 + 1)\mu) + (p + h_1 + h_2)EB_0,$$

where the random variables B_1 and B_0 are given by (9)–(10).

The formula for the average costs of a base-stock policy (y_1, y_2) also shows what the average backlog and average stock levels are. The term EB_0 denotes the average backlog at the end of a period. The amount $y_1 - EB_1 - (l_1 + 1)\mu + EB_0$ is the average physical stock of echelon 1 (= stockpoint 1) at the end of a period; this is the amount for which a cost h_1 is paid per unit of product. The amount $y_2 - (l_2 + 1)\mu + EB_0$ is the average physical stock of echelon 2 at the end of a period; this is the amount for which a cost h_2 is paid per unit of product. Further, the average stock in the pipeline between stockpoint 2 and stockpoint 1 is $l_1\mu$ (the throughput of the pipeline is equal to the mean demand and each unit of product is l_1 periods in the pipeline). This implies that the average physical stock in stockpoint 2 at the end of a period is equal to

$$\begin{aligned} [y_2 - (l_2 + 1)\mu + EB_0] - [y_1 - EB_1 - (l_1 + 1)\mu + EB_0] - l_1\mu \\ = y_2 - y_1 - l_2\mu + EB_1 = E\{((y_2 - y_1) - D_{t_0, t_0+l_2-1})^+\}. \end{aligned} \quad (11)$$

For the average costs in subsystem 1, under a base-stock policy y_1 , $y_1 \in \mathbb{R}$, we find the following alternative expression (via (6)):

$$G_1(y_1) = h_1(y_1 - (l_1 + 1)\mu) + (p + h_1 + h_2)EB_0^{(1)}, \quad (12)$$

where the random variable $B_0^{(1)}$ represents the backlog in subsystem 1:

$$B_0^{(1)} = (D_{t_0+l_2, t_0+l_2+l_1} - y_1)^+. \quad (13)$$

Formula (12) shows that $G_1(y_1)$ is a newsboy function. Notice that $B_0^{(1)}$ is related to B_1 and B_0 in the following way: $B_0^{(1)} = (B_0 | B_1 = 0)$.

2.3.2. Newsboy Equations. We now determine the partial derivatives of $G_1(y_1)$ and $G_2(y_1, y_2)$. The derivative of $G_1(y_1)$ is denoted by $g_1(y_1)$. By (12),

$$g_1(y_1) = h_1 + (p + h_1 + h_2) \frac{\delta}{\delta y_1} \{EB_0^{(1)}\}.$$

It is easily seen that

$$\frac{\delta}{\delta y_1} \{EB_0^{(1)}\} = -P\{B_0^{(1)} > 0\}.$$

Substitution of this property into the previous equation shows that

$$g_1(y_1) = h_1 - (p + h_1 + h_2)P\{B_0^{(1)} > 0\}, \quad (14)$$

where $B_0^{(1)}$ is given by (13).

For the function $G_2(y_1, y_2)$, we are interested in the partial derivative with respect to the last component y_2 . Hence, we define

$$g_2(y_1, y_2) := \frac{\delta}{\delta y_2} \{G_2(y_1, y_2)\}, \quad y_1, y_2 \in \mathbb{R}.$$

We find that

$$\begin{aligned} g_2(y_1, y_2) &= h_2 - h_1 \frac{\delta}{\delta y_2} \{EB_1\} + (p + h_1 + h_2) \frac{\delta}{\delta y_2} \{EB_0\} \\ &= h_2 + h_1 P\{B_1 > 0\} - (p + h_1 + h_2)P\{B_1 > 0 \text{ and } B_0 > 0\}. \end{aligned} \quad (15)$$

Here, the second step follows from the following properties.

$$\begin{aligned} \frac{\delta}{\delta y_2} \{EB_1\} &= -P\{B_1 > 0\}, \\ \frac{\delta}{\delta y_2} \{EB_0\} &= -P\{B_1 > 0 \text{ and } B_0 > 0\}. \end{aligned}$$

These properties are easily verified. The result in (15) constitutes the basis for the following lemma.

Lemma 4. *Let the echelon cost functions $c_n(\cdot)$ be given by (1)–(2). Then*

$$g_2(y_1, y_2) = (h_1 + h_2) - (p + h_1 + h_2)P\{B_0 > 0\} - P\{B_1 = 0\}g_1(y_1), \quad y_1, y_2 \in \mathbb{R},$$

with B_1 and B_0 given by (9)–(10).

Proof. It holds that

$$\begin{aligned} P\{B_1 > 0\} &= 1 - P\{B_1 = 0\}, \\ P\{B_1 > 0 \text{ and } B_0 > 0\} &= P\{B_0 > 0\} - P\{B_1 = 0 \text{ and } B_0 > 0\} \\ &= P\{B_0 > 0\} - P\{B_0 > 0 | B_1 = 0\}P\{B_1 = 0\}. \end{aligned}$$

By substitution of these expressions into Equation (15), we obtain (use the property that $B_0^{(1)} = (B_0 | B_1 = 0)$, and (14)):

$$\begin{aligned} g_2(y_1, y_2) &= h_2 + h_1(1 - P\{B_1 = 0\}) \\ &\quad - (p + h_1 + h_2)(P\{B_0 > 0\} - P\{B_0 > 0 | B_1 = 0\}P\{B_1 = 0\}) \\ &= (h_1 + h_2) - (p + h_1 + h_2)P\{B_0 > 0\} \\ &\quad - P\{B_1 = 0\}[h_1 - (p + h_1 + h_2)P\{B_0 > 0 | B_1 = 0\}] \\ &= (h_1 + h_2) - (p + h_1 + h_2)P\{B_0 > 0\} \\ &\quad - P\{B_1 = 0\}\left[h_1 - (p + h_1 + h_2)P\{B_0^{(1)} > 0\}\right] \\ &= (h_1 + h_2) - (p + h_1 + h_2)P\{B_0 > 0\} - P\{B_1 = 0\}g_1(y_1). \quad \square \end{aligned}$$

Things bring us at the point to derive *newsboy equations* for the optimal base-stock levels S_1 and S_2 . Suppose that the demand distribution function F is continuous on $(0, \infty)$, and that there is no probability mass in 0, i.e., $F(0) = 0$. Then $g_1(y_1)$ is a continuous function, and as an optimal base-stock level is a minimal point of $G_1(y_1)$, S_1 will be a zero point of $g_1(y_1)$, i.e., $g_1(S_1) = 0$. This leads immediately to a newsboy equation for S_1 ; see Part (i) of Theorem 3. Next, by Lemma 4,

$$g_2(S_1, y_2) = (h_1 + h_2) - (p + h_1 + h_2)P\{B_0 > 0\}, \quad y_2 \in \mathbb{R},$$

where B_0 is given by (9)–(10) with y_1 replaced by S_1 . One can easily verify that this function is continuous as a function of y_2 . Because S_2 is a minimizing point of $G_2(S_1, y_2)$, it will be a zero point of $g_2(S_1, y_2)$, i.e., $g_2(S_1, S_2) = 0$. This leads immediately to a newsboy equation for S_2 ; see Part (ii) of the following theorem. The equation for S_2 is called a newsboy equation because it constitutes a generalization of the well-known newsboy equation for a single-stage system. Theorem 3 is stated to hold for a continuous demand distribution F , but, in fact, it holds if both $g_1(y_1)$ and $g_2(S_1, y_2)$ has a zero point.

Theorem 3 (cf. van Houtum and Zijm [62], Section 4). *Newsboy equations for the optimal base-stock levels—Let the echelon cost functions $c_n(\cdot)$ be given by (1)–(2), and let F be continuous on $(0, \infty)$ with $F(0) = 0$. Then*

(i) *The optimal base-stock level S_1 for echelon 1 is such that*

$$P\{B_0^{(1)} = 0\} = \frac{p + h_2}{p + h_1 + h_2},$$

with

$$B_0^{(1)} = (D_{t_0+l_2, t_0+l_2+l_1} - S_1)^+.$$

(ii) *Under a given optimal base-stock level S_1 for echelon 1, the optimal base-stock level S_2 for echelon 2 is such that*

$$P\{B_0 = 0\} = \frac{p}{p + h_1 + h_2},$$

with

$$B_1 = (D_{t_0, t_0+l_2-1} - (S_2 - S_1))^+, \\ B_0 = (B_1 + D_{t_0+l_2, t_0+l_2+l_1} - S_1)^+.$$

This theorem says that, when S_1 is determined, then it is pretended that stockpoint 2 can always deliver (i.e., the analysis is limited to subsystem 1) and the value for S_1 is chosen such that the no-stockout probability at stage 1 is equal to $(p + h_2)/(p + h_1 + h_2)$. Next, when S_2 is determined, then the full system is considered, the base-stock level for echelon 1 is fixed at S_1 , and the value for S_2 is chosen such that the no-stockout probability at the most downstream stage 1 is equal to $p/(p + h_1 + h_2)$. With this S_2 , the demand over a longer lead time has to be covered, but we are allowed to have a lower no-stockout probability in the full system than in subsystem 1.

Like for a single-stage system, our generalized newsboy equations show the effect of the ratios of the parameters for inventory holding and penalty costs on the optimal base-stock levels. In addition, they reveal how physical stock is positioned in the chain as a function of the way value is being built up in the chain. This is seen as follows. The echelon holding cost parameters h_1 and h_2 are, in general, proportional to the values added at stages 1 and 2, respectively. W.l.o.g., we may norm the total added value such that $h_1 + h_2 = 1$. In that case, h_n , $n = 1, 2$, is equal to the fraction of the added value in stage n over the total added value in the chain. Let us look at the values for S_1 and S_2 as a function of h_2 , i.e., the fraction of added value at stage 2. The larger h_2 , the closer $(p + h_2)/(p + h_1 + h_2) = (p + h_2)/(p + 1)$ comes to 1, and, thus, the larger S_1 . The point S_2 is such that we have a no-stockout probability $p/(p + h_1 + h_2) = p/(p + 1)$ for the full system. This fraction is independent of h_2 .

As S_1 is increasing as a function of h_2 , S_2 will be decreasing (a larger S_1 implies that a slightly smaller value for S_2 is sufficient to obtain that $P\{B_0 = 0\} = p/(p+1)$), and, thus, the difference $S_2 - S_1$ is decreasing as well. The average physical stock in stockpoint 2 at the end of a period equals $E\{((S_2 - S_1) - D_{t_0, t_0+t_2-1})^+\}$ (cf. (11)) and is also decreasing as a function of h_2 . The average physical stock in stockpoint 1 is likely to be increasing (because of the increased S_1 and only slightly decreased S_2 ; however, we have no proof for this property). In the extreme case that $h_2 = 1$, and thus $h_1 = 0$, there is no added value at all at stage 1. Then we may choose $S_1 = \infty$, in which case there is no safety stock held in stockpoint 2. This property holds in general when $h_1 = 0$.

Corollary 1. *There exists an optimal base-stock policy under which no safety stock is held in stockpoint 2 in case $h_1 = 0$.*

Proof. Suppose that $h_1 = 0$. Then, by Part (i) of Theorem 3, S_1 may be chosen equal to $S_1 = \infty$. This implies that, in each period, all goods arriving in stockpoint 2 are immediately forwarded to stockpoint 1, and, thus, there is never stock present in stockpoint 2 at the end of a period. $\square$

3. Multiechelon, Serial Systems, and Computational Procedures

The whole analysis of §2 is easily generalized to serial systems with $N \geq 2$ stages. For the generalization of the optimality of base-stock policies and the decomposition result, see the remarks at the end of §2.2 (just before Remark 1). In this section, we present the cost formulas and newsboy equations as obtained for the N -stage system under linear inventory holding and penalty costs; see §3.1. After that, in §3.2, we describe both exact and efficient approximate computational procedures for the optimal base-stock levels and optimal costs.

3.1. Analytical Results

We first describe our model for the multiechelon, serial system, and introduce additional notation. We make the same assumptions as in §2, however, we now have $N (\geq 2)$ stages, which are numbered from downstream to upstream as stages $1, 2, \dots, N$. Periods are numbered $0, 1, \dots$. Lead times are deterministic, and the lead time for stage n is denoted by l_n . The cumulative lead time for stages $i, n \leq i \leq N$, together is denoted by L_n ; $L_n = \sum_{i=n}^N l_i$, and, for notational convenience, $L_{N+1} := 0$. The cumulative demand over periods $t_1, \dots, t_2$, $0 \leq t_1 \leq t_2$, is denoted by D_{t_1, t_2} , F is the generic distribution function for one-period demand, and μ denotes the mean demand per period.

For the costs, we assume linear inventory holding and penalty costs. A cost of H_n , $n = 2, \dots, N$, is charged for each unit that is in stock in stockpoint n at the end of a period and for each unit in the pipeline from the n th to the $(n-1)$ th stockpoint. A cost of H_1 is charged for each unit that is in stock in stockpoint 1 at the end of a period, and a penalty $p > 0$ is charged per unit of backlog at stockpoint 1 at the end of a period. We assume that $H_1 \geq H_2 \geq \dots \geq H_N \geq 0$; for notational convenience, $H_{N+1} = 0$. Next, we define $h_n := H_n - H_{n+1}$, $n = 1, \dots, N$, as the additional inventory holding-cost parameters. Notice that $h_n \geq 0$ for all n . Under this cost structure and given levels x_n for the echelon stocks at the end of a period, the total inventory holding and backordering costs at the end of that period are equal to $\sum_{n=1}^N c_n(x_n)$, where $c_n(x_n)$ denotes the costs attached to echelon n (cf. (1)–(2) for $N = 2$):

$$c_1(x_1) = h_1 x_1 + (p + H_1) x_1^-,$$

$$c_n(x_n) = h_n x_n, \quad 2 \leq n \leq N.$$

Optimal base-stock levels follow from the minimization of average costs of a base-stock policy in subsystems. Subsystem n , $n = 1, \dots, N$, is defined as the system consisting of the stages $1, \dots, n$, and with infinite supply at stage $n+1$ (= external supplier of raw materials

in case $n = N$). As costs we have the echelon cost functions $c_i(\cdot)$ for the echelons $i = 1, \dots, n$. Notice that subsystem N is identical to the full system. A base-stock policy for subsystem n is denoted by $(y_1, \dots, y_n)$, with $y_i \in \mathbb{R}$ for all $i = 1, \dots, n$, and the corresponding average costs are denoted by $G_n(y_1, \dots, y_n)$. For this function, a similar expression may be derived as for the average costs of a two-echelon system in Theorem 2. We define $B_i^{(n)}$ as the shortfall as faced by stockpoint i , $1 \leq i \leq n$, and $B_0^{(n)}$ as the backlog at the end of an arbitrary period. For these variables, one easily derives similar recursive expressions as in (9)–(10). This leads directly to the following theorem.

Theorem 4 (cf. van Houtum and Zijm [62], van Houtum et al. [65]). *Let $1 \leq n \leq N$. For subsystem n , the average costs of a base-stock policy $(y_1, \dots, y_n)$, with $y_i \in \mathbb{R}$ for all $i = 1, \dots, n$, are equal to*

$$G_n(y_1, \dots, y_n) = \sum_{i=1}^n h_i (y_i - \mathbb{E}B_i^{(n)} - (l_i + 1)\mu) + (p + H_1)\mathbb{E}B_0^{(n)},$$

with

$$B_n^{(n)} = 0, \quad (16)$$

$$B_i^{(n)} = (B_{i+1}^{(n)} + D_{t_0+L_{i+2}, t_0+L_{i+1}-1} - (y_{i+1} - y_i))^+, \quad 1 \leq i \leq n-1, \quad (17)$$

$$B_0^{(n)} = (B_1^{(n)} + D_{t_0+L_2, t_0+L_1} - y_1)^+ \quad (18)$$

(the equation for $B_i^{(n)}$, $1 \leq i \leq n-1$, vanishes in case $n = 1$).

An optimal base-stock level S_1 for stage 1 is obtained as a minimizer of the convex function $G_1(y_1)$. Next, under a given S_1 , an optimal base-stock level S_2 for stage 2 is obtained as a minimizer of the function $G_2(S_1, y_2)$, which is known to be convex as a function of y_2 ; and so on. The optimal base-stock levels may also be obtained from partial derivatives. Define

$$g_n(y_1, \dots, y_n) := \frac{\delta}{\delta y_n} \{G_n(y_1, \dots, y_{n-1}, y_n)\}, \quad 1 \leq n \leq N, \quad y_i \in \mathbb{R} \text{ for all } i = 1, \dots, n.$$

Similar to Lemma 4, one can derive that

$$g_n(y_1, \dots, y_n) = \sum_{i=1}^n h_i - (p + H_1)\mathbb{P}\{B_0^{(n)} > 0\} - \sum_{i=1}^{n-1} \mathbb{P}\{B_i^{(n)} = 0\} g_i(y_1, \dots, y_i), \quad (19)$$

where the $B_i^{(n)}$ are given by (16)–(18) (in this formula the last sum vanishes in case $n = 1$). Under a continuous demand distribution F , $g_1(y_1)$ has a zero point, $g_2(S_1, y_2)$ has a point S_2 such that $g_2(S_1, S_2) = 0$, and so on. Then the last sum in (19) becomes equal to 0, and we get the following newsboy equations.

Theorem 5 (cf. van Houtum and Zijm [62], Theorem 5.1). *Newsboy equations for the optimal base-stock levels—Let F be continuous on $(0, \infty)$ with $F(0) = 0$. For $n = 1, 2, \dots, N$, under given optimal base-stock levels $S_1, \dots, S_{n-1}$ for the stages $1, \dots, n-1$, S_n is such that*

$$\mathbb{P}\{B_0^{(n)} = 0\} = \frac{p + H_{n+1}}{p + H_1},$$

where $B_0^{(n)}$ is given by the recursive formulas (16)–(18) with y_i replaced by S_i for all i .

3.2. Computational Procedures

In case of a continuous demand distribution F with $F(0) = 0$, an optimal base-stock policy $(S_1, \dots, S_N)$ and the corresponding average costs can be determined as follows. First, for $n = 1, \dots, N$, S_n may be determined by the newsboy equation in Theorem 5. In general,

this newsboy equation cannot be solved analytically. Computational procedures can be developed, however. Suppose one has a computational procedure to compute $P\{B_0^{(n)} = 0\}$ for a given arbitrary S_n . Then, an S_n that solves the newsboy equation is easily computed via bisection search. Once optimal base-stock levels have been determined for all stages, the optimal average costs $G_N(S_1, \dots, S_N)$ follow from Theorem 4. Here, one needs a method to obtain the expected values of the $B_i^{(N)}$, $0 \leq i \leq N$. For both the computation of the optimal base-stock levels and the corresponding optimal costs, it suffices if one is able to evaluate the shortfalls/backlogs $B_i^{(n)}$ as given by (16)–(18). That is what we focus on in the rest of this subsection.

The shortfalls/backlogs $B_i^{(n)}$ may be determined recursively after a sufficiently fine discretization of the one-period demand distribution F . This is a first method. However, this method will be computationally inefficient in many cases, in particular, as N grows large. Therefore alternative procedures are desired. In §3.2.1, we describe an efficient, exact procedure for mixed Erlang demand, i.e., for the case that the one-period demand is a mixture of Erlang distributions with the same scale parameter. Such mixtures are relevant because the class of these mixtures is dense in the class of all distributions on $[0, \infty)$ (cf. Schassberger [53]). In §3.2.2, we describe a procedure based on two-moment fits. This is a fast, approximate procedure that is known to be accurate.

If the demand distribution F is *not* continuous, then Theorem 5 does not apply anymore, but Equation (19) still does. An optimal base-stock level for stage n is then found at the first point S_n where $g_n(S_1, \dots, S_{n-1}, S_n) \geq 0$. Similar computations apply as described above, and the same methods may be used for the computation of the shortfalls/backlogs $B_i^{(n)}$. Via discretization, one still obtains an exact approach. The method of §3.2.2 is also applicable without further changes. The method of §3.2.1 may be applied after a (two-moment) fit of a mixed Erlang distribution on the one-period demand. That step is an approximate step, and for the rest the method is exact. A special case of noncontinuous demand is obtained in the case of a discrete product. Then, the demand distribution F is discrete as well, and base-stock and inventory levels may be limited to discrete values—in which case, Theorem 4 and Equation (19) are still valid. In this case, a direct recursive computation of the distributions of the shortfalls/backlogs $B_i^{(n)}$ may be efficient. For further details on this discrete product case, see §4.2.

3.2.1. Exact Procedure for Mixed Erlang Demands. The exact procedure as described here stems from van Houtum et al. [66], where for a generalized system with fixed replenishment intervals per stage, evaluation of shortfalls/backlogs of the same form as in (16)–(18) is needed. This procedure is closely related to the exact procedure described in van Houtum and Zijm [63], but the procedure as described here leads to simpler formulas and is easier to implement. The key idea behind the procedure is that we define a class of mixed Erlang distributions that is closed under the two basic operations in the expressions for the shortfalls/backlogs: Convolution and the so-called truncated shift.

Let us first define the class of mixed Erlang distributions that we use. We take $\lambda > 0$ as a given, and define a class of mixed Erlang random variables $\mathcal{C}_\lambda$. Let $X_{k,\lambda}$ be an Erlang distribution with $k \in \mathbb{N}_0$ phases and scale parameter λ . $X_{k,\lambda}$ may be interpreted as the sum of k independent, exponentially distributed random variables with parameter λ . Notice that we allow that $k = 0$. The distribution function of $X_{k,\lambda}$ is denoted by $E_{k,\lambda}$. For $k \in \mathbb{N}_0$,

$$E_{k,\lambda}(x) = 1 - \sum_{j=0}^{k-1} \frac{(\lambda x)^j}{j!} e^{-\lambda x}, \quad x \geq 0,$$

and $E_{k,\lambda}(x) = 0$ for all $x < 0$ (the sum $\sum_{j=0}^{k-1}$ is empty for $k = 0$). Let X be a pure mixture of the random variables $X_{k,\lambda}$, described by a discrete distribution $\{q_k\}_{k \in \mathbb{N}_0}$ on $\mathbb{N}_0$; i.e., $X = X_{k,\lambda}$ with probability q_k for all $k \in \mathbb{N}_0$. The distribution function of X is given by

$F_X(x) = \sum_{k=0}^{\infty} q_k E_{k,\lambda}(x)$, $x \in \mathbb{R}$. Finally, we define random variable Y as the sum of a deterministic variable $d \geq 0$ and a pure mixture X ; i.e., $Y = d + X$, and its distribution function is given by $F_Y(x) = \mathbf{P}\{d + X \leq x\} = F_X(x - d)$, $x \in \mathbb{R}$; this distribution is obtained by a shift of F_X to the right over a distance d . The class $\mathcal{C}_\lambda$ consists of all Y s that can be constructed in this way. Each $Y \in \mathcal{C}_\lambda$ is uniquely determined by a $d \geq 0$ and a discrete distribution $\{q_k\}_{k \in \mathbb{N}_0}$.

The first operation that we recognize in (16)–(18) is a convolution; i.e., $B_{i+1}^{(n)} + D_{t_0+L_{i+2}, t_0+L_{i+1}-1}$ is a convolution of the random variables $B_{i+1}^{(n)}$ and $D_{t_0+L_{i+2}, t_0+L_{i+1}-1}$, and $D_{t_0+L_{i+2}, t_0+L_{i+1}-1}$ itself is a convolution of l_{i+1} one-period demands; and similarly for $B_1^{(n)} + D_{t_0+L_2, t_0+L_1}$. Let $Y \in \mathcal{C}_\lambda$ with parameters d and $\{q_k\}_{k \in \mathbb{N}_0}$, $\tilde{Y} \in \mathcal{C}_\lambda$ with parameters $\tilde{d}$ and $\{\tilde{q}_k\}_{k \in \mathbb{N}_0}$, and $\hat{Y} := Y + \tilde{Y}$. Then, the sum $\hat{Y}$ may be written as $\hat{Y} = \hat{d} + \hat{X}$, where $\hat{d} = d + \tilde{d}$ and $\hat{X} = X + \tilde{X}$. Here, X is the pure mixture of Erlangs with discrete distribution $\{q_k\}_{k \in \mathbb{N}_0}$, and $\tilde{X}$ is the pure mixture given by $\{\tilde{q}_k\}_{k \in \mathbb{N}_0}$. It is easily seen that $\hat{X}$ is also a pure mixture of Erlangs; its distribution $\{\hat{q}_k\}_{k \in \mathbb{N}_0}$ is obtained via the convolution of $\{q_k\}_{k \in \mathbb{N}_0}$ and $\{\tilde{q}_k\}_{k \in \mathbb{N}_0}$: $\hat{q}_k = \sum_{j=0}^k q_{k-j} \tilde{q}_j$, $k \in \mathbb{N}_0$. Hence, $\hat{Y} \in \mathcal{C}_\lambda$. So, $\mathcal{C}_\lambda$ is closed under convolutions, and we have expressions to compute the parameters of an element that is obtained via a convolution.

The second operation that we recognize in (16)–(18) is a so-called *truncated shift*. Let Y be an arbitrary random variable (i.e., not necessarily an element of $\mathcal{C}_\lambda$), $a \in \mathbb{R}$, and $\hat{Y} := (Y - a)^+$. If $a \leq 0$, then $\hat{Y} = (-a) + Y$, and, thus, the distribution of $\hat{Y}$ is obtained by a shift to the right of the distribution of Y over a distance $-a$. If $a > 0$, then the distribution of $\hat{Y}$ is obtained by a shift to the left of the distribution of Y over a distance a , where the probability mass that would arrive in the negative range is absorbed in 0. Therefore, $\hat{Y}$ is said to be a truncated shift of Y . Suppose, now that $Y \in \mathcal{C}_\lambda$ with parameters d and $\{q_k\}_{k \in \mathbb{N}_0}$, let $a \in \mathbb{R}$, and define $\hat{Y} := (Y - a)^+$. Let X be the pure mixture of Erlangs given by $\{q_k\}_{k \in \mathbb{N}_0}$ (so, $Y = d + X$). We distinguish two cases: $a \leq d$ and $a > d$. If $a \leq d$, then $\hat{Y} = (Y - a)^+ = (d + X - a)^+ = (d - a) + X$, and, thus, $\hat{Y} \in \mathcal{C}_\lambda$ with parameters $d - a$ and $\{q_k\}_{k \in \mathbb{N}_0}$. Suppose now that $a > d$. Then

$$\hat{Y} = (X - (a - d))^+ = (X_{k,\lambda} - (a - d))^+ \quad \text{with probability } q_k, \quad k \in \mathbb{N}_0. \quad (20)$$

For each $k \in \mathbb{N}_0$, the k phases of $X_{k,\lambda}$ are equivalent to the first k interarrival times of a Poisson process with parameter λ , and $(X_{k,\lambda} - (a - d))^+$ depends on how many interarrival times have been completed at time instant $a - d$. With probability $[(\lambda(a - d))^j / j!] e^{-\lambda(a - d)}$, j phases of the Poisson process have been completed at time $a - d$, $j \in \mathbb{N}_0$. If $j < k$ phases have been completed, then there still are $k - j$ phases to go at time instant $a - d$, and, thus, then $(X_{k,\lambda} - (a - d))^+ = X_{k-j,\lambda}$. If $j \geq k$, then no phases are left, and $(X_{k,\lambda} - (a - d))^+ = 0$. Hence

$$(X_{k,\lambda} - (a - d))^+ = \begin{cases} X_{j,\lambda} & \text{with prob. } r_{k,j} = \frac{(\lambda(a - d))^{k-j}}{(k-j)!} e^{-\lambda(a - d)}, \quad j = 1, \dots, k; \\ 0 & \text{with prob. } r_{k,0} = 1 - \sum_{j=0}^{k-1} \frac{(\lambda(a - d))^j}{j!} e^{-\lambda(a - d)}. \end{cases} \quad (21)$$

Combining this result and (20) shows that

$$\hat{Y} = X_{j,\lambda} \quad \text{with probability } \hat{q}_j = \sum_{k=j}^{\infty} q_k r_{k,j}, \quad j \in \mathbb{N}_0.$$

As we see, $\hat{Y}$ is a pure mixture of Erlangs in this case. This implies that $\hat{Y} \in \mathcal{C}_\lambda$. So, $\mathcal{C}_\lambda$ is also closed under truncated shifts, and we have expressions to compute the parameters of an element that is obtained via a truncated shift.

Suppose now that the one-period demand D_{t_0, t_0} belongs to $\mathcal{C}_\lambda$ for some $\lambda > 0$; i.e., that $F = \sum_{k=0}^{\infty} q_k E_{k, \lambda}(x - d)$, $x \in \mathbb{R}$, where d is a nonnegative, real-valued constant and $\{q_k\}_{k \in \mathbb{N}_0}$ is a discrete distribution on $\mathbb{N}_0$. To obtain a continuous F with $F(0) = 0$, we require that $q_0 = 0$. Then each of the demand variables $D_{t_0+L_{i+2}, t_0+L_{i+1}-1}$ and $D_{t_0+L_2, t_0+L_1}$ in (17)–(18) belongs to $\mathcal{C}_\lambda$ because they are convolutions of one-period demands. The shortfall $B_n^{(n)}$ in (18) is equal to $X_{0, \lambda}$ (and, thus, belongs to $\mathcal{C}_\lambda$). Next, for each $i = n-1, n-2, \dots, 1$, the distribution of $B_i^{(n)}$ is obtained via a convolution, leading to the distribution of $B_{i+1}^{(n)} + D_{t_0+L_{i+2}, t_0+L_{i+1}-1}$, followed by a truncated shift. Finally, $B_0^{(n)}$ is obtained via a convolution, leading to the distribution of $B_1^{(n)} + D_{t_0+L_2, t_0+L_1}$, followed by a truncated shift. In addition to these computations, it is simple to obtain the no-stockout probability $P\{B_0^{(n)} = 0\}$ and/or expected values of the shortfalls/backlogs.

This completes the description of the exact computational procedure for the mixed Erlang demand case. Such a mixture is assumed to be given for this procedure. In practice, however, often only the first two moments of the one-period demand are given, and then a two-moment fit may be applied first: A so-called Erlang($k-1, k$) distribution can be fitted if the coefficient of variation of the demand is smaller than or equal to one, and a so-called Erlang($1, k$) distribution otherwise (these fits are further explained in §3.2.2). In principle, more moments may be fitted as desired, yielding a larger mixture.

The more general class of *phase-type distributions* is likewise closed under convolutions and truncated shifts. So, an exact procedure can also be derived for phase-type distributions, although computations become much more complicated.

Finally, it is relevant to note that the shortfalls/backlogs $B_i^{(n)}$ are equivalent to waiting times in a so-called *appointment system* (Vanden Bosch and Dietz [67], Wang [68]). Suppose you have a single server in which $n+1$ customers arrive. The customers are numbered $n, n-1, \dots, 1, 0$, and they arrive at predetermined arrival times $0, y_n - y_{n-1}, \dots, y_2 - y_1, y_1$. The service times for the customers $n, n-1, \dots, 2, 1$ are given by the random variables $D_{t_0+L_{n+1}, t_0+L_n-1}, D_{t_0+L_n, t_0+L_{n-1}-1}, \dots, D_{t_0+L_3, t_0+L_2-1}, D_{t_0+L_2, t_0+L_1}$. Then, $B_i^{(n)}$ is the waiting time of customer i , $0 \leq i \leq n$ (cf. van Houtum and Zijm [63]). In fact, the exact procedure of this section may also be applied for the evaluation of waiting times in an appointment system if all service times belong to $\mathcal{C}_\lambda$ for a given $\lambda > 0$. The shortfalls/backlogs $B_i^{(n)}$ are also equivalent to waiting times in a multistage serial production system with planned lead times. For those systems, even a similar structure for the optimal policy and a decomposition result for the optimal planned lead times is obtained; see Gong et al. [38].

3.2.2. Efficient, Approximate Procedure Based on Two-Moment Fits. If one is satisfied with accurate approximations, then one may use the simple approximate procedure based on two-moment fits as described and tested in van Houtum and Zijm [62].

A two-moment fit may be applied to any nonnegative random variable X as follows. Let its mean μ_X (> 0) and coefficient of variation c_X (> 0) be given. Then, a mixture of two Erlangs may be fitted on X such that this mixture has the same first two moments as X (i.e., also the mean and coefficient of variation of this mixture are equal to μ_X and c_X , respectively). Let this mixture be denoted by $\hat{X}$. Then, $\hat{X} = X_{k_1, \lambda_1}$ with probability q_1 and $\hat{X} = X_{k_2, \lambda_2}$ with probability $q_2 = 1 - q_1$.

The type of mixture that may be fitted on X depends on the value of c_X . We give three types of mixtures as described by Tijms [60]. If $c_X \leq 1$, then we may fit an Erlang($k-1, k$) distribution, in which case, $k_1 = k-1$ and $k_2 = k$ for some $k \geq 2$ and $\lambda_1 = \lambda_2 = \lambda$. The Erlang($k-1, k$) distribution is a mixture of two Erlang distributions with the same scale parameter. The $k \geq 2$ is chosen such that $1/k < c_X^2 \leq 1/(k-1)$. Next, q_1 and λ are taken equal to

$$q_1 = \frac{1}{1 + c_X^2} \left[k c_X^2 - \sqrt{k(1 + c_X^2) - k^2 c_X^2} \right], \quad \lambda = \frac{k - q_1}{\mu_X}.$$

If $c_X \geq 1$, then we may fit a hyperexponential or an Erlang($1, k$) distribution. Which of these two distributions is used may depend on further information that is available on X , e.g., on the shape of its probability density function (see also Tijms [60]). A hyperexponential distribution is a mixture of two exponential distributions, i.e., $k_1 = k_2 = 1$. In this case, multiple choices for $\lambda_1, \lambda_2, q_1$ are possible, and one choice that works is given by

$$\lambda_1 = \frac{2}{\mu_X} \left(1 + \sqrt{\frac{c_X^2 - 1/2}{c_X^2 + 1}} \right), \quad \lambda_2 = \frac{4}{\mu_X} - \lambda_1, \quad q_1 = \frac{\lambda_1(\lambda_2\mu_X - 1)}{\lambda_2 - \lambda_1}.$$

An Erlang($1, k$) distribution is a mixture of an exponential distribution and an Erlang distribution with the same scale parameter. Then $k_1 = 1$ and $\lambda_1 = \lambda_2 = \lambda$. The k_2 is set as the smallest $k_2 \geq 3$ for which $(k_2^2 + 4)/(4k_2) \geq c_X^2$. Next, q_1 and λ are taken equal to

$$q_1 = \frac{2k_2c_X^2 + k_2 - 2 - \sqrt{k_2^2 + 4 - 4k_2c_X^2}}{2(k_2 - 1)(1 + c_X^2)}, \quad \lambda = \frac{q_1 + k_2(1 - q_1)}{\mu_X}.$$

To approximate the shortfalls/backlogs $B_i^{(n)}$ in (16)–(18), we take the following steps. First, we determine the first two moments of $B_n^{(n)} + D_{t_0+L_{n+1}, t_0+L_n} = D_{t_0+L_{n+1}, t_0+L_n}$, and we fit a mixture of two Erlangs on these first two moments. Given this fit, $B_{n-1}^{(n)}$ is a truncated shift of $D_{t_0+L_{n+1}, t_0+L_n}$, and via the observations made in §3.2.1 (among others, Equation (21)), it is straightforward to obtain the first two moments of $B_{n-1}^{(n)}$. Next, the first two moments of $B_{n-1}^{(n)} + D_{t_0+L_n, t_0+L_{n-1}}$ can be determined, and a mixture of two Erlangs may be fitted on these first two moments. This process is continued until a mixed Erlang distribution is obtained for $B_1^{(n)} + D_{t_0+L_2, t_0+L_1}$. From that last fit, it is straightforward to determine $EB_0^{(n)}$ or $P\{B_0^{(n)} = 0\} = P\{B_1^{(n)} + D_{t_0+L_2, t_0+L_1} \leq y_1\}$. (In this procedure, in case the two-moment fit is applied to a nonnegative random variable X that consists of a deterministic part $d > 0$ and a nonnegative variable $\tilde{X}$, i.e., $X = d + \tilde{X}$, one may consider to take this deterministic part explicitly into account; i.e., one can apply the fit on $\tilde{X}$ instead of X .)

In van Houtum and Zijm [62], the optimal policy and optimal costs of a multiechelon, serial system have been computed by both the approximate method based on two-moment fits and an exact method that is equivalent to the method of §3.2.1. A test bed has been defined in which holding cost parameters, lead times, the standard deviation of one-period demand, and the number of stages were varied, and an Erlang($k-1, k$) distribution has been assumed for the one-period demand (so that the exact method is applicable). The approximate method has appeared to be very accurate. The approximate procedure had a relative accuracy of 1% for the optimal base-stock levels and a relative accuracy of 2% for the optimal costs.

In case a higher accuracy is desired, the approximate method may be further improved by applying fits on the first three or even more moments; for three-moment fits, see Osogami and Harchol-Balter [47]. In the discrete product case (see also §4.2), one can use two-moments fits of discrete distribution as developed by Adan et al. [1].

4. Exact Solutions for Serial and Assembly Systems

In this section, we describe several generalizations/extensions of the multiechelon, serial system for which the optimal solution is known. First, in §§4.1–4.4, we describe modeling variants that we can easily deal with: Continuous review (and time) instead of periodic review, a discrete instead of continuous product, discounted instead of average costs, and the case with a γ -service level constraint instead of backordering costs. After that, in §4.5, we discuss the reduction of general assembly systems to serial systems. Next, in §§4.6–4.7, we describe the main results for serial systems with two different forms of batching: A fixed batch size per stage and a fixed replenishment interval per stage. Finally, some other extensions are discussed in §4.8.

4.1. Continuous Review

In §§2 and 3, we have assumed periodic review, but there is (almost) a full equivalence between periodic-review and continuous-review multiechelon systems: see Chen [10], Chen and Zheng [12], and Gallego and Zipkin [35]. Here, we demonstrate that equivalence for the two-echelon, serial system of §2.

Suppose we have the same two-echelon system as in §2, but now with continuous time and continuous review, i.e., we consider a time interval $[0, \infty)$ and ordering decisions may be taken at any time instant $t \in [0, \infty)$. Demands are assumed to occur according to a compound Poisson process. Hence, the demand process is memoryless, which is similar to i.i.d. demands in the periodic-review case. The total demand in an interval $(t_1, t_2]$ is denoted by D_{t_1, t_2} . So, D_{t_1, t_2} denotes the demand over a time interval with length $t_2 - t_1$; this is slightly different from the periodic-review case, where D_{t_1, t_2} was used to denote total demand over the periods $t_1, \dots, t_2$ and, thus, corresponds to a length $t_2 - t_1 + 1$. The lead times l_1 and l_2 for the stockpoints 1 and 2 may be arbitrary, positive, real-valued numbers. Finally, the echelon cost functions $c_1(x_1)$ and $c_2(x_2)$ are now assumed to be cost *rate* functions.

For the continuous-review system, we define a cycle for each time instant $t_0 \in [0, \infty)$ in a similar way as for the periodic-review case. We get a similar picture as in Figure 3, but now, decision 2 is taken at time instant t_0 , and decision 1 is taken at time instant $t_0 + l_2$, where the level z_1 is limited from above by $z_2 - D_{t_0, t_0+l_2}$ (in the periodic-review case, z_1 was bounded from above by $z_2 - D_{t_0, t_0+l_2-1}$; the change in this expression is due to the change in the definition of demands D_{t_1, t_2}). Decision 2 directly affects the echelon 2 cost rate at time instant $t_0 + l_2$, and decision 1 directly affects the echelon 1 cost rate at time instant $t_0 + l_2 + l_1$. These costs are given by exactly the same formulas as in the periodic-review case, i.e., by (3) and (4), respectively (notice, however, that the interpretation of D_{t_0, t_0+l_2} and $D_{t_0+l_2, t_0+l_2+l_1}$ is slightly different now).

Next, for each $t_0 \in [0, \infty)$, we define the same relaxed single-cycle problem as in the periodic-review case; the only difference is that in the definition of problem $(RP(t_0))$ the demand variable D_{t_0, t_0+l_2-1} is replaced by D_{t_0, t_0+l_2} . This relaxed single-cycle problem is solved in the same way as before. Therefore, we again find that there is an optimal base-stock policy (S_1, S_2) for problem $(RP(t_0))$, and the optimal base-stock levels follow from the minimization of convex functions $G_1(y_1)$ and $G_2(S_1, y_2)$; these functions are defined by (6) and (7), with D_{t_0, t_0+l_2-1} replaced by D_{t_0, t_0+l_2} in (7). Subsequently, for the infinite-horizon problem, it is optimal to follow base-stock policy (S_1, S_2) at each time instant, and, thus, base-stock policy (S_1, S_2) is also optimal for that problem. Finally, under linear holding and penalty costs, we obtain the same formulas as in the periodic-review case, but with D_{t_0, t_0+l_2-1} replaced by D_{t_0, t_0+l_2} in Equation (9) for B_1 . Theorem 2 is still valid, and the newsboy equations of Theorem 3 hold as long as zero points exist for the functions $g_1(y_1)$ and $g_2(S_1, y_2)$. As the demand process is a compound Poisson process, the distribution functions for D_{t_0, t_0+l_2} and $D_{t_0+l_2, t_0+l_2+l_1}$ have a positive probability mass in zero and, thus, it is not guaranteed that zero points exist for $g_1(y_1)$ and $g_2(S_1, y_2)$. This last issue constitutes a minor difference between the continuous-review and the periodic-review case. For the rest, all results are essentially the same.

4.2. Discrete Product

In §§2 and 3, we assumed that ordered amounts and inventory levels are continuous variables, mainly because that smooths the analysis. This assumption is natural for a continuous product for which customers may demand any real-valued amount. Further, the assumption makes sense for a discrete product with a sufficiently high mean demand and customers that may demand any integer-valued amount. However, for a discrete product with a low mean demand, it is more logical to limit order sizes and inventory levels to integer values. The analysis and results for the two-echelon system then change as follows. All cost functions, such as $G_1(y_1)$ in (6) and $G_2(y_1, y_2)$ in (7), are limited to discrete domains $\mathbb{Z}$ and $\mathbb{Z}^2$,

respectively. All results in the Lemmas 1–4 and Theorems 1–2 are still valid, where now the discrete variant of the definition of a convex function has to be taken, and the optimal base-stock levels S_1 and S_2 are obtained by the minimization of one-dimensional functions on $\mathbb{Z}$: $S_1 = \arg \min_{y_1 \in \mathbb{Z}} G_1(y_1)$ and $S_2 = \arg \min_{y_2 \in \mathbb{Z}} G_2(S_1, y_2)$. The newsboy equations of Theorem 3, which hold under linear inventory holding and penalty costs, become newsboy inequalities in this case (cf. Doğru et al. [23]). An optimal base-stock level S_1 for echelon 1 is obtained at the lowest $y_1 \in \mathbb{Z}$ for which

$$P\{B_0^{(1)} = 0\} \geq \frac{p + h_2}{p + h_1 + h_2},$$

with $B_0^{(1)} = (D_{t_0+l_2, t_0+l_2+l_1} - y_1)^+$. Define $\varepsilon(S_1)$ as the difference between the left and right side of this inequality at the point S_1 . Then, $\varepsilon(S_1) \geq 0$ and, in general, $\varepsilon(S_1)$ will be small. Next, an optimal base-stock level S_2 for echelon 2 is obtained at the lowest $y_2 \in \mathbb{Z}$ for which

$$P\{B_0 = 0\} \geq \frac{p}{p + h_1 + h_2} + P\{B_1 = 0\}\varepsilon(S_1), \quad (22)$$

with

$$\begin{aligned} B_1 &= (D_{t_0, t_0+l_2-1} - (y_2 - S_1))^+, \\ B_0 &= (B_1 + D_{t_0+l_2, t_0+l_2+l_1} - S_1)^+. \end{aligned}$$

The second term on the right side of (22) is nonnegative, and, thus, under the optimal base-stock policy (S_1, S_2) , the no-stockout probability in the full system is at least equal to $p/(p + h_1 + h_2)$. (For the generalization of these newsboy inequalities to serial systems with two or more stages Doğru et al. [23].)

4.3. Discounted Costs

Instead of minimizing average costs, one may be interested in minimizing discounted costs with a discount factor β , $0 < \beta < 1$. In practice, using discounted costs becomes relevant if the total lead time of a multiechelon system is long. The analysis hardly changes under discounted costs, because we can show on the basis of the two-echelon system of §2. Cycles are defined in precisely the same way as under average costs. For the cycle costs C_{t_0} , however, the echelon 2 costs $C_{t_0+l_2, 2}$ have to be multiplied by a factor β^{l_2} , and the echelon 1 costs $C_{t_0+l_2+l_1, 1}$ by a factor $\beta^{l_2+l_1}$ as they are charged l_2 and $l_2 + l_1$ periods after period t_0 . Equivalently, in the single-cycle analysis, one may replace the echelon cost functions $c_1(x_1)$ and $c_2(x_2)$ by the modified functions $\tilde{c}_1(x_1) = \beta^{l_2+l_1}c_1(x_1)$ and $\tilde{c}_2(x_2) = \beta^{l_2}c_2(x_2)$. Under the presence of the discount factor, all convexity properties remain valid, and, thus, all main results hold again. Base-stock policies are optimal again. Under linear inventory holding and penalty costs, again, newsboy equations are obtained. For the optimal base-stock level S_1 , the newsboy equation in Theorem 3(i) is still valid. For the optimal base-stock level S_2 , we obtain the same newsboy equation as in Theorem 3(ii), but with the newsboy fractile $p/(p + h_1 + h_2)$ replaced by $(p - h_2(1 - \beta^{l_1})/\beta^{l_1})/(p + h_1 + h_2)$. Hence, the presence of the discount factor β has no effect on S_1 , and it has a decreasing effect on S_2 (this is due to the decreased importance of echelon 1 costs $C_{t_0+l_2+l_1, 1}$ relative to echelon 2 costs $C_{t_0+l_2, 2}$ in a cycle).

4.4. γ -Service-Level Constraint

As stated before, when analyzing multiechelon systems, often linear inventory holding and penalty costs are assumed for the costs. The penalty costs are the counterpart for inventory holding costs, and optimal policies find a balance between these two types of costs. As an alternative for the penalty costs, one may assume a target service level, and then the

objective is to minimize the inventory holding costs subject to a service-level constraint. Both types of models are related because models with penalty costs are Lagrange relaxations of models with service-level constraints; see van Houtum and Zijm [64] for an exposition of this relationship. The penalty costs that we have assumed in §§2.3 and 3.1 are of the so-called γ -type, and, therefore, the results of these sections can be extended to models with a so-called γ -service-level constraint. This is described below.

Consider the multiechelon serial system of §3.1. We still assume linear inventory holding costs, but we assume a γ -service-level constraint (which is equivalent to an average backlog constraint) instead of the linear penalty costs. The γ -service level is also known as the *modified fill rate*, and is closely related to the regular fill rate ($= \beta$ -service level). For high service levels (more precisely, as long as demand is very rarely backordered for more than one period), both measures are virtually identical. Let γ_0 be the target γ -service level. We make the additional assumption that the demand distribution F has a *connected support*, i.e., F is strictly increasing from 0 to 1 on an interval $[a, b]$, with $0 \leq a < b$ (b is allowed to be ∞). Under a base-stock policy $(y_1, \dots, y_N)$, the average backlog at the end of a period equals $EB_0^{(N)}$ (see Theorem 4), and the γ -service level is equal to

$$\gamma(y_1, \dots, y_N) = 1 - \frac{EB_0^{(N)}}{\mu};$$


$B_0^{(N)}$ is given by (16)–(18). An optimal policy for the γ -service-level model is obtained as follows. First, if the optimal policy $(S_1, \dots, S_N)$ for the penalty cost model (with penalty cost parameter p) has a γ -service-level $\gamma(S_1, \dots, S_N) = \gamma(p) = \gamma_0$, then $(S_1, \dots, S_N)$ is optimal for the service-level problem with target service-level γ_0 (cf. Everett [27]; see also Porteus [49] (Appendix B) and van Houtum and Zijm [64]). Second, the function $\gamma(p)$ is nondecreasing in p (cf. Everett [27], van Houtum and Zijm [64]). Third, under the assumption that F has connected support, one can show that the optimal base-stock levels $S_1, \dots, S_N$ are continuous in p ; thus, $\gamma(p)$ is also continuous in p . Moreover, $\gamma(p) \uparrow 1$ as $p \rightarrow \infty$. Therefore, the service-level problem with target $\gamma_0 < 1$ may be solved by repeatedly solving the penalty cost problem, tuning the penalty cost p until the γ -service-level $\gamma(p)$ of the optimal policy equals γ_0 . This solves the service-level problem and implies that *the class of base-stock policies is also optimal for the service-level problem with a γ -service-level constraint*.

For a multiechelon model with a target β -service level ($=$ fill rate) or a target α -service level ($=$ no-stockout probability), the relaxed model is a model with β -type or α -type penalty costs, respectively. Then, the resulting echelon cost functions $c_n(x_n)$ are not convex anymore, and the approach of §2.2 does not work anymore to prove the optimality of base-stock policies. In fact, it is likely that the structure of optimal policies is more complicated for these types of service-level constraints. Nevertheless, it still may make sense to take the class of base-stock policies as a given and to optimize within this class, e.g., Boyaci and Gallego [6].

4.5. Assembly Systems

All results and insights presented so far are for serial multiechelon systems. They also apply to multiechelon systems with a pure assembly/convergent structure, in which each stage has one or more predecessors and uses input materials of all predecessors to assemble/produce one output product. This is due to Rosling [51], who showed that the analysis of pure assembly systems is equivalent to the analysis of a serial system (independently, in Langenhoff and Zijm [45], this result has been obtained for a two-echelon assembly system). We show this equivalence for a system in which a final product is obtained by the assembly of two components that are made from raw material; see Figure 5. The components are numbered 1 and 2 and the lead time for component $n = 1, 2$ is $l_n \in \mathbb{N}$. The final product has index 0 and a lead time $l_0 \in \mathbb{N}_0$. W.l.o.g., we assume that one unit of the final product is made from one unit of Component 1 and one unit of Component 2. For the demand process, we have the same assumptions and notation as for the serial system in §2. For the costs,

FIGURE 5. An assembly system with two components.



we assume convex echelon cost functions $c_n(x_n)$, $n = 0, 1, 2$; in addition, we assume that $c_1(x_1)$ is nondecreasing.

If both components have equal lead times, i.e., if $l_1 = l_2$, then the two components may be replaced by one new virtual component of which each unit consists of one unit of Component 1 and one unit of Component 2, and the assembly system reduces to a two-echelon serial system.

From now on, we assume that the component lead times are different; w.l.o.g., assume that $l_2 > l_1$. For the resulting assembly system, an optimal policy may be derived along the same lines as in §2.2 for the two-echelon, serial system. At the beginning of each period $t_0 \in \mathbb{N}_0$, a single cycle starts and consists of the following three connected decisions:

- **Decision 2:** This decision concerns the order placed for Component 2 at the beginning of period t_0 , by which echelon inventory position 2 is increased up to z_2 . This decision leads to echelon 2 costs $c_2(z_2 - D_{t_0, t_0+l_2})$ at the end of period $t_0 + l_2$ and the resulting echelon stock 2 at the beginning of period $t_0 + l_2$ is $z_2 - D_{t_0, t_0+l_2-1}$.
- **Decision 1:** This decision concerns the order placed for component 1 at the beginning of period $t_0 + l_2 - l_1$, by which echelon inventory position 1 is increased up to z_1 . This decision leads to echelon 1 costs $c_1(z_1 - D_{t_0+l_2-l_1, t_0+l_2})$ at the end of period $t_0 + l_2$, and the resulting echelon stock 1 at the beginning of period $t_0 + l_2$ is $z_1 - D_{t_0+l_2-l_1, t_0+l_2-1}$.
- **Decision 0:** This decision concerns the order placed for the final product at the beginning of period $t_0 + l_2$, by which echelon inventory position 0 is increased up to z_0 . When this decision is taken, we are limited from above by the echelon stocks of the two components at that moment, i.e., $z_0 \leq \min\{z_2 - D_{t_0, t_0+l_2-1}, z_1 - D_{t_0+l_2-l_1, t_0+l_2-1}\}$. Decision 0 leads to echelon 0 costs $c_0(z_0 - D_{t_0+l_2, t_0+l_2+l_0})$ at the end of period $t_0 + l_2 + l_0$.

We may now introduce an additional constraint based on the observation that it is never useful to order more for Component 1 than what is available in the parallel pipeline for Component 2. More precisely, the level z_1 to which echelon inventory position 1 is increased by decision 1 may be limited by the echelon stock 2 at that moment plus the amounts that will arrive at stockpoint 2 at the beginning of the periods $t_0 + l_2 - l_1, \dots, t_0 + l_2$, i.e., by $z_2 - D_{t_0, t_0+l_2-l_1-1}$. If we take z_1 equal to $z_2 - D_{t_0, t_0+l_2-l_1-1}$, then echelon stock 2 and echelon stock 1 are both equal to $z_2 - D_{t_0, t_0+l_2-1}$ at the beginning of period $t_0 + l_2$. If we would take z_1 larger than $z_2 - D_{t_0, t_0+l_2-l_1-1}$, we would know beforehand that at the beginning of period $t_0 + l_2$, a portion of the arriving order at stockpoint 1 has to wait one or more periods for companion units in stockpoint 2. That portion would only lead to a larger echelon stock 1, and, thus, to equal or increased costs because $c_1(x_1)$ is nondecreasing. Hence, for decision 1, we introduce the additional constraint $z_1 \leq z_2 - D_{t_0, t_0+l_2-l_1-1}$. As a result, the constraint for decision 0 simplifies to $z_0 \leq z_1 - D_{t_0+l_2-l_1, t_0+l_2-1}$, and the decision structure for our assembly system becomes identical to the decision structure for a serial

system with three stages and lead times $l_0, l_1, l_2 - l_1$. Therefore, the optimal policy for our assembly system can be derived along the same lines as for that equivalent serial system (the cost structure in our assembly system is slightly different from the standard cost structure in a three-stage serial system, but it is still such that we have convex direct expected costs in the relaxed single-cycle problem). We again find that base-stock policies are optimal, and the optimal base-stock levels follow from the minimization of convex cost functions. In the special case of linear inventory holding and penalty costs, we obtain newsboy equations that are identical to the newsboy equations for a three-stage serial system with lead times $l_0, l_1, l_2 - l_1$, additional holding cost parameters h_0, h_1, h_2 , and penalty cost parameter p .

The description above shows that the reduction of an assembly system to a serial system follows from a basic observation. Hence, this reduction is easily applied to many extensions of the Clark-Scarf system, among which the extensions in §§4.6–4.8.

4.6. Fixed Batch Size per Stage

In many supply chains, there may be setup times and costs involved each time that an order is placed. Setup costs may be modeled directly by fixed ordering costs. This leads to a serial system with a fixed ordering cost per stage, as studied by Clark and Scarf [15]. These fixed ordering costs cannot be captured by convex cost functions $c_n(x_n)$, and, thus, the analysis of §2 does not work anymore. In fact, an optimal solution seems to be complicated in this case; an exception is the case with a fixed ordering cost at the most upstream stage only (see also §4.8).

An alternative way to limit the number of orders per stage is by the introduction of a fixed batch size per stage, a fixed replenishment interval per stage, or a combination of both. These limitations may be determined at the first decision level of the hierarchical approach as discussed at the beginning of §1. In this subsection, we discuss the case with a fixed batch size per stage.

Consider the multiechelon, serial system as described in §3, and assume that a fixed batch size Q_n applies for stage n , $n = 1, \dots, N$. This means that stage n is allowed to order at the beginning of each period, but the size of each order has to be an integer multiple of Q_n . There are no fixed ordering costs. The fixed batch size Q_{n+1} for stage $n + 1$ is assumed to be an integer multiple of the fixed batch size for stage n , $n = 1, \dots, N - 1$. This is known as the integer-ratio constraint. This constraint facilitates the analysis and reflects that the further upstream we are in a supply chain, the higher the setup times and costs tend to be, and, thus, larger batch sizes are desired. We also assume that at time 0, the physical stock in stage n is an integer multiple of Q_{n-1} , $n = 2, \dots, N$. For this system, Chen [10] (see also Chen [9]) derived the following optimal policy structure, via the approach that we used in §2.2. Each stage n , $n = 1, \dots, N$, has to control its echelon inventory position by an (s, Q) -policy with fixed batch size Q_n and a reorder level s_n that follows from the minimization of a one-dimensional convex function. This policy is called a multiechelon (s, Q) -policy, and is a generalized form of a base-stock policy. Under a base-stock policy, each stage aims to bring its echelon inventory position back to the same point at the beginning of each period, while each stage aims to bring its echelon inventory position back to the interval $(s, s + Q]$ under a multiechelon (s, Q) -policy. For the case with linear inventory holding and penalty costs, Doğru et al. [24] generalized the cost formulas of Theorem 4 and the newsboy equations of Theorem 5, which now hold for the reorder levels s_n . In fact, for each $n = 1, \dots, N$, the newsboy equation itself as given in Theorem 5 does not change; there are only a few changes in the recursive formulas (16)–(18) for the backlogs $B_0^{(n)}$.

4.7. Fixed Replenishment Interval per Stage

An alternative way to limit the number of orders per stage is by fixed replenishment intervals. Fixed replenishment intervals facilitate freight consolidations and logistics/production

scheduling and are, therefore, often observed in practice (cf. Graves [39]). In this subsection, we summarize the main results for such systems.

Consider the multiechelon, serial system as described in §3, and assume that a fixed replenishment interval T_n is specified for stage n , $n = 1, \dots, N$. In this case, orders may have any size, but stage n is only allowed to order at the beginning of every T_n periods. The replenishment interval T_{n+1} of stage $n+1$ is assumed to be an integer multiple of the replenishment interval T_n of stage n , $n = 1, \dots, N-1$ (integer-ratio constraint). In addition, we assume that the replenishment epochs are timed such that arriving materials at one stockpoint can be forwarded immediately to the next stockpoint if desired (synchronization constraint). This system has been analyzed in van Houtum et al. [66], along essentially the same lines as in §2.2. The main difference is constituted by the definition of a cycle. Consider, for example, a system with $N = 2$ stages. Then, a cycle is defined for each period t_0 in which stage 2 is allowed to order. An order by stage 2 in such a period t_0 directly affects the echelon 2 costs in the periods $t_0 + l_2, t_0 + l_2 + 1, \dots, t_0 + l_2 + T_2 - 1$, and it limits the levels to which echelon inventory position 1 may be increased in the periods $t_0 + l_2, t_0 + l_2 + T_1, \dots, t_0 + l_2 + kT_1$, where $k = T_2/T_1$. Further, each order by stage 1 in one of these periods $t' = t_0 + l_2, t_0 + l_2 + T_1, \dots, t_0 + l_2 + kT_1$ has a direct effect on the echelon 1 costs in the periods $t' + l_1, t' + l_1 + 1, \dots, t' + l_1 + T_1 - 1$. A cycle now consists of $k + 1$ decisions, one decision for stage 2 and k decisions for stage 1, and the cycle costs consist of the echelon 2 costs in the periods $t_0 + l_2, t_0 + l_2 + 1, \dots, t_0 + l_2 + T_2 - 1$ and the echelon 1 costs in the periods $t_0 + l_2 + l_1, t_0 + l_2 + l_1 + 1, \dots, t_0 + l_2 + l_1 + T_2 - 1$. Based on this definition of a cycle, all main results of the Clark-Scarf model have been generalized in van Houtum et al. [66]. In this case, we find a multiechelon (T, S) -policy as optimal policy; i.e., at the beginning of every T_n periods, stage n orders according to a base-stock policy with level S_n . For the newsboy equations, we now have to look at the average no-stockout probability over multiple periods, but we keep the same newsboy fractiles.

It is also possible to use both fixed batch sizes and fixed replenishment intervals. Serial systems with that combination have been analyzed by Chao and Zhou [8]. They combined the insights of Chen [10] and van Houtum et al. [66], and showed that the structure of the optimal policy is obtained by the combination of multiechelon (s, Q) - and (T, S) -policies.

For a cost comparison between serial systems with fixed batch sizes and serial systems with fixed replenishment intervals, we refer to Feng and Rao [32]. For a system with linear inventory holding costs, linear penalty costs, and fixed ordering costs, they compared the optimal multiechelon (T, S) -policy to the optimal multiechelon (s, Q) -policy. Multiechelon (s, Q) -policies lead to lower costs in general, but the differences in costs are relatively small. Hence, multiechelon (T, S) -policies are easily more attractive in situations in which freight consolidations and other coordination issues are important.

4.8. Other Extensions

There are a few more multiechelon, serial systems for which the structure of the optimal policy has been derived. Chen and Song [11] derived the optimal policy for a serial system with Markov-modulated demand, and Gallego and Özer [33] for a serial system with a specific form of advance demand information. In both cases, generalized forms of base-stock policies are optimal. Generalized base-stock policies may also be optimal for serial systems with an additional feature for the most upstream case.

Consider, for example, the two-echelon, serial system of §2 with a fixed capacity C for the upstream stage. Due to this fixed capacity, the upstream stage is never allowed to order more than C units in any period. For this system, a (modified) base-stock policy with parameters (S_1, S_2) is optimal (cf. Zijm and van Houtum [69]). This result is obtained as follows. Define cycles, cycle costs, and the relaxed single-cycle problem in a similar way as in §2.2. For the downstream stage of the relaxed single-cycle problem, one can show that a base-stock policy with a level S_1 is optimal. Next, one can conclude that it is optimal for stage 1 to follow

this base-stock policy in all periods. What remains is an infinite-horizon problem for stage 2 with a convex cost function $G_2(S_1, y_2)$ that denotes the costs attached to a period t_0 if the inventory position of echelon 2 in that period is increased to level y_2 . This problem fits in the single-stage, capacitated inventory model as analyzed by Federgruen and Zipkin [30, 31]. Hence, for echelon 2, a so-called modified base-stock policy is optimal, i.e., at the beginning of each period, echelon 2 has to increase its echelon inventory position to a level S_2 if the fixed capacity allows this, and, otherwise, the echelon inventory position is increased as far as possible by an order of size C . The difference between S_2 and the actual level to which echelon inventory position 2 is increased is called a shortfall and its steady-state distribution is identical to the steady-state waiting time in an equivalent $D|G|1$ queue (cf. Tayur [58], Zijm and van Houtum [69]). By exploiting this observation, the results in Theorems 2 and 3 are easily generalized. For a multiechelon, serial system with a fixed capacity constraint at the most upstream stage, the optimality of base-stock policies is obtained in the same way.

Similarly, the optimal policy for a multiechelon, serial system with a fixed ordering cost for the upstream stage is obtained. In this case, all stages except the most upstream one has to follow a base-stock policy, and for the most upstream stage, it is optimal to follow an (s, S) -policy (cf. Clark and Scarf [15]). The policy for the most upstream stage follows from the fact that an (s, S) -policy is optimal for a single-stage inventory system with fixed ordering costs (cf. Scarf [52]).

Finally, Shang and Song [54] (see also Boyaci et al. [7]) obtained interesting results for the multiechelon, serial system by the definition of lower- and upper-bound subsystems for the subsystems $1, \dots, N$ for the case with linear inventory holding and penalty costs. The upper-bound subsystems have a newsboy solution and have been shown to lead to lower bounds S_n^l for the optimal base-stock levels S_n . The lower-bound subsystems also have a newsboy solution and lead to upper bounds S_n^u for the optimal base-stock levels S_n . The weighted averages $(S_n^l + S_n^u)/2$ have appeared to be rather accurate approximations for the optimal base-stock levels S_n . An advantage of these approximations is that they are easy to compute. An alternative newsboy-type approximation has been developed by Gallego and Özer [34]. In Shang and Song [55], the bounds of Shang and Song [54] have been generalized to serial systems with a fixed batch size per stage; for a connection between these bounds and the newsboy equations for the optimal base-stock/reorder levels, see Dođru et al. [24].

5. Distribution and General Systems


The theory presented in the previous sections shows that generalized base-stock policies are optimal for a variety of multiechelon systems with a pure serial or assembly/convergent structure, that optimal base-stock levels follow from the minimization of convex, one-dimensional functions, and that optimal base-stock levels satisfy newsboy equations for many systems. In §§5.1 and 5.2, we discuss systems with a pure distribution/divergent structure. Nice results may be derived for such systems under the so-called *balance assumption*. Without the balance assumption, however, the structure of the optimal policy may be complicated, and alternative approaches have to be taken in that case. In §5.3, we discuss general systems, with a mixed convergent-divergent structure. That are the systems that often occur in practice. We briefly discuss the approaches that have been developed for such systems.

5.1. A Basic Distribution System

In this subsection, we first extend the analysis of the two-echelon, serial system to a very basic distribution system. While doing that, we will introduce the balance assumption. As we shall see, the balance assumption, or, better, imbalance between inventories of different local stockpoints, is the key problem in the analysis of distribution systems.

Consider the distribution/divergent system depicted in Figure 6. In this system, there is one central stockpoint supplied by an external supplier, and two successive stockpoints

FIGURE 6. A two-echelon distribution system with two local warehouses.



supplied by this central stockpoint. Such a system may occur in a production environment, in which an intermediate product is used in two different final products. Alternatively, we obtain such a structure in a distribution network in which a product is kept on stock in a central warehouse and two different local warehouses. From now on, we use the terminology that is common for the latter environment.

For our distribution system, we make similar assumptions for the two-echelon, serial system of §2. The local warehouses are numbered 1 and 2, and we also denote them as stockpoints 1 and 2. The central warehouse is denoted as stockpoint 0. We have periods numbered $0, 1, \dots$. The central warehouse has a deterministic lead time $l_0 \in \mathbb{N}$, and local warehouse n has a deterministic lead time $l_n \in \mathbb{N}_0$, $n = 1, 2$. Demands at local warehouse $n = 1, 2$ in different periods are independent and identically distributed on $[0, \infty)$, and the demands at one local warehouse are independent of the demands at the other local warehouse. The cumulative demand at local warehouse n over periods $t_1, \dots, t_2$, $0 \leq t_1 \leq t_2$, is denoted by $D_{t_1, t_2}^{(n)}$, and the total demand at both warehouses together over those periods is denoted by $D_{t_1, t_2} = D_{t_1, t_2}^{(1)} + D_{t_1, t_2}^{(2)}$.

The costs are described by convex echelon cost functions $c_n(x_n)$. A special cost structure is constituted by linear inventory holding and penalty costs. Under that structure, a cost h_0 (≥ 0) is charged for each unit on stock in the central warehouse at the end of a period and for each unit in the pipelines from the central warehouse to the local warehouses. A cost $h_0 + h_n$ ($h_n \geq 0$) is charged for each unit on stock in local warehouse n at the end of a period, and a penalty cost p_n is charged per unit of backordered demand at local warehouse n at the end of a period, $n = 1, 2$. Let x_n be echelon stock n at the end of a period. Then, the total inventory holding and penalty costs at the end of a period can be shown to be equal to $\sum_{n=0}^2 c_n(x_n)$ with


$$c_0(x_0) = h_0 x_0,$$

$$c_n(x_n) = h_n x_n + (p_n + h_n + h_0) x_n^-, \quad n = 1, 2.$$

The objective is to minimize the average costs over the infinite horizon. We denote this problem as problem (P).

For the analysis, we follow the same steps as in §2.2. We start with the definition of cycles and cycle costs. Let $IL_{t,n}$ and $IP_{t,n}$ denote echelon stock n and echelon inventory position n at the beginning of period t (just before demand occurs), and let $C_{t,n}$ be the costs attached to echelon n at the end of period t . A cycle starts with an order placed by the central warehouse at the beginning of a period $t_0 \in \mathbb{N}_0$. This decision is called decision 0. By this decision, $IP_{t_0,0}$ becomes equal to some level z_0 . First of all, this decision determines

FIGURE 7. The consequences of the decisions 0 and 1-2.



the echelon 0 costs at the end of period $t_0 + l_0$:

$$E\{C_{t_0+l_0,0} | IP_{t_0,0} = z_0\} = E\{c_0(z_0 - D_{t_0,t_0+l_0})\}.$$

Second, by this decision, echelon stock 0 at the beginning of period $t_0 + l_0$ becomes equal to $IL_{t_0+l_0,0} = z_0 - D_{t_0,t_0+l_0-1}$, and this directly limits the levels to which one can increase the echelon inventory positions $IP_{t_0+l_0,n}$, $n = 1, 2$, of the local warehouses at that moment. The latter decision is denoted as decision 1-2. Suppose that by this decision, $IP_{t_0+l_0,n}$ becomes equal to z_n , $n = 1, 2$. The sum $z_1 + z_2$ of these levels is limited from above by $z_0 - D_{t_0,t_0+l_0-1}$. Decision 1-2 directly affects the echelon n costs at the end of period $t_0 + l_0 + l_n$.

$$E\{C_{t_0+l_0+l_n,n} | IP_{t_0+l_0,n} = z_n\} = E\{c_n(z_n - D_{t_0+l_0,t_0+l_0+l_n}^{(n)})\}, \quad n = 1, 2.$$

The cycle costs C_{t_0} are equal to $C_{t_0+l_0,0} + \sum_{n=1}^2 C_{t_0+l_0+l_n,n}$. When the decisions 0 and 1-2 are taken, there is also bounding from below, but this is ignored for the moment. For a visualization of these decisions and the cycle costs; see Figure 7.

The second step of the analysis consists of the definition of the relaxed single-cycle problem. We obtain the following relaxed problem.

$$\begin{aligned} (\text{RP}(t_0)) : \quad & \text{Min} \quad EC_{t_0} = EC_{t_0+l_0,0} + \sum_{n=1}^2 EC_{t_0+l_0+l_n,n} \\ & \text{s.t.} \quad EC_{t_0+l_0,0} = E\{c_0(z_0 - D_{t_0,t_0+l_0})\}, \\ & \quad EC_{t_0+l_0+l_n,n} = E\{c_n(z_n - D_{t_0+l_0,t_0+l_0+l_n}^{(n)})\}, \quad n = 1, 2, \\ & \quad z_1 + z_2 \leq IL_{t_0+l_0,0}, \\ & \quad IL_{t_0+l_0,0} = z_0 - D_{t_0,t_0+l_0-1}. \end{aligned}$$

Problem $(\text{RP}(t_0))$ is a two-stage stochastic dynamic programming problem. Decision 0 is described by z_0 and is not limited at all. The resulting direct expected costs are equal to $E\{c_0(z_0 - D_{t_0,t_0+l_0})\}$. Decision 1-2 is described by z_1 and z_2 , and, via the constraint $z_1 + z_2 \leq IL_{t_0+l_0,0}$, its decision space depends on $IL_{t_0+l_0,0}$. Hence, we use $IL_{t_0+l_0,0}$ to describe the state of the system when decision 1-2 is taken. This state depends on decision 2 via the relation $IL_{t_0+l_0,0} = z_0 - D_{t_0,t_0+l_0-1}$. Decision 1-2 results in direct expected costs $\sum_{n=1}^2 E\{c_n(z_n - D_{t_0+l_0,t_0+l_0+l_n}^{(n)})\}$.

We see decision 1-2 as a decision that decides on two issues simultaneously.

- the aggregate level $z_{loc} = z_1 + z_2$ to which the echelon inventory positions $IP_{t_0+l_0,n}$, $n = 1, 2$, together are increased; and
- the allocation of this total amount z_{loc} to the echelons 1 and 2, which is described by z_1 and z_2 .

Let us first look at the allocation part. Suppose that a total amount $z_{loc} = x$, $x \in \mathbb{R}$, is being allocated. Then, z_1 and z_2 are obtained by the following allocation problem:

$$(\text{AP}(x)) : \min \sum_{n=1}^2 \mathbb{E} \left\{ c_n (z_n - D_{t_0+l_0, t_0+l_0+l_n}^{(n)}) \right\}$$

$$\text{s.t. } z_1 + z_2 = x.$$

The optimal solution of problem $(\text{AP}(x))$ is denoted by $z_n^*(x)$, $n = 1, 2$, and the optimal costs are denoted by $G_{loc}(x)$. The functions $z_n^*(x)$ are called *optimal allocation functions*. Because $c_n(\cdot)$ is convex, also $\mathbb{E}\{c_n(z_n - D_{t_0+l_0, t_0+l_0+l_n}^{(n)})\}$ is convex as a function of z_n , and one can show that the optimal costs $G_{loc}(x)$ are convex as a function of x . Let S_{loc} be a point in which $G_{loc}(x)$ is minimized; this point is such that $S_n = z_n^*(S_{loc})$ minimizes $\mathbb{E}\{c_n(z_n - D_{t_0+l_0, t_0+l_0+l_n}^{(n)})\}$, $n = 1, 2$ (we allow that S_{loc} and the S_n 's are infinite). Decision 1-2 is taken optimally by increasing the sum of the echelon inventory positions $n = 1, 2$ to level $x = \min\{IL_{t_0+l_0, 0}, S_{loc}\}$, i.e., according to a base-stock policy with level S_{loc} , and by allocating according to the optimal allocations $z_n^*(x)$, $n = 1, 2$.

Given the optimal solution for decision 1-2, we obtain total cycle costs

$$G_0(z_0) = \mathbb{E}\{c_0(z_0 - D_{t_0, t_0+l_0}) + G_{loc}(\min\{z_0 - D_{t_0, t_0+l_0-1}, S_{loc}\})\}$$

as a result of the level z_0 to which $IP_{t_0, 0}$ is increased. Also, this function may be shown to be convex. Hence, for decision 0 it is optimal to follow a base-stock policy with level S_0 , where S_0 is a minimizing point of $G_0(z_0)$. The optimal costs of problem $(\text{RP}(t_0))$ are given by $G_0(S_0)$. Notice that the optimal policy for problem $(\text{RP}(t_0))$ is described by the base-stock levels S_{loc} and S_0 and the optimal allocation functions $z_n^*(x)$, $n = 1, 2$.

We now arrive at the third step of the analysis. The optimal costs $G_0(S_0)$ constitute a lower bound LB for the optimal costs C_P of the infinite-horizon problem (P). Next, suppose that we apply the optimal policy of problem $(\text{RP}(t_0))$ in each period of problem (P). Then, for echelon inventory position 0 and the sum of the echelon inventory positions $n = 1, 2$, we can follow base-stock policies with levels S_0 and S_{loc} , respectively; i.e., for these echelon inventory positions, the ordering behavior is precisely as in problem $(\text{RP}(t_0))$. However, the allocation of the amount $x = \min\{IL_{t, 0}, S_{loc}\}$ to echelons 1 and 2 at the beginning of period t may be problematic for some $t \in \mathbb{N}_0$. We would like to allocate $z_1^*(x)$ and $z_2^*(x)$, respectively, but it may happen that one level is below the current echelon inventory position. We demonstrate this by a possible sample path.

First, suppose that our distribution system is such that we have strictly increasing functions $z_n^*(x)$, $n = 1, 2$. Next, suppose that at the beginning of some period t , the echelon stock of the central warehouse is precisely equal to S_{loc} ; i.e., $IL_{t, 0} = S_{loc}$. Then, at the beginning of period t , the echelon inventory positions 1 and 2 are increased to levels $z_1^*(S_{loc}) = S_1$ and $z_2^*(S_{loc}) = S_2$, respectively, and no physical stock is left in the central warehouse. Next, suppose that in period t , zero demand occurred at local warehouse 1, and a positive demand d_2 occurs at local warehouse 2. Then, at the beginning of period $t + 1$, the echelon inventory positions of echelons 1 and 2 before ordering are equal to $\widetilde{IP}_{t+1, 1} = S_1$ and $\widetilde{IP}_{t+1, 1} = S_2 - d_2$, respectively. Next, suppose that the order placed by the central warehouse in period $t - l_0 + 1$ was zero (because the total demand in period $t - l_0$ was zero), then nothing arrives in the central warehouse in period $t + 1$ and, thus, $IL_{t+1, 0} = S_{loc} - d_2$. We now would like to allocate $z_1^*(IL_{t+1, 0})$ and $z_2^*(IL_{t+1, 0})$ to echelons 1 and 2, respectively. However,

$$z_1^*(IL_{t+1, 0}) < z_1^*(S_{loc}) = S_1 = \widetilde{IP}_{t+1, 1},$$

i.e., echelon inventory position 1 before ordering is larger than the level to which echelon inventory position 1 should be increased according to the optimal policy for problem $(\text{RP}(t_0))$. We say that there is *imbalance* between the echelon inventory positions 1 and 2.

Here, we described one situation that leads to imbalance. In general, it may occur if there is a big demand in one local warehouse, while there is a small demand in the other local warehouse, and not much stock is available at the central warehouse to balance the inventories again.

Because of a possible imbalance, the allocation cannot be executed according to the functions $z_n^*(x)$, $n = 1, 2$ in all periods. In the periods with imbalance, one can balance the echelon inventory positions as much as possible. If for local warehouse 1, the current inventory position is above the desired level according to the functions $z_n^*(x)$, then this is done by keeping echelon inventory position 1 at the current level and allocating the rest to echelon 2, and vice versa. This is known as myopic allocation. By following this rule, we obtain a feasible policy for problem (P) that leads to an upper bound UB for C_P ; this UB may be determined via simulation. We call this policy the LB heuristic. The distance between UB and C_P denotes how well the LB heuristic performs. This distance $UB - C_P$, and also the distance $UB - LB$, will be small if imbalance occurs in relatively few periods only and if the imbalance is rather limited in those periods.

Clearly, due to the phenomenon of imbalance, the analysis of §2.2 for the two-echelon, serial system cannot be directly generalized to our basic distribution system. However, the generalization is possible if we assume that the echelon inventory positions $n = 1, 2$ are always balanced after allocation in all periods. This is equivalent to allowing that an echelon inventory position $n = 1, 2$ is decreased by the allocation, i.e., the corresponding local warehouse receives a negative shipment from the central warehouse. This assumption is called the *balance assumption*. Under the balance assumption, the optimal policy of problem (RP(t_0)) is also optimal for problem (P). This implies that then a base-stock policy, in combination with the optimal allocation functions $z_n^*(x)$, $n = 1, 2$, is optimal, and the optimal base-stock levels and the functions $z_n^*(x)$ can be determined sequentially (cf. Federgruen and Zipkin [28, 29]). The latter property generalizes the decomposition result. In addition, under linear inventory holding and penalty costs, the newsboy equations for the optimal base-stock levels can be generalized (Diks and de Kok [19], Doğru et al. [23]).

5.2. Literature on Distribution Systems

The research on distribution systems has a long history. Clark and Scarf [14] recognized already that base-stock policies are not optimal in general (i.e., without the balance assumption). Eppen and Schrage [25] introduced the balance assumption for a two-echelon, distribution system consisting of a stockless central warehouse and multiple local warehouses (they called that assumption the “allocation assumption”). For a two-echelon, distribution system with a stock-keeping central warehouse, the optimality of base-stock policies under the balance assumption and the decomposition result were derived by Federgruen and Zipkin [28, 29]. Diks and de Kok [19] extended these results to multiechelon, distribution systems. In this literature, linear inventory holding and penalty costs mainly were considered; it is straightforward to extend these results to general convex cost functions $c_n(\cdot)$. Under linear inventory holding and penalty costs, newsboy equations for the optimal base-stock levels have been derived for a general distribution system with continuous demand by Diks and de Kok [19] and for a two-echelon distribution system with discrete demand by Doğru et al. [23].

The above results give useful insights, however, the balance assumption is not always justified. Hence, it is relevant to know how well base-stock policies with optimal allocation functions perform if the balance assumption is not made, i.e., how well the LB heuristic as defined above performs. In Doğru et al. [22], the performance of the LB heuristic has been evaluated in a very large test bed of more than 5,000 instances for two-echelon distribution systems with symmetric and asymmetric local warehouses and with linear inventory holding and penalty costs. Notice that the optimal costs C_P can be determined by stochastic dynamic programming, but because of the curse of dimensionality, this is only possible for small-size instances with discrete demand. For that reason, $(UB - LB)/LB$ instead of $(UB - C_P)/C_P$

was used as the measure for the performance of the LB heuristic. It appeared that the LB heuristic performs well in many instances, but a large gap $(UB - LB)/LB$ may also easily occur, and even large gaps of more than 100% were found for some instances. Large gaps mainly occur if the demands at the local warehouses have high coefficients of variation, if the central warehouse has a long lead time (which limits the speed to react on an imbalance situation), and if there is one local warehouse with a low mean demand and a very low additional holding cost parameter and another local warehouse with a higher mean demand and a much larger additional holding cost parameter. These results extend earlier results by Zipkin [70].

In a subsequent study, Dođru [21] (Chapter 4) computed the optimal policy via stochastic dynamic programming for a two-echelon distribution system with discrete demand distributions on small, finite supports. He compared the decisions taken under the optimal policy to the decisions taken under the LB heuristic for instances with large $(UB - C_P)/C_P$ ratios. This showed that in these instances, the allocation functions $z_n^*(\cdot)$ as used by the LB heuristic are fine, but that the aggregate base-stock level S_{loc} is too high or the S_0 is somewhat too low (both lead to a too-low average physical stock in the central warehouse). This suggests that in instances for which the LB heuristic performs poorly, a much better heuristic may be obtained by slightly adapting the base-stock levels S_0 and S_{loc} . One may even go further, and enumerate over all possible values of S_0 and S_{loc} and pick the combination with the lowest costs. That results in the DS heuristic as proposed by Gallego et al. [36], in a continuous review setting with Poisson demand processes. For this DS heuristic, small gaps between the average costs of the DS heuristic and the lower-bound LB were found. The experiments in both Dođru [21] and Gallego et al. [36] show that it makes sense to use base-stock policies in combination with the optimal allocation functions $z_n^*(\cdot)$. However, in several cases, we cannot use the levels of the LB heuristic, and we have to try other combinations. The latter increases the computational complexity, especially for systems with multiple echelon levels and many stockpoints.

Another way to cope with possible imbalance problems is by the assumption of alternative allocation rules. One such rule is FCFS allocation in distribution systems with continuous review, as used, for example, by Axsäter [2] and Sherbrooke [56]. In addition, one assumes base-stock policies. Then, the problem is to evaluate the system under a given base-stock policy and to optimize the base-stock levels. There has been much research in this direction; for an overview, see Axsäter [4]. Gallego et al. [36] executed an experiment in which a system with optimal allocation has been compared to a system with FCFS allocation. Optimal allocation always performed better, but the differences in costs were relatively small. Hence, FCFS allocation is a sensible option for systems with serious imbalance problems under the LB heuristic (distribution systems with low demand rates probably belong to this category). Other alternative allocation rules have been studied by Axsäter et al. [5] and Güllü et al. [42].

For systems without imbalance problems, the LB heuristic is appropriate. Variants of the LB heuristic have been developed to increase the speed of computational procedures. This was done by the assumption of linear instead of optimal allocation rules, and is useful for large-scale systems with multiechelon levels and many stockpoints; for research in this direction, see Diks and de Kok [20] and van der Heijden et al. [61], and the references therein.

5.3. General Systems and Connection with Practice

So far, we have treated multiechelon systems with a pure serial, a pure assembly/convergent, or a pure distribution/divergent structure. These systems are applicable in practice, for example, when a company is responsible for only a small part of the supply chain with such a pure structure and wants to control that part by multiechelon models. However, many other practical situations exist with a mixture of convergent and divergent structures. That leads to multiechelon models that are hard to solve to optimality, or to models with many stockpoints. For such models, a few interesting concepts have been developed.

There is one concept for general networks of stockpoints based on the principles for pure convergent and pure divergent systems as described in §§4.5 and 5.1. This concept is denoted as synchronized base-stock policies; for an extensive description, see de Kok and Fransoo [16]. The base-stock policies are called synchronized as the control of components that go into the same end-products are coordinated according to the insights for convergent systems. This concept has been applied at Philips Electronics to support weekly collaborative planning of operations by Philips Semiconductors and one of its customers, Philips Optical Storage; see de Kok et al. [18]. A second concept has been developed by Ettl et al. [26]. They use a continuous-review, base-stock policy for each stockpoint and assume FCFS allocation rules; this is in line with the research on continuous-review distribution systems with FCFS allocation as mentioned in §5.2. This concept has been applied at IBM; see Lin et al. [46]. A third concept for general networks has been described by Graves and Willems [40, 41] and extends earlier work by Inderfurth [43], Inderfurth and Minner [44], and Simpson [57]. This concept builds on base-stock policies, bounded demands, and decoupling of a supply chain into subsystems via safety stocks. It is mainly developed for supply chain design and has been applied at Eastman Kodak.

All three concepts have led to huge cost savings at the companies where they were applied, and, thus, these concepts have been successful already. Nevertheless, further research is desired to improve and extend them. In the first two concepts, several approximate steps are made in the evaluation of base-stock policies and optimization of base-stock levels to obtain efficient solution procedures for large networks. In the third concept, simplifying assumptions are made for the same purpose. First of all, it is relevant to study the effect of these approximations/assumptions on the quality of the generated solutions, i.e., on the distance between the generated solutions and optimal solutions (where in the case of the third concept optimal solutions for the model without simplifying assumptions are meant). Second, it would be interesting to compare these concepts for a setting in which all three concepts can be applied. Third, in all three concepts, no capacity constraints and batching rules are taken into account. If the hierarchical approach as discussed at the beginning of §1 is adopted, then one may deal with capacity issues at the first decision level via appropriately set batching rules, and at the second level decisions may be supported by multiechelon models that respect these batching rules. This suggests to incorporate insights from serial systems with fixed batch sizes and fixed replenishment intervals, cf. §§4.6 and 4.7. If the first-level decisions lead to capacity constraints (or, better workload control rules) for single or multiple items, those constraints have to be taken into account as well; although this will be hard. In fact, even single-product multiechelon models with a capacity constraint per stage are already hard (e.g., Glasserman and Tayur [37], Parker and Kapuscinski [48], and the references therein). Fourth, the first two concepts are appropriate for operational planning, but in practice they will be applied in a rolling horizon setting, and the effect of that deserves special attention.

6. A Classification of Multiechelon Systems and Conclusion

As we have seen in the previous sections, there are several multiechelon systems for which many nice results are obtained. For those systems (generalized) base-stock policies are optimal and a decomposition result applies for the optimal base-stock or reorder levels. In addition, for many of these systems, newsboy equations have been derived. Also, these systems are where newsvendor bounds (cf. Shang and Song [54, 55]) are most likely to work. We call these systems “*nice*” systems, and they are listed in the upper part of Table 1, where we distinguish two subclasses: systems for which newsboy equations have been derived and systems for which they have not been derived (at least, not yet; we believe that they do exist for these systems). The nice systems have in common that all main results are obtained via a single-cycle analysis, for which a stochastic dynamic program with a finite number of stages has to be solved. For these systems, successive cycles are more or less decoupled.

TABLE 1. A classification of multiechelon systems.

Nice systems

Systems for which newsboy equations have been derived:

- Standard serial system (§3.1)
- Assembly system (§4.5, Rosling [51])
- Serial system with a fixed batch size per stage (§4.6, Chen [10])
- Serial system with a fixed replenishment interval per stage (§4.7, van Houtum et al. [66])
- Distribution system under the balance assumption (§5.1)
- Serial system with a capacity constraint at the most upstream stage (§4.8, Zijm and van Houtum [69])

Systems for which no newsboy equations have been derived (at least, not yet):

- Serial system with fixed batch sizes and fixed replenishment intervals (§4.7, Chao and Zhou [8])
- Serial system with advanced demand information (§4.8, Gallego and Özer [33])
- Serial system with Markov-modulated demand (§4.8, Chen and Song [11])
- Serial system with a fixed ordering cost at the most upstream stage (§4.8, Clark and Scarf [15])

Complicated systems

- Distribution system without balance assumption (§5.2)
- Distribution systems with FCFS allocation (§5.2)
- Systems with a mixed convergent-divergent structure (§5.3)
- Systems with a capacity constraint at each stage (§5.3)
- Systems with a fixed ordering cost at each stage (§4.6, Clark and Scarf [15])

In the lower part of Table 1, we have listed a number of systems that we call “*complicated*” *systems*. For these systems, there is a kind of coupling (or, dependence) between successive cycles. The structure of optimal policies cannot be derived via a single-cycle analysis. Also, that structure may be rather complicated and, thus, unattractive for practical purposes. For these systems, it may be sensible (and justified) to assume (generalized) base-stock policies, as in the concepts for general systems that we discussed in §5.3. But there is no decomposition result anymore, and, thus, optimal base-stock levels have to be determined in an alternative way. In fact, even an evaluation of a base-stock policy may already be complicated.

The distinction between nice and complicated systems is delicate (as delicate as between product-form and nonproduct-form networks in the area of queueing networks). Apart from the issues raised at the end of §5.3, future research may be devoted to that distinction as well. That may lead to a bigger set of nice systems and improved insights for heuristic solutions for complicated systems.

References

[1] I. J. B. F. Adan, M. J. A. van Eenige, and J. A. C. Resing. Fitting discrete distributions on the first two moments. *Probability in the Engineering and Informational Sciences* 9:623–632, 1996.

[2] S. Axsäter. Simple solution procedures for a class of two-echelon inventory problems. *Operations Research* 38:64–69, 1990.

[3] S. Axsäter. *Inventory Control*. Kluwer, Boston, MA, 2000.

[4] S. Axsäter. Supply chain operations: Serial and distribution inventory systems, Ch. 10. A. G. de Kok and S. C. Graves, eds. *Supply Chain Management: Design, Coordination and Operation*. Handbooks in OR & MS. Elsevier, Amsterdam, The Netherlands, 2003.

[5] S. Axsäter, J. Marklund, and E. A. Silver. Heuristic methods for centralized control of one-warehouse, N -retailer inventory systems. *Manufacturing & Service Operations Management* 4:75–97, 2002.

[6] T. Boyaci and G. Gallego. Serial production/distribution systems under service constraints. *Manufacturing & Service Operations Management* 3:43–50, 2001.

- [7] T. Boyaci, G. Gallego, K. H. Shang, and J. S. Song. Erratum to bounds in Serial production/distribution systems under service constraints. *Manufacturing & Service Operations Management* 5:372–374, 2003.
- [8] X. Chao and S. X. Zhou. Optimal policies for multi-echelon inventory system with batch ordering and periodic batching. Working paper, North Carolina State University, Raleigh, NC, 2005.
- [9] F. Chen. Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Science* 44:S221–S234, 1998.
- [10] F. Chen. Optimal policies for multi-echelon inventory problems with batch ordering. *Operations Research* 48:376–389, 2000.
- [11] F. Chen and J. S. Song. Optimal policies for multiechelon inventory problems with Markov-modulated demand. *Operations Research* 49:226–234, 2001.
- [12] F. Chen and Y. S. Zheng. Lower bounds for multi-echelon stochastic inventory problems. *Management Science* 40:1426–1443, 1994.
- [13] A. J. Clark. A dynamic, single-item, multi-echelon inventory model. Research report, RAND Corporation, Santa Monica, CA, 1958.
- [14] A. J. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management Science* 6:475–490, 1960.
- [15] A. J. Clark and H. Scarf. Approximate solutions to a simple multi-echelon inventory problem, K. J. Arrow, S. Karlin, and H. Scarf, eds. *Studies in Applied Probability and Management Science*. Stanford University Press, Stanford, CA, 88–100, 1962.
- [16] A. G. de Kok and J. C. Fransoo. Planning supply chain operations: Definition and comparison of planning concepts, Ch. 12. A. G. de Kok and S. C. Graves, eds. *Supply Chain Management: Design, Coordination and Cooperation*. Handbooks in OR & MS. Elsevier, Amsterdam, The Netherlands, 2003.
- [17] A. G. de Kok and S. C. Graves, eds. *Supply Chain Management: Design, Coordination and Cooperation*. Handbooks in OR & MS. Elsevier, Amsterdam, The Netherlands, 2003.
- [18] A. G. de Kok, F. Janssen, J. van Doremalen, E. van Wachem, M. Clerkx, and W. Peeters. Philips Electronics synchronizes its supply chain to end the bullwhip effect. *Interfaces* 35:37–48, 2005.
- [19] E. B. Diks and A. G. de Kok. Optimal control of a divergent multi-echelon inventory system. *European Journal of Operational Research* 111:75–97, 1998.
- [20] E. B. Diks and A. G. de Kok. Computational results for the control of a divergent N -echelon inventory system. *International Journal of Production Economics* 59:327–336, 1999.
- [21] M. K. Doğru. Optimal control of one-warehouse multi-retailer systems: An assessment of the balance assumption. Ph.D. thesis, Technische Universiteit Eindhoven, Eindhoven, The Netherlands, 2006.
- [22] M. K. Doğru, A. G. de Kok, and G. J. van Houtum. A numerical study on the effect of the balance assumption in one-warehouse multi-retailer inventory systems. Working paper, Technische Universiteit Eindhoven, Eindhoven, The Netherlands, 2006.
- [23] M. K. Doğru, A. G. de Kok, and G. J. van Houtum. Newsvendor characterizations for one-warehouse multi-retailer systems with discrete demand. Working paper, Technische Universiteit Eindhoven, Eindhoven, The Netherlands, 2006.
- [24] M. K. Doğru, G. J. van Houtum, and A. G. de Kok. Newsboy equations for optimal reorder levels of serial inventory systems with fixed batch sizes. Working paper, Technische Universiteit Eindhoven, Eindhoven, The Netherlands, 2006.
- [25] G. Eppen and L. Schrage. Centralized ordering policies in a multi-warehouse system with lead times and random demand. L. B. Schwartz, ed., *Multi-Level Production/Inventory Control Systems: Theory and Practice*. North-Holland, Amsterdam, The Netherlands, 51–67, 1981.
- [26] M. Ettli, G. E. Feigin, G. Y. Lin, and D. D. Yao. A supply network model with base-stock control and service requirements. *Operations Research* 48:216–232, 2000.
- [27] H. Everett, III. Generalized Lagrange multiplier method for solving problems of optimum allocation of resources. *Operations Research* 11:399–417, 1963.
- [28] A. Federgruen and P. H. Zipkin. Allocation policies and cost approximations for multilocation inventory systems. *Management Science* 30:69–84, 1984.
- [29] A. Federgruen and P. H. Zipkin. Computational issues in an infinite horizon, multi-echelon inventory model. *Operations Research* 32:818–836, 1984.

- [30] A. Federgruen and P. H. Zipkin. An inventory model with limited production capacity and uncertain demands, I. The average cost criterion. *Mathematics of Operations Research* 11:193–207, 1986.
- [31] A. Federgruen and P. H. Zipkin. An inventory model with limited production capacity and uncertain demands, II. The discounted cost criterion. *Mathematics of Operations Research* 11:208–216, 1986.
- [32] K. Feng and U. S. Rao. Echelon-stock (R, nT) control in two-stage serial stochastic inventory systems. *Operations Research Letters*. Forthcoming. 2006.
- [33] G. Gallego and Ö. Özer. Optimal replenishment policies for multiechelon inventory problems under advance demand information. *Manufacturing & Service Operations Management* 5:157–175, 2003.
- [34] G. Gallego and Ö. Özer. A new algorithm and a new heuristic for serial supply systems. *Operations Research Letters* 33:349–362, 2005.
- [35] G. Gallego and P. H. Zipkin. Stock positioning and performance estimation in serial production-transportation systems. *Manufacturing & Service Operations Management* 1:77–88, 1999.
- [36] G. Gallego, Ö. Özer, and P. H. Zipkin. Bounds, heuristics, and approximations for distribution systems. *Operations Research*. Forthcoming. 2006.
- [37] P. Glasserman and S. R. Tayur. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science* 41:263–281, 1995.
- [38] L. Gong, A. G. de Kok, and J. Ding. Optimal leadtimes planning in a serial production system. *Management Science* 40:629–632, 1994.
- [39] S. C. Graves. A multiechelon model with fixed replenishment intervals. *Management Science* 42:1–18, 1996.
- [40] S. C. Graves and S. P. Willems. Optimizing strategic safety stock placement in supply chains. *Manufacturing & Service Operations Management* 2:68–83, 2000.
- [41] S. C. Graves and S. P. Willems. Erratum: Optimizing strategic safety stock placement in supply chains. *Manufacturing & Service Operations Management* 5:176–177, 2003.
- [42] R. Güllü, G. J. van Houtum, F. Z. Sargut, and N. K. Erkip. Analysis of a decentralized supply chain under partial cooperation. *Manufacturing & Service Operations Management* 7:229–247, 2005.
- [43] K. Inderfurth. Safety stock optimization in multi-stage inventory systems. *International Journal of Production Economics* 24:103–113, 1991.
- [44] K. Inderfurth and S. Minner. Safety stocks in multi-stage inventory systems under different service levels. *European Journal of Operational Research* 106:57–73, 1998.
- [45] L. J. G. Langenhoff and W. H. M. Zijm. An analytical theory of multi-echelon production/distribution systems. *Statistica Neerlandica* 44:149–174, 1990.
- [46] G. Lin, M. Ettl, S. Buckley, S. Bagchi, D. D. Yao, B. L. Naccarato, R. Allan, K. Kim, and L. Koenig. Extended-enterprise supply-chain management at IBM Personal Systems Group and other divisions. *Interfaces* 30:7–25, 2000.
- [47] T. Osogami and M. Harchol-Balter. Closed form solutions for mapping general distributions to quasi-minimal PH distributions. *Performance Evaluation* 63:524–552, 2006.
- [48] R. P. Parker and R. Kapuscinski. Optimal policies for a capacitated two-echelon inventory system. *Operations Research* 52:739–755, 2004.
- [49] E. L. Porteus. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Palo Alto, CA, 2002.
- [50] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York, 1994.
- [51] K. Rosling. Optimal inventory policies for assembly systems under random demand. *Operations Research* 37:565–579, 1989.
- [52] H. Scarf. The optimality of (S, s) policies in the dynamic inventory problem, Ch. 13. K. Arrow, S. Karlin, and P. Suppes, eds. *Mathematical Methods in the Social Sciences*. Stanford University Press, Palo Alto, CA, 1960.
- [53] R. Schassberger. *Warteschlangen*. Springer, Berlin, 1973.
- [54] K. H. Shang and J. S. Song. Newsvendor bounds and heuristic for optimal policies in serial supply chains. *Management Science* 49:618–638, 2003.
- [55] K. H. Shang and J. S. Song. Supply chains with economies of scale: Single-stage heuristic and approximations. Working paper, Duke University, Durham, NC, 2005.

- [56] C. C. Sherbrooke. METRIC: A multi-echelon technique for recoverable item control. *Operations Research* 16:122–141, 1968.
- [57] K. F. Simpson. In-process inventories. *Operations Research* 6:863–871, 1958.
- [58] S. R. Tayur. Computing the optimal policy for capacitated inventory models. *Communications in Statistics-Stochastic Models* 9:585–598, 1993.
- [59] S. R. Tayur, R. Ganeshan, and M. Magazine, eds. *Quantitative Models for Supply Chain Management*. Kluwer, Boston, MA, 1999.
- [60] H. C. Tijms. *Stochastic Models: An Algorithmic Approach*. Wiley, New York, 1994.
- [61] M. C. van der Heijden, E. B. Diks, and A. G. de Kok. Stock allocation in general multi-echelon distribution systems with (R, S) order-up-to policies. *International Journal of Production Economics* 49:157–174, 1997.
- [62] G. J. van Houtum and W. H. M. Zijm. Computational procedures for stochastic multi-echelon production systems. *International Journal of Production Economics* 23:223–237, 1991.
- [63] G. J. van Houtum and W. H. M. Zijm. Incomplete convolutions in production and inventory models. *OR Spektrum* 19:97–107, 1997.
- [64] G. J. van Houtum and W. H. M. Zijm. On the relation between service and cost models for general inventory systems. *Statistica Neerlandica* 54:127–147, 2000.
- [65] G. J. van Houtum, K. Inderfurth, and W. H. M. Zijm. Materials coordination in stochastic multiechelon systems. *European Journal of Operational Research* 95:1–23, 1996.
- [66] G. J. van Houtum, A. Scheller-Wolf, and J. Yi. Optimal control of serial inventory systems with fixed replenishment intervals. *Operations Research*. Forthcoming, 2006.
- [67] P. M. Vanden Bosch and D. C. Dietz. Scheduling and sequencing arrivals to an appointment system. *Journal of Service Research* 4:15–25, 2001.
- [68] P. P. Wang. Sequencing and scheduling N customers for a stochastic server. *European Journal of Operational Research* 119:729–738, 1999.
- [69] W. H. M. Zijm and G. J. van Houtum. On multi-stage production/inventory systems under stochastic demand. *International Journal of Production Economics* 35:391–400, 1994.
- [70] P. H. Zipkin. On the imbalance of inventories in multi-echelon systems. *Mathematics of Operations Research* 9:402–423, 1984.
- [71] P. H. Zipkin. *Foundations of Inventory Management*. Irwin/McGraw Hill, New York, 2000.