Operations Research

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

Asymptotically Optimal Inventory Control for Assemble-to-Order Systems with Identical Lead Times

Martin I. Reiman, Qiong Wang

To cite this article:
Martin I. Reiman, Qiong Wang (2015) Asymptotically Optimal Inventory Control for Assemble-to-Order Systems with Identical Lead Times. Operations Research 63(3):716-732. http://dx.doi.org/10.1287/opre.2015.1372

OPERATIONS RESEARCH
Vol. 63, No. 3, May–June 2015, pp. 716–732
ISSN 0030-364X (print) | ISSN 1526-5463 (online)
http://dx.doi.org/10.1287/opre.2015.1372

Asymptotically Optimal Inventory Control for Assemble-to-Order Systems with Identical Lead Times

Martin I. Reiman
Alcatel-Lucent Bell Labs, Murray Hill, New Jersey 07974,
marty@research.bell-labs.com

Qiong Wang
Industrial and Enterprise Systems Engineering, University of Illinois at Urbana–Champaign, Urbana, Illinois 61801,
qwang04@illinois.edu

Optimizing multiproduct assemble-to-order (ATO) inventory systems is a long-standing difficult problem. We consider ATO systems with identical component lead times and a general "bill of materials." We use a related two-stage stochastic program (SP) to set a lower bound on the average inventory cost and develop inventory control policies for the dynamic ATO system using this SP. We apply the first-stage SP optimal solution to specify a base-stock replenishment policy, and the second-stage SP recourse linear program to make allocation decisions. We prove that our policies are asymptotically optimal on the diffusion scale, so the percentage gap between the average cost from its lower bound diminishes to zero as the lead time grows.

Subject classifications: assemble-to-order; inventory management; stochastic linear program; stochastic control; asymptotic optimality; diffusion scale.
Area of review: Stochastic Models.
History: Received May 2013; revisions received July 2014, January 2015; accepted February 2015.

1. Introduction

The assemble-to-order (ATO) inventory system, where multiple components are used to produce multiple products, is a classical and much studied model in inventory theory. Demand for the products is random, while components are obtained from an uncapacitated supplier after a deterministic (component dependent) lead time. The components that are used in each product are specified in the "bill of materials" (BOM) matrix. A key assumption is that assembly is instantaneous, so that inventory is kept at the component rather than the product level. Any unfulfilled demand is backlogged and each backlogged product incurs a product-dependent constant backlog cost per unit of time. Each component in inventory incurs a component-specific inventory holding cost per unit of time. The objective is to find a control policy, which can conveniently be described via replenishment (ordering from supplier) and allocation (assembling to meet demand) policies, that minimizes the long-run average expected total cost.

Figure 1 shows two representative ATO systems that are standard in the literature. In the W system (e.g., Doğru et al. 2010), two products are assembled from three components, with both products using the same amount of the common component 0. In the M system (e.g., Doğru et al. 2014, Lu and Song 2005, Nadar et al. 2014), two components are used to assemble three products. The same amount of component i is used by products 0 and i ($i = 1, 2$). The N system (Lu et al. 2014) is a special W system with either no component 1 or 2. This paper studies ATO systems with general structure, but will occasionally use W and M systems for illustration.

Although the problem of optimizing inventory control in single-product ATO systems was solved a long time ago (Karlin and Scarf 1958, Rosling 1989), optimizing systems with multiple products has been recognized as a tremendously more difficult problem (e.g., see Song and Zipkin 2003). The difficulty arises from the need to optimize allocation of components to demands for different products, which in principle depends on the entire status of the pipeline, including not only the total numbers of ordered components that are yet to arrive, but also the exact times when such replenishment orders were placed. Tracking this information would require a huge state space that grows exponentially with the lead time. Except in a few rare cases in which special parameter values along with simple structure render the allocation decision inconsequential (Doğru et al. 2010; Lu et al. 2010; Lu et al. 2014; Reiman and Wang 2012), currently there appears to be no computationally feasible way to obtain an optimal policy.

Most prior work has focused on optimizing within particular policy classes and has yielded suboptimal policies in

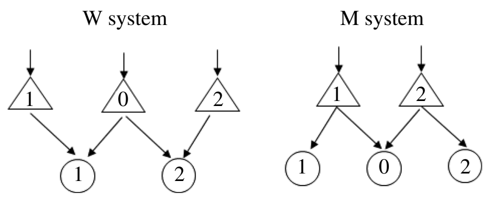

Figure 1. Two sample ATO systems.

general. To various degrees, most existing approaches rely on a first-in-first-out (FIFO) scheme to simplify allocation decisions (see, e.g., Lu and Song 2005, Lu et al. 2003 for the continuous review cases; Agrawal and Cohen 2001, Akçay and Xu 2004, Hausman et al. 1998, Zhang 1997 for discrete review cases). The FIFO scheme is implemented with commitment, i.e., a component is allocated to the demand that arrives earlier even when the corresponding product cannot be assembled immediately because other required components are not available. Among these papers it is worth singling out Lu and Song (2005), which finds the optimal base-stock levels for FIFO allocation with an arbitrary BOM. For large systems, computing these optimal base-stock levels is difficult and this issue has been addressed in recent work (van Jaarsveld and Scheller-Wolf 2015). There have been a few departures from FIFO in recent developments. For instance, "no hold-back" policies proposed in Lu et al. (2010) require that all demands should be served when components are available, and hence violate the definition of FIFO, which holds components for demands that arrive earlier regardless of whether they can be served. However, optimization over this class of policies, including the replenishment decision, is not considered in Lu et al. (2010), or any subsequent work that we are aware of.

Another way to simplify this problem is to focus on ATO systems with exponentially distributed component supply time. Exploiting the memoryless property, one can obtain a Markov decision process with a reasonable state space and characterize the optimal policies (Benjaafar and ElHafsi 2006, Nadar et al. 2014). Nevertheless, here we solve a different problem that does not have the memoryless property, so the simplification does not apply.

Doğru et al. (2010) take a different line of attack on this problem. Instead of first defining a feasible heuristic policy and proceeding to optimize relevant parameters, they start from relaxing the feasibility constraints and asking what is the best outcome that could be achieved. Guided by this outcome, they work backward to design policies that concede little ground on performance to accommodate feasibility constrains. Specifically, they consider ATO systems with identical lead times, and develop a stochastic program (SP) to set a lower bound on the inventory cost. They prescribe a base-stock policy for replenishment, where the base-stock levels are set by the first-stage optimal solution of a two-stage SP. Imitating the optimal solution of the second-stage SP recourse problem, they define a priority allocation policy for the so-called W system. Using the W system as a test bed, they demonstrate the effectiveness of the approach by showing it is exactly optimal in two special cases and outperforms other schemes in a broad set of numerical experiments.

Although the approach taken in Doğru et al. (2010) points to a promising new path forward, the development is far from over. The allocation policy there only applies to the W system, and hence needs to be generalized to systems with general BOMs. It is also desirable to justify the approach by proving that it satisfies some well-established performance criterion. In addition, several technical gaps in Doğru et al. (2010) need to be filled. The lower bound SP there is stated in terms of an infimum, which may not be attained. That paper presents a transformed SP, with a minimum that is attained, but only for the W system. Base-stock levels in Doğru et al. (2010) are set by a different SP from the lower bound SP, giving rise to the need to reconcile the two SPs.

In this paper, we solve these outstanding issues and deliver a complete asymptotic framework for optimizing ATO systems with identical lead times. The framework resembles the "four-step" method introduced in Harrison (1988) that has since been applied repeatedly to problems involving control of stochastic processing networks (e.g., Harrison and Wein 1990, Harrison 1996, Harrison and López 1999, Plambeck and Ward 2006, and many others). In our problem context, these four steps are the following:

1. Introduce a two-stage SP whose optimal solution is a lower bound on the average inventory cost that can be achieved by any feasible inventory policy.
2. Solve the SP.
3. Use the SP optimal solution to design a dynamic control policy for the associated ATO system.
4. Prove that the policy developed in step 3 is asymptotically optimal on the diffusion scale.

We resort to the use of the lower bound SP provided in Doğru et al. (2010) for step 1. Whereas Doğru et al. (2010) present a transformed SP to guarantee the minimum is reached for the W system, we provide, in Theorem 2, a transformed SP that works for all BOMs. This result makes a solution of the SP possible, and hence fits in step 2. Based on this development, we extend the base-stock policy in Doğru et al. (2010) into a family of policies and thus bridge the difference between the aforementioned two SPs. We also develop an allocation principle that spawns a family of feasible allocation policies, which include the priority policy for the W system as a special case and thus completes and generalizes the partial results in Doğru et al. (2010) on step 3. Step 4 is not discussed in Doğru et al. (2010). We fill the blank by developing an asymptotic analysis that proves our control scheme is optimal on the diffusion scale, i.e., as the lead time grows, the percentage difference of the resulting inventory cost from its lower bound diminishes to zero.

Asymptotic analysis has long been applied to a variety of stochastic control problems. Plambeck and Ward advanced its use to optimizing ATO production/inventory systems (Plambeck and Ward 2006). Their control policy involves a one-time pricing decision that determines a stochastic demand process and a capacity decision that determines a stochastic production process. The difference between demand and production yields a component shortage process, which is an input to their allocation policy, designed to minimize the costs of delaying revenues and holding inventories. They develop an asymptotic analysis in the "high-volume" region in which demand arrival rates at the same price levels are scaled up proportionally. Using the criterion of total discounted profit, they prove that their policy is asymptotically optimal on the diffusion scale, i.e., when arrival rates increase, the percentage difference of the discounted profit from its maximum value approaches zero. Their allocation policy and asymptotic analysis in the high-volume region are adopted by Lu et al. (2014) in their study of inventory control in the aforementioned N systems.

Our problem context differs from Plambeck and Ward (2006) by focusing on replenishment instead of production. In the latter case, the inventory cost is driven by the limitation on the flexibility to change production capacities to fit with demand volumes instantaneously. Hence, a high-volume asymptotic regime (Plambeck and Ward 2006), created by scaling up demand arrival rates, is suitable for asymptotic analysis. In our problem, components are ordered from outside vendors, and the delay between placing such orders and receiving them gives rise to the inventory cost. Correspondingly, we develop an asymptotic analysis for long lead times. Since keeping demand arrival rates fixed and extending the lead time has the same effect as keeping the lead time fixed and raising the arrival rates, our asymptotic regime can be easily transformed into the high-volume regime in Plambeck and Ward (2006). (The reverse is not true since there is no notion of lead time in the ATO production/inventory models of the latter.)

In comparison with the asymptotic analysis in Plambeck and Ward (2006), we have made major breakthroughs on the following two fronts. First, to guarantee asymptotic optimality, the allocation policy in Plambeck and Ward (2006) requires demands to be served only at discrete points of time with a sufficient separation between them. The idea is to use the interim between two consecutive allocation points to accumulate backlogs and inventories, used as "safety stocks" to overshadow randomness of the system at the time of allocation. This big-step approach excludes continuous allocation, so most of the time, components are held back from all product demands, including the most valuable ones. In comparison, allocation policies developed in this paper allow demands to be served continuously over time while achieving the same level of asymptotic optimality as Plambeck and Ward (2006) (albeit for a different objective function).

Second, the ATO system in Plambeck and Ward (2006) is a production/inventory system in which component supply is controlled by an open-loop production policy. As a consequence, random deviations of the amount of a component produced from the optimal amount can accumulate over time, as does the resulting financial cost. This does not pose a problem to the task in Plambeck and Ward (2006), which is cross-functional planning that involves pricing, production, and sequencing decisions to optimize the net present value (NPV) of the total profit. The discount rate embedded in the objective function tapers off deviations faster than they can accumulate. However, for our problem, the traditional formulation is to minimize the long-run average inventory cost. In the absence of discounting to bound future costs, we need the expected inventory cost of our systems to converge *uniformly* to its lower bound over an *infinite* time horizon. We prove that under our inventory control policies, this condition is achievable on the diffusion scale.

The rest of this paper is organized as follows. We define the inventory control problem in §2. We discuss the use of the SP to set a lower bound on the inventory cost and to develop inventory policies in §3. We carry out the asymptotic analysis in §4. A clear future goal is to extend the results of this paper to systems with nonidentical lead times. To this end, we discuss related challenges in §5. The electronic companion (available as supplemental material at http://dx.doi.org/10.1287/opre.2015.1372) contains most of the proofs.

2. The Inventory Control Problem

We consider an ATO system that has m products and n components. The bill of materials is given by the matrix A, of which element a_{ji} represents the amount of component j ($1 \leqslant j \leqslant n$) needed to assemble one unit of product i ($1 \leqslant i \leqslant m$), and the jth column, $\mathbf{A}_j$ ($1 \leqslant j \leqslant n$), gives the use of component j by all products. We assume that a_{ji} ($1 \leqslant i \leqslant m$, $1 \leqslant j \leqslant n$) are nonnegative integers. (We can thus handle any rational quantities by appropriate definition of a component.) We denote by $\bar{a}$ the largest element of A and $\underline{a}$ the smallest nonzero element. As alluded to in the introduction, we assume that all components have the same replenishment lead time L. Although this paper focuses on continuous-review models, our analysis can be extended to periodic-review models with little effort.

There are three semi-infinite time intervals that arise in our model and analysis. Demand for products arrives over the interval $[0, \infty)$. Orders for components can be placed over the interval $[-L, \infty)$. Finally, we do our cost accounting over the interval $[L, \infty)$, which we call our optimization horizon. Demand is modeled by the vector process $\mathcal{D} = \{\mathcal{D}(t), t \geqslant 0\}$, where

$$\mathcal{D}(t) = (\mathcal{D}_1(t), \ldots, \mathcal{D}_m(t)), \quad t \geqslant 0,$$

$\mathcal{D}_i(t)$ is the amount of demand for product i ($1 \leqslant i \leqslant m$) that arrives within the interval $[0, t]$, and $\mathcal{D}(0^-) = 0$. (All sample paths are taken to be right continuous.) We assume that $\mathscr{D}$ is compound Poisson: the number of orders is a Poisson process $\Lambda = \{\Lambda(t), t \geqslant 0\}$ with

$$E[\Lambda(1)] = \lambda,$$

and there is an associated i.i.d. sequence of random vectors that give order sizes. A generic element of this sequence is denoted by $\mathbf{S} = (S_1, \ldots, S_m)$, where S_i is the order size for product i $(1 \leqslant i \leqslant m)$. Although the order size vectors are independent, the components S_i $(1 \leqslant i \leqslant m)$ can be dependent. We assume that $\mathbf{S}$ has a finite moment of order $(2+\delta)$, i.e.,

$$\eta_i \equiv E[S_i^{2+\delta}] < \infty, \quad 1 \leqslant i \leqslant m,$$

where $\delta > 0$ can be arbitrary small. Thus demands arriving per unit time have finite means

$$\boldsymbol{\mu} = (\mu_1, \ldots, \mu_m), \quad \text{where } \mu_i = \lambda E[S_i],\ 1 \leqslant i \leqslant m,$$

and a finite covariance matrix Σ. We denote the variances by σ_{ii}, $1 \leqslant i \leqslant m$.

The demand that arrives at a particular time t (if any) is denoted by

$$\mathbf{d}(t) \equiv \mathscr{D}(t) - \mathscr{D}(t^-), \quad t \geqslant 0,$$

and demand that arrives between two distinct time points is denoted by

$$\mathbf{D}(t_1, t_2) \equiv \mathscr{D}(t_2) - \mathscr{D}(t_1), \quad t_2 > t_1 \geqslant 0.$$

With a slight abuse of notation, let

$$\mathbf{D}(t) \equiv \mathbf{D}(t-L, t), \quad t \geqslant L,$$

denote demand that arrives within the lead time immediately before time t. Since the arrival process is compound Poisson, $\mathbf{D}(t)$ has the same distribution for all $t \geqslant L$. Let $\mathbf{D} = (D_1, \ldots, D_m)$ be a random vector that has this distribution. We refer to $\mathbf{D}$ as the lead time demand, and note that

$$E[\mathbf{D}] = L\boldsymbol{\mu} \quad \text{and} \quad E[(\mathbf{D} - E[\mathbf{D}])(\mathbf{D} - E[\mathbf{D}])'] = L\Sigma. \quad (1)$$

As previously mentioned, the control policy, which we denote by p, can conveniently be described via a replenishment policy and an allocation policy. A replenishment policy gives rise to the process $\{\mathscr{R}(t), t \geqslant -L\}$, where

$$\mathscr{R}(t) = (\mathscr{R}_1(t), \ldots, \mathscr{R}_n(t)),$$

and $\mathscr{R}_j(t)$ represents the amount of component j $(1 \leqslant j \leqslant n)$ ordered during $[-L, t]$. Let $\mathscr{R}(-L^-) = 0$ and define

$$\mathbf{r}(t) \equiv \mathscr{R}(t) - \mathscr{R}(t^-), \quad t \geqslant -L,$$
$$\mathbf{R}(t_1, t_2) \equiv \mathscr{R}(t_2) - \mathscr{R}(t_1), \quad t_2 > t_1 \geqslant -L, \quad \text{and}$$
$$\mathbf{R}(t) \equiv \mathbf{R}(t-L, t), \quad t \geqslant 0,$$

to be orders placed at time t, during the period $(t_1, t_2]$, and within the lead time immediately before time t, respectively. Since one cannot order a negative quantity, each element of $\mathscr{R}(t)$ is nondecreasing over t $(t \geqslant -L)$. Recall that we allow replenishment orders starting at $t = -L$.

An allocation policy gives rise to the process $\{\mathscr{Z}(t), t \geqslant 0\}$, where

$$\mathscr{Z}(t) = (\mathscr{Z}_1(t), \ldots, \mathscr{Z}_m(t)),$$

and $\mathscr{Z}_i(t)$ is the amount of product i $(1 \leqslant i \leqslant m)$ served during $[0, t]$. Let $\mathscr{Z}(0^-) = 0$ and define

$$\mathbf{z}(t) \equiv \mathscr{Z}(t) - \mathscr{Z}(t^-), \quad t \geqslant 0,$$
$$\mathbf{Z}(t_1, t_2) \equiv \mathscr{Z}(t_2) - \mathscr{Z}(t_1), \quad t_2 > t_1 \geqslant 0, \quad \text{and}$$
$$\mathbf{Z}(t) \equiv \mathbf{Z}(t-L, t), \quad t \geqslant L,$$

to be demand served at time t, during the period $(t_1, t_2]$, and within the lead time immediately before time t, respectively. Since one can only serve a positive amount of demand, $\mathscr{Z}(t)$ is element-wise nondecreasing over time.

The event sequence at any time is as follows: arrival of new demands, receipt of previously ordered components, allocation of available components to serve demands, and placement of new orders. Not all events happen at each time. Unsatisfied demands are backlogged and unused components stay in inventory. Let

$$\mathbf{B}(t) = (B_1(t), \ldots, B_m(t)) \quad \text{and} \quad \mathbf{I}(t) = (I_1(t), \ldots, I_n(t))$$

be the backlog and inventory levels at t $(t \geqslant 0)$ after these events. Then

$$\mathbf{B}(t) = \mathbf{B}(t^-) + \mathbf{d}(t) - \mathbf{z}(t) \quad \text{and} \quad (2)$$
$$\mathbf{I}(t) = \mathbf{I}(t^-) + \mathbf{r}(t-L) - A\mathbf{z}(t), \quad (3)$$

where we define $\mathbf{B}(0^-) = 0$ and $\mathbf{I}(0^-) = 0$. Between two distinct time points $0 \leqslant t_1 < t_2$,

$$\mathbf{B}(t_2) = \mathbf{B}(t_1) + \mathbf{D}(t_1, t_2) - \mathbf{Z}(t_1, t_2) \quad \text{and} \quad (4)$$
$$\mathbf{I}(t_2) = \mathbf{I}(t_1) + \mathbf{R}(t_1 - L, t_2 - L) - A\mathbf{Z}(t_1, t_2). \quad (5)$$

Specializing these conditions to $t_1 = t - L$ and $t_2 = t$,

$$\mathbf{B}(t) = \mathbf{B}(t-L) + \mathbf{D}(t) - \mathbf{Z}(t) \quad \text{and} \quad (6)$$
$$\mathbf{I}(t) = \mathbf{I}(t-L) + \mathbf{R}(t-L) - A\mathbf{Z}(t). \quad (7)$$

For discussions below, we define $\mathbf{B}^-(t)$ and $\mathbf{I}^-(t)$ as backlog and inventory levels at time t $(t \geqslant 0)$ after demand arrival and components receipt but before the allocation of components. Hence,

$$\mathbf{B}(t) = \mathbf{B}^-(t) - \mathbf{z}(t), \quad t \geqslant 0,$$
$$\mathbf{I}(t) = \mathbf{I}^-(t) - A\mathbf{z}(t), \quad t \geqslant 0.$$

We define

$$\mathbf{Q}(t) \equiv A\mathbf{B}^-(t) - \mathbf{I}^-(t) = A\mathbf{B}(t) - \mathbf{I}(t), \quad t \geqslant 0, \quad (8)$$

as the component shortage at time t: there are more components j on hand than the amount needed to clear the existing backlog if $Q_j(t) \leqslant 0$ and less than enough if $Q_j(t) > 0$ $(1 \leqslant j \leqslant n)$.

Let h_j be the cost of holding a unit of component j $(1 \leqslant j \leqslant n)$ in inventory per unit of time and b_i be the cost of keeping a unit of demand for product i $(1 \leqslant i \leqslant m)$ in backlog per unit of time. We assume that $h_j > 0$ $(1 \leqslant j \leqslant n)$ and $b_i > 0$ $(1 \leqslant i \leqslant m)$. Satisfying a unit of demand for product i $(1 \leqslant i \leqslant m)$ removes from the system a cost of

$$c_i = b_i + \sum_{j=1}^{n} a_{ji} h_j$$

per unit of time. We refer to c_i as the unit inventory cost of product i $(1 \leqslant i \leqslant m)$. Let

$$\mathbf{b} = (b_1, \ldots, b_m), \quad \mathbf{h} = (h_1, \ldots, h_n), \quad \text{and} \quad \mathbf{c} = (c_1, \ldots, c_m).$$

Hence the total expected inventory plus backlog cost at time t is

$$\mathbf{h} \cdot \mathbf{I}(t) + \mathbf{b} \cdot \mathbf{B}(t).$$

The goal of inventory management is to develop a policy p to minimize the following long-run average expected total inventory cost:

$$\mathscr{C}^p = \limsup_{T \to \infty} \frac{1}{T} \int_{L}^{T+L} E[\mathbf{b} \cdot \mathbf{B}(t) + \mathbf{h} \cdot \mathbf{I}(t)]\, dt. \tag{9}$$

Note that the integral in (9) is over the interval $[L, T+L]$, which is consistent with our optimization horizon, rather than $[0, T]$ as in Equation (8) of Doğru et al. (2010). The difference between these intervals becomes immaterial when the limit $T \to \infty$ is taken. To be feasible, the policy cannot serve more demand than the amount arrived or the amount allowed by the available supply of required components, i.e., for any $0 \leqslant t_1 < t_2$,

$$\mathbf{Z}(t_1, t_2) \leqslant \mathbf{B}(t_1) + \mathbf{D}(t_1, t_2), \tag{10}$$

$$A\mathbf{Z}(t_1, t_2) \leqslant \mathbf{I}(t_1) + \mathbf{R}(t_1 - L, t_2 - L). \tag{11}$$

The policy needs to be nonanticipating, i.e., $\mathbf{r}(t)$ and $\mathbf{z}(t)$ can depend only on information available at t, given by $\mathbf{B}(0^-)$, $\mathbf{I}(0^-)$, $\{\mathscr{D}(s), 0 \leqslant s \leqslant t\}$, $\{\mathscr{Z}(s), 0 \leqslant s < t\}$, and $\{\mathscr{R}(s), -L \leqslant s < t\}$.

3. Stochastic Program: Lower Bound and Policy Development

In §3.1, we transform the lower bound SP in Doğru et al. (2010) into another SP of which the minimum can always reached. Based on this transformation, we develop inventory policies in §3.2.

3.1. The SP Lower Bound

The one-period ATO model considered in Song and Zipkin (2003) corresponds to the following SP:

$$C_+^* = \min_{\mathbf{y} \geqslant 0} C_+(\mathbf{y})$$

$$C_+(\mathbf{y}) = \mathbf{h} \cdot \mathbf{y} + \mathbf{b} \cdot E[\mathbf{D}] - E[\varphi_+(\mathbf{y}; \mathbf{D})], \tag{12}$$

$$\varphi_+(\mathbf{y}; \mathbf{D}) = \max_{\mathbf{z} \geqslant 0} \{\mathbf{c} \cdot \mathbf{z} \mid \mathbf{z} \leqslant \mathbf{D}, A\mathbf{z} \leqslant \mathbf{y}\}.$$

This is also a simple version of "newsvendor network" of Harrison and Van Mieghem (1999). This SP involves minimizing the inventory cost for a particular point in time in the absence of any past history. Hence the replenishment policy reduces to a one-time decision, $\mathbf{y} = (y_1, \ldots, y_n)$, where y_j is the order quantity of component j $(1 \leqslant j \leqslant n)$. The allocation decision also simplifies to the choice of a vector $\mathbf{z} = (z_1, \ldots, z_m)$, where z_i is the amount of product i demand $(1 \leqslant i \leqslant m)$ to serve. Demand is given by a random vector $\mathbf{D}$. Cost parameters $\mathbf{b}$, $\mathbf{h}$, and $\mathbf{c}$ are analogous to those in §2. Components are allocated after observing all demands, an arrangement that is impossible for a dynamic ATO system where demand arrivals continue into the indefinite future.

Although intuitively it may seem that the SP (12) provides, at each point in time, a relaxation of the dynamic inventory control problem, Doğru et al. (2010) show that this is not the case. In particular, in the dynamic system, past decisions may lead to backlogs of lower value products, allowing the manager to divert components ordered to clear these backlogs to serve more valuable new arrivals if needed. The SP in (12) has no such flexibility. Doğru et al. (2010) introduced the following relaxed version of (12), where the initial backlog ($\boldsymbol{\alpha}$) can be chosen optimally:

$$\underline{C} = \inf_{\boldsymbol{\alpha} \geqslant 0} \Phi(\boldsymbol{\alpha}) \tag{13}$$

$$\Phi(\boldsymbol{\alpha}) \equiv \inf_{\mathbf{y} \geqslant 0} \big\{\mathbf{h} \cdot \mathbf{y} + \mathbf{b} \cdot E[(\boldsymbol{\alpha} + \mathbf{D})] - E[\varphi_+(\mathbf{y}; \boldsymbol{\alpha} + \mathbf{D})]\big\},$$

$$\varphi_+(\mathbf{y}; \boldsymbol{\alpha} + \mathbf{D}) = \max_{\mathbf{z} \geqslant 0} \big\{\mathbf{c} \cdot \mathbf{z} \mid \mathbf{z} \leqslant \boldsymbol{\alpha} + \mathbf{D}, A\mathbf{z} \leqslant \mathbf{y}\big\}. \tag{14}$$

By Theorem 2.1 in Doğru et al. (2010), if $\mathbf{D}$ has the same distribution as the lead time demand, then $\underline{C}$ in (13) is a lower bound on the cost objective in (9). Below is a restatement of their result.

Theorem 1. *Let p be any feasible inventory control policy. Let $\mathscr{C}^p$ be the resulting long-run average total expected inventory cost as defined in (9). Let $\underline{C}$ be given by (13). Then*

$$\underline{C} \leqslant \mathscr{C}^p. \tag{15}$$

It is easy to verify that $\Phi(\boldsymbol{\alpha})$ is nonincreasing in $\boldsymbol{\alpha}$. Because of the unbounded support of $\mathbf{D}$, the infimum in (13) may not be attained at finite values of $\boldsymbol{\alpha}$ and $\mathbf{y}$ even though $\underline{C}$ is finite. To deal with this issue we consider an alternative relaxation of (12). Instead of optimizing the initial backlog level, we allow $\mathbf{y}$ and $\mathbf{z}$ to be negative, giving rise to the following SP:

$$C^* = \inf_{\mathbf{y}\in\mathbf{R}^n} C(\mathbf{y}) \tag{16}$$

$$C(\mathbf{y}) \equiv \mathbf{h}\cdot\mathbf{y} + \mathbf{b}\cdot E[\mathbf{D}] - E[\varphi(\mathbf{y};\mathbf{D})],$$

$$\varphi(\mathbf{y};\mathbf{D}) \equiv \max_{\mathbf{z}\in\mathbf{R}^m}\{\mathbf{c}\cdot\mathbf{z} \mid \mathbf{z}\leqslant\mathbf{D},\, A\mathbf{z}\leqslant\mathbf{y}\}.$$

The analogy of not having these nonnegativity constraints in the SP for the inventory system is that the manager is allowed to undo past ordering and allocation decisions that she may regret after observing new demands. Thus, it is not surprising that C^* is also a lower bound on $\mathscr{C}^p$. Moreover, the following result, whose proof is in the electronic companion, shows that C^* attains the same value as $\underline{C}$, but unlike the latter SP, the solution is attained at a finite value, $\mathbf{y}^*$ (possibly not unique), which is bounded by a finite constant that depends only on the A, $\mathbf{b}$, and $E[\mathbf{D}]$. Incidentally, the optimal solution of the original SP (12), $\mathbf{y}^o$, which is restricted to be positive, can also be bounded by the same constant from above.

THEOREM 2. *There exists* $\mathbf{y}^*\in\mathbf{R}^n$ *such that*

$$C(\mathbf{y}^*) = C^* = \underline{C}.$$

Let $\zeta_j \equiv \min_{i:\,a_{ji}>0}\{b_i/a_{ji}\}$ $(1\leqslant j\leqslant n)$ *and define*

$$M \equiv \max_{1\leqslant j\leqslant n}\left\{\mathbf{A}_j\cdot E[\mathbf{D}] + \frac{\mathbf{b}\cdot E[\mathbf{D}]}{\min(\zeta_j, h_j)}\right\}.$$

For any $\mathbf{y}^*$ *such that* $C(\mathbf{y}^*) = C^*$,

$$|y_j^*| \leqslant M, \quad 1\leqslant j\leqslant n,$$

and for any $\mathbf{y}^o$ *that optimizes* (12),

$$0\leqslant y_j^o \leqslant M, \quad 1\leqslant j\leqslant n.$$

According to the theorem, the lower bound in (13), $\underline{C}$, can be determined by an alternative SP

$$\underline{C} = \min_{\mathbf{y}\in\mathbf{R}^n} C(\mathbf{y}) \tag{17}$$

$$C(\mathbf{y}) \equiv \mathbf{h}\cdot\mathbf{y} + \mathbf{b}\cdot E[\mathbf{D}] - E[\varphi(\mathbf{y};\mathbf{D})]$$

$$\varphi(\mathbf{y};\mathbf{D}) \equiv \max_{\mathbf{z}\in\mathbf{R}^m}\{\mathbf{c}\cdot\mathbf{z} \mid \mathbf{z}\leqslant\mathbf{D},\; A\mathbf{z}\leqslant\mathbf{y}\}.$$

The SP has complete recourse, and the finiteness of the optimal solution allows its use for defining inventory control policies, as we will see next.

3.2. Policy Development

Based on the SPs discussed above, we formulate a family of base-stock replenishment policies and define an allocation principle that admits a set of feasible allocation policies. Both apply to systems with general BOMs and identical lead times.

3.2.1. Base-Stock Replenishment Policies. A base-stock policy keeps the inventory position of each component, i.e., the total inventory (on hand and in pipeline) in excess of the amount needed to clear existing backlogs, at a constant level. Doğru et al. (2010) suggest a base-stock policy that uses the first-stage optimal solution of the one-period ATO model (12), $\mathbf{y}^o$, as base-stock levels. Since $\mathbf{D}$ in the SP is defined to have the same distribution as demands over a lead time in the corresponding ATO system, the policy naturally imitates the optimal order solution of the SP.

Theorem 2 in the last section allows us to set base-stock levels at $\mathbf{y}^*$, the first-stage optimal solution of the lower bound SP (17), which is more directly related to the minimum inventory cost. As is justified by the asymptotic analysis in §4, we propose the following family of base-stock policies that uses weighted average of $\mathbf{y}^*$ and $\mathbf{y}^o$ as base-stock levels, i.e.,

$$\mathbf{y}^\gamma = \gamma\mathbf{y}^* + (1-\gamma)\mathbf{y}^o, \quad 0\leqslant\gamma\leqslant 1. \tag{18}$$

Both $\mathbf{y}^*$ $(\gamma=1)$ and $\mathbf{y}^o$ $(\gamma=0)$ are possible base-stock levels. As we show in Theorem 4, any $\gamma\in[0,1]$ yields an asymptotically optimal policy. Based on our numerical results so far, we recommend using $\gamma=0$, so that $\mathbf{y}^o$ is used for base-stock levels. When (12) and/or (17) has multiple optimal solutions, $\mathbf{y}^\gamma$ is nonunique. In this case, any $\mathbf{y}^\gamma$ can be employed, consistently for all time, as base-stock levels. Unlike $\mathbf{y}^o$, $\mathbf{y}^*$ does not need to be positive, so a component can have a negative base-stock level. In this case, the manager does not order a component until its inventory position (on hand plus in transit minus backlog) falls below the base-stock level, and starts ordering the component to keep its inventory position exactly at this constant level.

3.2.2. Allocation Principle. Based on the second-stage recourse LP (14), Doğru et al. (2010) propose a priority policy for allocating components in the W system, but stop short of defining an allocation policy for systems with general BOMs. This task, which is far less obvious, is completed by our development below.

Recall from (8) that $\mathbf{Q}(t)$ is the shortage of on-hand inventory for clearing existing backlogs at time t. Notice that $\mathbf{Q}(t)$ (and also $\mathbf{B}(t)$ and $\mathbf{I}(t)$ below) depend on the choice of base-stock levels from (18), even though for simplicity of presentation, we do not add additional notation to explicitly indicate this connection. Shortages are carried by product backlogs. Given shortage levels $\mathbf{Q}(t)$, the optimal backlog levels that minimize the total inventory cost are

$$\mathbf{B}^*(t) = \arg\min\{\mathbf{c}\cdot\mathbf{B} \mid \mathbf{B}\geqslant\mathbf{0},\, A\mathbf{B}\geqslant\mathbf{Q}(t)\}. \tag{19}$$

We refer to $\mathbf{B}^*(t)$ as the backlog targets. To avoid wild fluctuations of the backlog targets in response to small perturbations of component shortage, and moreover, to guarantee our approach to be asymptotically optimal, we require $\mathbf{B}^*(t)$ to be uniformly Lipschitz continuous with respect to $\mathbf{Q}(t)$, which follows immediately from Hoffman's lemma (Hoffman 1952) (also see Schrijver 1998, Theorem 10.5) if the optimal solution of (19) is a singleton for all possible values of $\mathbf{Q}(t)$. If the optimal solution is not unique, Hoffman's lemma only implies that the solution sets are Lipschitz continuous, and we need to solve the following quadratic programming problem to select the solution with minimum Euclidean norm:

$$\min_{\mathbf{B}\geqslant\mathbf{0}}\{\|\mathbf{B}\| \mid A\mathbf{B}\geqslant\mathbf{Q}(t),\, \mathbf{c}\cdot\mathbf{B}\leqslant\psi^*\} \tag{20}$$

where ψ^* is the optimal objective value of (19) and $\|\mathbf{B}\|$ denotes Euclidean norm of $\mathbf{B}$. In general, Lipschitz continuity of sets does not imply Lipschitz continuity of the minimum norm selection; cf. Aubin and Frankowska (2008). However, for the special case of (20), Han et al. (2012, Theorem 4.1.d) show that the optimal solution is unique and Lipschitz continuous.

From (19), if $\mathbf{Q}(t)\leqslant\mathbf{0}$ (no component shortage), then $\mathbf{B}^*(t)=\mathbf{0}$, a target that can be easily reached by clearing all existing backlogs. However, when $Q_j(t)>0$ for some j ($1\leqslant j\leqslant n$), we may not attain the corresponding backlog target because it is optimized based on aggregated component shortage with no regard to existing backlogs of individual products. For instance, in the W system, if $Q_1(t)\leqslant 0$ and $Q_0(t)>0$ (here for notational convenience, we deviate from the general formulation slightly by referring to the common component as component 0), then $B_1^*(t)=0$ and $B_2^*(t)=Q_0(t)>0$ by (19). But if product 2 has no existing backlog, the shortage of component 0 is entirely in the backlog of product 1 ($B_1(t)$), which exceeds its target and the excess cannot be immediately removed. To bridge this gap between optimality of backlog targets and feasibility of allocation policies, we define the following allocation principle.

Allocation Principle. *No product should have its backlog level strictly above its target if all required components have sufficient inventories, i.e., an allocation policy must yield*

$$[B_i(t)-B_i^*(t)]^+\wedge\Big[\min_{j:\,a_{ji}>0}\{(I_j(t)-a_{ji}+1)^+\}\Big]=0,\qquad 1\leqslant i\leqslant m. \tag{21}$$

In addition, any product whose backlog level is below or at the target is not served, i.e.,

$$z_i(t)\leqslant[B_i^-(t)-B_i^*(t)]^+,\quad 1\leqslant i\leqslant m. \tag{22}$$

In general, our allocation principle can be implemented by the following procedure: An inventory manager monitors component shortage (or excess), which changes upon arrival of new demand(s) or receipt of ordered component(s). The change triggers a reset of backlog targets by re-solving (19). The manager then uses all available components to reduce the excess of existing backlogs (which includes newly arrived demands) over their targets until there is no overage to be eliminated or no more required components are available. In the latter case, the manager can clear backlogs in any sequence that she finds appropriate, e.g., giving higher priority to products with higher c_is, as long as no product is served if its backlog is below, or has been reduced to, its targeted level.

In some systems, this implementation procedure simplifies to an intuitive rule. For instance, in the W system, which includes the N system as a special case, following this procedure leads to the same priority rule of Doğru et al. (2010), i.e., clear as many backlogs as possible while giving the priority of using component 0 to product 1 (assuming $c_1\geqslant c_2$ without loss of generality). This is because in the W system, (19) sets backlog targets at

$$\begin{aligned}\mathbf{B}^*(t)&=\arg\min\{c_1B_1+c_2B_2 \mid B_i\geqslant 0,\ B_i\geqslant Q_i(t),\\ &\qquad\qquad B_1+B_2\geqslant Q_0(t),\ i=1,2\}\\ &=(Q_1^+(t),\,Q_2^+(t)\vee(Q_0(t)-Q_1^+(t))^+).\end{aligned} \tag{23}$$

To satisfy (21), product 1 is served until component 1 runs out ($B_1(t)=B_1^*(t)=Q_1^+(t)$) or component 0 has no inventory left ($I_0(t)=0$, recall that $a_{01}=1$). Product 2's backlog is cleared if and only if component 2 is available ($B_2(t)>Q_2^+(t)$) and doing so will not cause shortage of component 0 to be a bottleneck for serving product 1 ($B_2(t)>(Q_0(t)-Q_1^+(t))^+$).

In other systems, the implementation can be simplified to qualitatively different rules that apply to different parameter regions. Take the M system as an example. As is explained in detail in Doğru et al. (2014), when $c_0<\max(c_1,c_2)$, the principle reduces to a static priority rule under which products with higher c_is have the precedence to be served. When $\max(c_1,c_2)\leqslant c_0\leqslant c_1+c_2$, a state-dependent priority applies, in which case product 0 has the priority when only one of the other two products compete with it for a component, but has to yield the priority if both of the other two products are waiting to be served. When $c_0>c_1+c_2$, the implementation simplifies to a reservation policy, i.e., product i is not served unless enough component i ($i=1,2$) has been set aside to clear existing backlog of product 0.

To put our approach into perspective, we compare it with the allocation policy in Plambeck and Ward (2006). The latter also uses an LP to determine backlog targets as the basis for allocating components. Their LP is similar to ours except it includes an additional constraint (Plambeck and Ward 2006, Equation (15)), which in our notation would be

$$B_i^*(t)\leqslant B_i(t),\quad 1\leqslant i\leqslant m. \tag{24}$$

The new constraint ensures that no backlog target exceeds the current backlog level of the product, and hence by also satisfying other constraints, is feasible to reach. In Plambeck and Ward (2006) components are allocated only periodically. With the length of periods carefully chosen, enough backlogs can be accumulated during the interim so that (24) is rarely binding when the allocation takes place. Thus as the system scales, the probability that their backlog targets at times of allocation are the same ones as in (19) converges to 1. We serve demands continuously, in which case some backlog targets in (19) may exceed the current backlog levels, and thus are not reachable. However, the analysis in §4 will prove that under our allocation principle, the difference is asymptotically negligible.

3.2.3. Allocation Principle Under Base-Stock Replenishment. Without loss of generality, assume that at $t=0$, inventory positions are at base-stock levels $\mathbf{y}^\gamma$. Then to keep the inventory positions at the constant levels,

$$\mathbf{R}(t)=A\mathbf{D}(t), \quad t\geqslant L,$$

i.e., amounts ordered during the past lead time equals the amounts needed to satisfy demands that arrived during the same period. Since $\mathbf{R}(t)$ is the total inventory in transit at time t and

$$\mathbf{y}^\gamma=\mathbf{I}(t)+\mathbf{R}(t)-A\mathbf{B}(t), \quad t\geqslant L, \tag{25}$$

by definition, the on-hand inventory at time t satisfies

$$\mathbf{I}(t)=\mathbf{y}^\gamma+A\mathbf{B}(t)-A\mathbf{D}(t), \quad t\geqslant L. \tag{26}$$

Combining (26) with (8), component shortage process simplifies to

$$\mathbf{Q}(t)=A\mathbf{D}(t)-\mathbf{y}^\gamma, \quad t\geqslant L. \tag{27}$$

Thus by denoting $\mathbf{B}$ by $\mathbf{D}(t)-\mathbf{z}$, (19) can be transformed into

$$\begin{aligned}&\min_{\mathbf{B}\in\mathbf{R}^m}\{\mathbf{c}\cdot\mathbf{B}\mid\mathbf{B}\geqslant\mathbf{0},\,A\mathbf{B}\geqslant\mathbf{Q}(t)\}\\&\quad=\min_{\mathbf{z}\in\mathbf{R}^m}\big\{\mathbf{c}\cdot(\mathbf{D}(t)-\mathbf{z})\mid\mathbf{D}(t)-\mathbf{z}\geqslant\mathbf{0},\\&\qquad\qquad A(\mathbf{D}(t)-\mathbf{z})\geqslant A\mathbf{D}(t)-\mathbf{y}^\gamma\big\}\\&\quad=\mathbf{c}\cdot\mathbf{D}(t)-\max_{\mathbf{z}\in\mathbf{R}^m}\big\{\mathbf{c}\cdot\mathbf{z}\mid\mathbf{z}\leqslant\mathbf{D}(t),\,A\mathbf{z}\leqslant\mathbf{y}^\gamma\big\}.\end{aligned} \tag{28}$$

Since $\mathbf{D}(t)$ equals in distribution to $\mathbf{D}$, the transformation in (28) shows that (19) is equivalent to (17), which means our backlog targets are set by the recourse LP of the lower bound SP. In effect, one many consider the allocation either as distribution of components, which is optimized in (17), or as an allotment of component shortage, which is optimized by (19). Either way, if replenishment in an ATO system follows the first-stage optimal solution of (17) and the allocation outcome is at the level dictated by (19), then the inventory cost should reach the lower bound set by the SP. Lemma 1 formalizes this result.

Lemma 1 (Verification Lemma). *Any inventory policy that uses a base-stock replenishment policy with base-stock levels $\mathbf{y}^*$, where $\mathbf{y}^*$ is an optimal solution of (17), is optimal if the resulting backlog levels satisfy*

$$\mathbf{c}\cdot E[\mathbf{B}(t)]=\mathbf{c}\cdot E[\mathbf{B}^*(t)] \quad \text{for all } t\geqslant L. \tag{29}$$

This lemma is not quite a special case of the Verification Lemma (Lemma 2) in Reiman and Wang (2012) since it does not involve condition 3 (Equation (19)) there. The proof is immediate from (8), (26), (28), and (29):

$$\begin{aligned}&\mathbf{b}\cdot E[\mathbf{B}(t)]+\mathbf{h}\cdot E[\mathbf{I}(t)]\\&\quad=\mathbf{c}\cdot E[\mathbf{B}(t)]+\mathbf{h}\cdot(\mathbf{y}^*-AE[\mathbf{D}(t)])\\&\quad=\mathbf{c}\cdot E[\mathbf{B}^*(t)]+\mathbf{h}\cdot\mathbf{y}^*-\mathbf{h}\cdot(AE[\mathbf{D}])\\&\quad=\mathbf{b}\cdot E[\mathbf{D}]+\mathbf{h}\cdot\mathbf{y}^*-E\Big[\max_{\mathbf{z}}\{\mathbf{c}\cdot\mathbf{z}\mid\mathbf{z}\leqslant\mathbf{D},\,A\mathbf{z}\leqslant\mathbf{y}^*\}\Big],\end{aligned}$$

i.e., the integrand in (9) reaches the lower bound given by (17) at all $t\geqslant L$.

In general, it is not possible to satisfy (29) for all time. Nevertheless, the next section will prove that, for a broader set of base-stock levels ($\mathbf{y}^\gamma$, $\gamma\in[0,1]$), the allocation principle ensures the difference vanishes as the lead time increases, so our approach is asymptotically optimal.

4. Asymptotic Analysis

For the purposes of our asymptotic analysis, we introduce a family of ATO systems indexed by the lead time L. All parameters other than L are held fixed, while $L\to\infty$. (As was noted in the introduction, our results also apply to the high-volume limit considered in Plambeck and Ward (2006), where the lead time stays constant while the order arrival rate $\boldsymbol{\lambda}$ increases.)

Let $\mathscr{C}_L^\gamma$ be the long-run average expected total inventory cost under our base-stock replenishment policy with $\mathbf{y}^\gamma$ $(0\leqslant\gamma\leqslant1)$ as the base-stock levels and an allocation policy that satisfies the allocation principle. Let $\underline{C}_L$ be the lower bound. Then our main result (Theorem 4) states that

$$\lim_{L\to\infty}\frac{\mathscr{C}_L^\gamma}{\underline{C}_L}=1. \tag{30}$$

We will first show in Theorem 3 that $L^{-1/2}\underline{C}_L$ converges to a finite positive constant. Thus (30) holds if and only if

$$\lim_{L\to\infty}\frac{\mathscr{C}_L^\gamma-\underline{C}_L}{\sqrt{L}}=0, \tag{31}$$

and (31) is actually what we show. The scaling of the costs and various stochastic processes in this system is basically that of a (functional) central limit theorem. This reflects a simple fact: as the lead time L grows, the total demand over a lead time (when properly centered and scaled) converges to a normally distributed random variable.

Observe that (31) is more stringent than the "fluid-scale" asymptotic optimality criterion. The definition of the latter is the same as in (31) except that $\sqrt{L}$ is replaced by L. A policy that is optimal only on the fluid scale may have an average cost that differs from the exact optimum by a quantity on the order of $\sqrt{L}$ (or possibly larger). In this case, the ratio in (30) may not converge to unity.

Asymptotic optimality on the fluid scale is an easy criterion that is satisfied by every existing approach that we are aware of. However, to the best of our knowledge, no procedure has been proved to be asymptotically optimal on the diffusion scale for minimizing the long-run average cost of ATO inventory systems. In fact, thought experiments illustrate that FIFO allocation does not qualify: consider the inverse V system in which one component is used to serve two products with $c_1 > c_2$ (i.e., a W system with no components 1 and 2). Let $Q(t)$ be the shortage process under a given replenishment policy. Then the optimal backlog targets in (19) are

$$B_1^*(t) = 0 \quad \text{and} \quad B_2^*(t) = Q^+(t),$$

which allocates all shortage to product 2. As we will prove later, this target can be reached on the diffusion scale under our allocation principle. However, it is easy to see the same outcome is not possible under FIFO. With no preference given to either product, if the two have the same demand arrival processes, they will have the same expected backlog level, which is on the order of $\sqrt{L}$.

Recall that FIFO is a prevailing feature of existing approaches. Even periodic-review models that implement non-FIFO policies for intraperiod allocation use FIFO for interperiod allocation (e.g., Agrawal and Cohen 2001, Akçay and Xu 2004). This means all these approaches are not optimal on the diffusion scale in general.

No-holdback (NHB) allocation, as is discussed in Lu et al. (2010), is a broad concept that may admit both good and bad policies (the latter case may include not holding back components and giving priority to low-value demands). The particular implementation of NHB considered in Lu et al. (2010), first-ready-first-serve (FRFS), becomes FIFO when applied to the above inverse V systems, and thus is not asymptotically optimal. Numerical experiments in Doğru et al. (2014) illustrate that in some cases of the M system, even the best implementation of NHB under base-stock replenishment policies will fail to meet the criterion. Intuitively, when the backlog cost of product 0 is huge and those of products 1 and 2 are small, the backlog target in (19) for product 0 is zero. Theorem 4 shows that this target is achievable on the diffusion scale under our allocation principle, which features holding back (reserving) components in this case. However, if NHB is implemented, when all products have backlogs, product i can use up component i ($i = 1, 2$) if the other component is not available for product 0. Since a base-stock policy implements order-for-order replenishment, this can lead to an extended period during which the only times when both components are available simultaneously are those when replenishment orders triggered by a prior demand for product 0 are received. In this case, the assembly of product 0 follows the delayed arrival process of the demand, so the backlog is on the order of $\sqrt{L}$. Therefore the expected total inventory cost can differ significantly on the diffusion scale from its optimal value, as the numerical results in Doğru et al. (2014) show.

Based on the discussion in §3.2, we believe the allocation policy in Plambeck and Ward (2006), for maximizing the total discounted profit in ATO production/inventory systems, can be asymptotically optimal on the diffusion scale in our case when used with base-stock levels $\mathbf{y}^\gamma$.

To indicate their dependence on L, the various random variables and processes will assume the same definitions as before except that they may have a superscript (L) or be assigned an argument L to specify the system under discussion. Some variables are centered (i.e., taking the difference from its mean). As a rule of our notation, we use $\tilde{X}$ (attaching a tilde) to denote a scaled but not centered value of variable X, and $\hat{X}$ (attaching a hat) to denote a centered and scaled value of X. For instance, analogous to the definition of $\mathbf{D}$ in §2, $\mathbf{D}^{(L)}$ denotes the random vector that has the same distribution as demands arriving over a lead time in system L. Its centered and scaled version is given by

$$\hat{\mathbf{D}}^{(L)} \equiv \frac{\mathbf{D}^{(L)} - L\boldsymbol{\mu}}{\sqrt{L}}. \tag{32}$$

Following (1),

$$E[\hat{\mathbf{D}}^{(L)}] = 0 \quad \text{and} \quad E[\hat{\mathbf{D}}^{(L)}(\hat{\mathbf{D}}^{(L)})'] = \Sigma.$$

Furthermore, for the interest of our discussion, we scale without centering demand arrivals, orders placed, and demand served at time t ($t \geqslant 0$), between times t_1 and t_2 ($0 \leqslant t_1 < t_2$), and within a lead time immediately preceding time t ($t \geqslant 0$):

$$\tilde{\mathbf{d}}^{(L)}(t) \equiv \frac{\mathbf{d}^{(L)}(Lt)}{\sqrt{L}}, \qquad \tilde{\mathbf{D}}^{(L)}(t_1, t_2) \equiv \frac{\mathbf{D}^{(L)}(Lt_1, Lt_2)}{\sqrt{L}},$$

$$\tilde{\mathbf{D}}^{(L)}(t) \equiv \frac{\mathbf{D}^{(L)}(Lt)}{\sqrt{L}}, \qquad \tilde{\mathbf{r}}^{(L)}(t) \equiv \frac{\mathbf{r}^{(L)}(Lt)}{\sqrt{L}},$$

$$\tilde{\mathbf{R}}^{(L)}(t_1, t_2) \equiv \frac{\mathbf{R}^{(L)}(Lt_1, Lt_2)}{\sqrt{L}}, \qquad \tilde{\mathbf{R}}^{(L)}(t) \equiv \frac{\mathbf{R}^{(L)}(Lt)}{\sqrt{L}},$$

and

$$\tilde{\mathbf{z}}^{(L)}(t) \equiv \frac{\mathbf{z}^{(L)}(Lt)}{\sqrt{L}}, \qquad \tilde{\mathbf{Z}}^{(L)}(t_1, t_2) \equiv \frac{\mathbf{Z}^{(L)}(Lt_1, Lt_2)}{\sqrt{L}},$$

$$\tilde{\mathbf{Z}}^{(L)}(t) \equiv \frac{\mathbf{Z}^{(L)}(Lt)}{\sqrt{L}},$$

respectively. We also scale the backlog and inventory levels at each time by

$$\tilde{\mathbf{B}}^{(L)}(t) \equiv \frac{\mathbf{B}^{(L)}(Lt)}{\sqrt{L}} \quad \text{and} \quad \tilde{\mathbf{I}}^{(L)}(t) \equiv \frac{\mathbf{I}^{(L)}(Lt)}{\sqrt{L}}, \quad t \geqslant 0,$$

respectively. Combining these two quantities yields the scaled version of the shortage process, i.e.,

$$A\tilde{\mathbf{B}}^{(L)}(t) - \tilde{\mathbf{I}}^{(L)}(t) = \tilde{\mathbf{Q}}^{(L)}(t) \equiv \frac{\mathbf{Q}^{(L)}(Lt)}{\sqrt{L}} \quad t \geqslant 0.$$

The next two subsections present some preliminary lemmas. In §4.1, we discuss properties of scaled (and sometimes also centered) versions of demand processes. In §4.2, we address properties of our inventory policies. The discussion culminates in §4.3 where we prove that these policies are asymptotically optimal on the diffusion scale.

4.1. The Demand Process

We first consider demands arriving at a given point in time, that is, order sizes. As L increases, even though the distribution of the order size $\mathbf{S}$ does not change, the maximum order size over all arrivals that occur within a lead time will increase. However, the following lemma shows that its expected value is negligible on the leading order of average cost ($\sqrt{L}$) that is of interest to us.

Lemma 2. (*See the electronic companion for proof*) *Under the assumption that* $\mathbf{S}$ *has a finite moment of order* $2+\delta$ $(\delta>0)$,

$$E\Big[\sup_{t-1\leqslant\tau\leqslant t}\tilde{d}_i^{(L)}(\tau)\Big]\leqslant 3\lambda^{1/(2+\delta)}(1+\eta_i)L^{-\delta/(2(2+\delta))},\quad 1\leqslant i\leqslant m. \tag{33}$$

Besides the size of single arrivals, we are also interested in the total demand arriving between two distinct time points, $\mathbf{D}^{(L)}(Lt_1,Lt_2)$, $0\leqslant t_1<t_2$. Their centered and scaled values are

$$\hat{\mathbf{D}}^{(L)}(t_1,t_2)\equiv\frac{\mathbf{D}^{(L)}(Lt_1,Lt_2)-L(t_2-t_1)\boldsymbol{\mu}}{\sqrt{L}},\quad 0\leqslant t_1<t_2,$$

where

$$E[\hat{\mathbf{D}}^{(L)}(t_1,t_2)]=\mathbf{0}\quad\text{and}$$
$$E[\hat{\mathbf{D}}^{(L)}(t_1,t_2)\hat{\mathbf{D}}^{(L)}(t_1,t_2)']=(t_2-t_1)\Sigma.$$

Although we do not explicitly use this fact, it should be noted that a functional central limit theorem indicates that when $L\to\infty$, $\hat{\mathbf{D}}(0,t)$ converges to a Brownian motion. In the following lemma, we establish upper bounds on the centered-and-scaled version of the expected maximum demand that occurs within certain subintervals of a lead time. Note that although the result is stated with respect to the time interval $[0,1]$, because of stationarity this result holds for any interval of the form $[t,t+1]$ for $t\geqslant 0$.

Lemma 3. (*See the electronic companion for proof*) *For all* $i=1,\ldots,m$,

$$E\Big[\sup_{0\leqslant\tau\leqslant L^{-1/4}}|\hat{D}_i^{(L)}(0,\tau)|\Big]\leqslant(1+\sigma_{ii}^2)L^{-1/8}, \tag{34}$$

$$E\Big[\sup_{L^{-1/4}\leqslant\tau\leqslant 1}\big(|\hat{D}_i^{(L)}(0,\tau)|-\sqrt{L}\tau\kappa\big)^+\Big]\leqslant\frac{\sigma_{ii}^2}{\kappa}L^{-1/4}, \tag{35}$$

where κ *can be any positive constant.*

We also center and scale demand over a lead time, $\hat{\mathbf{D}}^{(L)}(t)\equiv\hat{\mathbf{D}}^{(L)}(t-1,t)$, $t\geqslant 1$. The distribution of $\hat{\mathbf{D}}^L(t)$ is the same as $\hat{\mathbf{D}}^{(L)}$ in (32).

4.2. Inventory Policy

Following our development in §3.2, for each system L, we prescribe a family of base-stock replenishment policies with base-stock levels

$$\mathbf{y}^{(L)\gamma}=\gamma\mathbf{y}^{(L)*}+(1-\gamma)\mathbf{y}^{(L)o},\quad \gamma\in[0,1].$$

Here $\mathbf{y}^{(L)*}$ and $\mathbf{y}^{(L)o}$ are, respectively, optimal solutions of the SPs

$$\min_{\mathbf{y}} C^{(L)}(\mathbf{y}), \tag{36}$$

where

$$C^{(L)}(\mathbf{y})=\mathbf{h}\cdot\mathbf{y}+\mathbf{b}\cdot(L\boldsymbol{\mu})-E[\varphi(\mathbf{y};\mathbf{D}^{(L)})]$$
$$\varphi(\mathbf{y};\mathbf{D}^{(L)})=\max_{\mathbf{z}}\{\mathbf{c}\cdot\mathbf{z}\mid\mathbf{z}\leqslant\mathbf{D}^{(L)},\,A\mathbf{z}\leqslant\mathbf{y}\},$$

and

$$\min_{\mathbf{y}\geqslant\mathbf{0}} C_+^{(L)}(\mathbf{y}), \tag{37}$$

where

$$C_+^{(L)}(\mathbf{y})=\mathbf{h}\cdot\mathbf{y}+\mathbf{b}\cdot(L\boldsymbol{\mu})-E[\varphi_+(\mathbf{y};\mathbf{D}^{(L)})]$$
$$\varphi_+(\mathbf{y};\mathbf{D}^{(L)})=\max_{\mathbf{z}\geqslant\mathbf{0}}\{\mathbf{c}\cdot\mathbf{z}\mid\mathbf{z}\leqslant\mathbf{D}^{(L)},\,A\mathbf{z}\leqslant\mathbf{y}\}.$$

We can center and scale decision variables by letting

$$\hat{\mathbf{y}}=\frac{\mathbf{y}-LA\boldsymbol{\mu}}{\sqrt{L}}\quad\text{and}\quad\hat{\mathbf{z}}=\frac{\mathbf{z}-L\boldsymbol{\mu}}{\sqrt{L}} \tag{38}$$

to transform the two SPs into

$$\min_{\hat{\mathbf{y}}}\hat{C}^{(L)}(\hat{\mathbf{y}}), \tag{39}$$

where

$$\hat{C}^{(L)}(\hat{\mathbf{y}})=\mathbf{h}\cdot\hat{\mathbf{y}}-E[\hat{\varphi}(\hat{\mathbf{y}};\hat{\mathbf{D}}^{(L)})]$$
$$\hat{\varphi}(\hat{\mathbf{y}};\hat{\mathbf{D}}^{(L)})=\max_{\hat{\mathbf{z}}}\{\mathbf{c}\cdot\hat{\mathbf{z}}\mid\hat{\mathbf{z}}\leqslant\hat{\mathbf{D}}^{(L)},\,A\hat{\mathbf{z}}\leqslant\hat{\mathbf{y}}\},$$

and

$$\min_{\hat{\mathbf{y}}\geqslant-\sqrt{L}A\boldsymbol{\mu}}\hat{C}_+^{(L)}(\hat{\mathbf{y}}), \tag{40}$$

where

$$\hat{C}_+^{(L)}(\hat{\mathbf{y}})=\mathbf{h}\cdot\hat{\mathbf{y}}-E[\hat{\varphi}_+(\hat{\mathbf{y}};\hat{\mathbf{D}}^{(L)})]$$
$$\hat{\varphi}_+(\hat{\mathbf{y}};\hat{\mathbf{D}}^{(L)})=\max_{\hat{\mathbf{z}}\geqslant-\sqrt{L}\boldsymbol{\mu}}\{\mathbf{c}\cdot\hat{\mathbf{z}}\mid\hat{\mathbf{z}}\leqslant\hat{\mathbf{D}}^{(L)},\,A\hat{\mathbf{z}}\leqslant\hat{\mathbf{y}}\},$$

respectively. One may easily verify that

$$\hat{C}^{(L)}(\hat{\mathbf{y}})=\frac{C^{(L)}(L^{1/2}\hat{\mathbf{y}}+LA\boldsymbol{\mu})}{\sqrt{L}}\quad\text{and}$$
$$\hat{C}_+^{(L)}(\hat{\mathbf{y}})=\frac{C_+^{(L)}(L^{1/2}\hat{\mathbf{y}}+LA\boldsymbol{\mu})}{\sqrt{L}}. \tag{41}$$

Therefore the optimal solution(s) of (39) and (40) are related to those of (36) and (37) as follows:

$$\hat{\mathbf{y}}^{(L)*} = \frac{\mathbf{y}^{(L)*} - LA\boldsymbol{\mu}}{\sqrt{L}} \quad \text{and} \quad \hat{\mathbf{y}}^{(L)o} = \frac{\mathbf{y}^{(L)o} - LA\boldsymbol{\mu}}{\sqrt{L}}. \tag{42}$$

Let

$$\underline{\hat{C}}^{(L)} \equiv \frac{\underline{C}_L}{\sqrt{L}} \quad \text{and} \quad \hat{C}_+^{(L)*} \equiv \min_{\hat{\mathbf{y}} \geqslant -\sqrt{L}A\boldsymbol{\mu}} \hat{C}_+^{(L)}(\hat{\mathbf{y}}).$$

When $\mathbf{y}^{(L)\gamma}$ ($\gamma \in [0,1]$) are used as base-stock levels,

$$\hat{\mathbf{y}}^{(L)\gamma} = \gamma \hat{\mathbf{y}}^{(L)*} + (1-\gamma)\hat{\mathbf{y}}^{(L)o}, \quad \gamma \in [0,1],$$

gives their differences from mean demands, measured on the $\sqrt{L}$ scale. The base-stock levels can grow without bound as the lead time increases whereas by Lemma 4, whose proof is in the electronic companion, the scaled differences from the means are bounded regardless of how long the lead time.

Lemma 4. *There exists a constant M such that for all $L > 0$,*

$$|\hat{y}_j^{(L)\gamma}| \leqslant M, \quad 1 \leqslant j \leqslant n, \ \gamma \in [0,1]. \tag{43}$$

We can similarly specify the allocation principle for the Lth system. Operating at the leading order, for each L, we solve the following (scaled) version of (19):

$$\tilde{\mathbf{B}}^{(L)*}(t) = \operatorname{argmin}\{\mathbf{c} \cdot \mathbf{B} \mid \mathbf{B} \geqslant \mathbf{0}, A\mathbf{B} \geqslant \tilde{\mathbf{Q}}^{(L)}(t)\}, \quad t \geqslant 0, \tag{44}$$

to set the backlog target. Recall that in (19), we apply a minimum norm selection to keep the optimal solution $\mathbf{B}^*(t)$ uniformly Lipschitz continuous with respect to $\mathbf{Q}(t)$. Here we continue the same approach to make $\tilde{\mathbf{B}}^{(L)*}(t)$ uniformly Lipschitz continuous with respect to $\tilde{\mathbf{Q}}^{(L)}(t)$. As a consequence, the following lemma (whose proof is in the electronic companion) shows that we may use demand fluctuations to bound the changes of backlog target over time.

Lemma 5. *There exists a constant g that depends only on A and $\mathbf{c}$, such that for any $t_2 > t_1 \geqslant 1$,*

$$\begin{aligned}&\left|\tilde{B}_i^{(L)*}(t_2) - \tilde{B}_i^{(L)*}(t_1)\right| \\ &\quad \leqslant g \sum_{l=1}^{m} \left|\hat{D}_l^{(L)}(t_1, t_2) - \hat{D}_l^{(L)}(t_1 - 1, t_2 - 1)\right|, \quad 1 \leqslant i \leqslant m,\end{aligned} \tag{45}$$

and for $t \geqslant 1$,

$$\begin{aligned}&\left|\tilde{B}_i^{(L)*}(t) - \tilde{B}_i^{(L)*}(t^-)\right| \\ &\quad \leqslant g \sum_{l=1}^{m} \left|\tilde{d}_l^{(L)}(t) - \tilde{d}_l^{(L)}(t-1)\right|, \quad 1 \leqslant i \leqslant m.\end{aligned} \tag{46}$$

4.3. Asymptotic Optimality

We introduce the “limit” stochastic program as follows. Let

$$\boldsymbol{\xi} = (\xi_1, \ldots, \xi_m)$$

be a normally distributed random vector with mean $\mathbf{0}$ and covariance Σ. For $\mathbf{y} \in \mathbf{R}^n$, let

$$\hat{C}(\mathbf{y}) \equiv \mathbf{h} \cdot \mathbf{y} - E[\varphi(\mathbf{y}; \boldsymbol{\xi})],$$

where

$$\varphi(\mathbf{y}; \boldsymbol{\xi}) \equiv \max_{\mathbf{z} \in \mathbf{R}^m} \{\mathbf{c} \cdot \mathbf{z} \mid \mathbf{z} \leqslant \boldsymbol{\xi}, A\mathbf{z} \leqslant \mathbf{y}\}.$$

Note that for given $(\mathbf{y}, \boldsymbol{\xi})$, $\varphi(\mathbf{y}; \boldsymbol{\xi})$ is always feasible, so

$$\hat{C}^* \equiv \inf_{\mathbf{y} \in \mathbf{R}^n} \hat{C}(\mathbf{y})$$

is also a two-stage SP with complete recourse. Since $\mathbf{y} = \mathbf{0}$ and $z_i = \xi_i^-$ $(1 \leqslant i \leqslant m)$ are feasible,

$$\hat{C}^* \leqslant \hat{C}(\mathbf{0}) = -E[\varphi(\mathbf{0}; \boldsymbol{\xi})] \leqslant -\sum_{i=1}^{m} c_i E[\xi_i^-] < \sum_{i=1}^{m} c_i E[|\xi_i|].$$

Therefore $\hat{C}^*$ is a positive constant, which, by Theorem 3, is the limit of the (scaled) lower bound on our cost objective.

Theorem 3. (*See the electronic companion for proof*) *As the lead time L increases, the optimal objective values of scaled original and lower bound SPs converge to that of the limit SP, i.e.,*

$$\lim_{L \to \infty} \hat{C}_+^{(L)*} = \lim_{L \to \infty} \underline{\hat{C}}^{(L)} = \hat{C}^*. \tag{47}$$

The first equality in (47) provides a key step in justifying the use of (18) to set base-stock levels. It shows that, on the diffusion scale, there is no difference in optimality between using the original SP (12) and the lower bound SP (17) since their objective values converge to the same limit. The second equality, as a result of the central limit theorem, shows that on the diffusion scale, the lower bound on our cost objective converges to a constant. Hence our policies attain the target in (30) if they are optimal on the diffusion scale (31). Below we establish the latter by first presenting a sufficient condition in an asymptotic verification lemma and then proving our policies do satisfy this condition.

As a parallel to the verification lemma in §3.2, the asymptotic verification lemma shows that regardless which base-stock levels in (18) are used, inventory control is asymptotically optimal on the diffusion scale if the allocation policy keeps *scaled* backlog levels at their targets in (44).

LEMMA 6 (ASYMPTOTIC VERIFICATION LEMMA). (*See the electronic companion for proof.*) *Any family of inventory policies that use base-stock replenishment with base-stock levels* $\mathbf{y}^{(L)\gamma}$, *defined in* (18), *in the system with lead time* L *is asymptotically optimal if*

$$\lim_{L\to\infty} \sup_{t\geqslant 1}\{|E[\tilde{\mathbf{B}}^{(L)}(t)]-E[\tilde{\mathbf{B}}^{(L)*}(t)]|\}=0. \tag{48}$$

It is worth mentioning that the above Asymptotic Verification Lemma can be satisfied by policies that do not strictly follow our allocation principle. In particular, we generalize (21) to

$$[B_i^{(L)}(t)-B_i^{*(L)}(t)]^+\wedge\Big[\min_{j:a_{ji}>0}\{(I_j^{(L)}(t)-a_{ji}+1-w_{ji}^{(L)})^+\}\Big]=0, \quad 1\leqslant i\leqslant m, \tag{49}$$

where $w_{ji}^{(L)}\geqslant 0$ is the amount of component j $(1\leqslant j\leqslant n)$ held back from product i $(1\leqslant i\leqslant m)$, which can be determined differently for different cases. Regardless of the method of their determination, as long as

$$\lim_{L\to\infty}\frac{w_{ji}^{(L)}}{\sqrt{L}}=0, \quad 1\leqslant i\leqslant m,\ 1\leqslant j\leqslant n, \tag{50}$$

the resulting difference from (21) is invisible on the diffusion scale, so asymptotic optimality will not be affected. The generalization allows additional flexibility of policy design that can improve performance in some special cases outside the asymptotic region. For instance, when c_1 vastly exceeds c_2 in a W system, reserving a small number of component 0 for product 1 can yield a lower cost than that of following the allocation principle strictly and hence not holding back the component (Doğru et al. 2010). In the absence of compelling evidence to set them otherwise, $w_{ji}=0$ $(1\leqslant j\leqslant n,\ 1\leqslant i\leqslant m)$ is a reasonable choice.

THEOREM 4. *Let* $\{p^*(L), L>0\}$ *denote a family of policies that use base-stock replenishment with base-stock levels* $\mathbf{y}^{(L)\gamma}$ *with* $\gamma\in[0,1]$, *and use an allocation policy that satisfies the allocation principle* (*with* (21) *replaced by* (49)) *along with* (50). *Then*

$$\lim_{L\to\infty}\frac{\mathscr{C}_L^{p^*(L)}}{\underline{C}_L}=1. \tag{51}$$

Before presenting the proof, an intuitive explanation is in order. We refer to the positive difference of a backlog level from its target as the excess, and the negative difference as the deficit. The opening part of the proof makes a simple but useful observation: the definition of backlog target (20) and the requirement of the allocation principle (49) (with condition (50)) imply that no product will have a nontrivial amount of excess unless some other product has a nontrivial amount of deficit. Hence our conclusion holds if the expected deficit of any product never passes a negligible level, i.e., when condition (54) in the proof below applies. To prove the condition, we first note that the backlog targets, $\mathbf{B}^{(L)}$, are on the same order $(\sqrt{L})$ as the component shortage, $\mathbf{Q}^{(L)}$. Demand arrival rates (with time scaled, but not space) are on the order of L. Since products with a deficit are not served, as L increases, the probability that a product has a deficit persisting for more than a lead time becomes asymptotically negligible; so do the expected deficits, a situation shown by the discussion of Case 1 in (56)–(59). For sample paths where the deficit persists for less than a lead time, taken care of in (60)–(66), Lemma 3 is used to show that, again, the expected deficit is asymptotically negligible.

The proof of the above arguments is facilitated by stationary demand processes and the base-stock replenishment policy. The latter renders the past states irrelevant after a lead time, so we can develop a bound on the expected deficit over the finite time interval $[0,2]$ and apply the bound uniformly to the infinite time horizon.

PROOF. Following Lemma 6, we prove the result by showing that (48) holds.

For given L and $t\geqslant 1$, let

$$\mathscr{S}_L^+(t)=\{i\colon \tilde{B}_i^{(L)}(t)>\tilde{B}_i^{(L)*}(t),\ 1\leqslant i\leqslant m\},$$
$$\mathscr{S}_L^-(t)=\{i\colon \tilde{B}_i^{(L)}(t)<\tilde{B}_i^{(L)*}(t),\ 1\leqslant i\leqslant m\}.$$

Since $\tilde{\mathbf{B}}^{(L)*}(t)$ is feasible for (20),

$$\mathbf{A}_j\cdot\tilde{\mathbf{B}}^{(L)*}(t)\geqslant\tilde{\mathbf{Q}}^{(L)}(t)=\mathbf{A}_j\cdot\tilde{\mathbf{B}}^{(L)}(t)-\tilde{\mathbf{I}}^{(L)}(t).$$

Therefore for all $j=1,\ldots,n$,

$$\sum_{i\in\mathscr{S}_L^+(t)} a_{ji}(\tilde{B}_i^{(L)}(t)-\tilde{B}_i^{(L)*}(t)) \leqslant \tilde{I}_j^{(L)}(t)+\sum_{i\in\mathscr{S}_L^-(t)} a_{ji}(\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t)). \tag{52}$$

Observe that for every $i'\in\mathscr{S}_L^+(t)$, there exists some j' such that $a_{j'i'}>0$ and

$$\sum_{i\in\mathscr{S}_L^+(t)} a_{j'i}(\tilde{B}_i^{(L)}(t)-\tilde{B}_i^{(L)*}(t)) \leqslant \sum_{i\in\mathscr{S}_L^-(t)} a_{j'i}(\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t))-\frac{1-a_{j'i'}-w_{j'i'}^{(L)}}{\sqrt{L}}. \tag{53}$$

If not, then (52) implies that

$$B_{i'}^{(L)}(t)>B_{i'}^{(L)*}(t) \quad\text{and}\quad I_j^{(L)}(t)-a_{ji'}+1-w_{j'i'}^{(L)}>0, \quad \text{for all } j \text{ such that } a_{ji'}>0,$$

which violates the allocation principle (49). Since all terms on the left-hand side of (53) are strictly positive and $i'\in\mathscr{S}_L^+(t)$, the inequality implies that

$$\frac{1-a_{j'i'}-w_{j'i'}^{(L)}}{\sqrt{L}}<a_{j'i'}(\tilde{B}_{i'}^{(L)}(t)-\tilde{B}_{i'}^{(L)*}(t))+\frac{1-a_{j'i'}-w_{j'i'}^{(L)}}{\sqrt{L}} \leqslant \sum_{i\in\mathscr{S}_L^-(t)} a_{j'i}(\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t)).$$

Using (50),

$$\lim_{L\to\infty}\frac{1-a_{j'i'}-w_{j'i'}^{(L)}}{\sqrt{L}}=0,$$

and we prove the theorem by showing that for any $\epsilon>0$, if L is large enough, then for any $i\in\mathcal{S}_L^-(t)$,

$$E[\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t)]<\epsilon,\quad \text{for all } t\geqslant 1. \tag{54}$$

For a given $i\in\mathcal{S}_L^-(t)$, let

$$t_i^{(L)}=\sup\{\tau\colon 0\leqslant\tau\leqslant t \text{ and } \tilde{B}_i^{(L)*}(\tau)<\tilde{B}_i^{(L)}(\tau)\}.$$

We can write

$$\begin{aligned}E\big[\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t)\big]&=E\big[(\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t))\mathbf{1}(t_i^{(L)}<t-1)\big]\\&\quad+E\big[(\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t))\mathbf{1}(t_i^{(L)}\geqslant t-1)\big]\end{aligned} \tag{55}$$

and prove (54) by considering the two situations on the right-hand side separately.

In cases where $t_i^{(L)}<t-1$, $\tilde{B}_i^{(L)}(\tau)\leqslant\tilde{B}_i^{(L)*}(\tau)$ for all $\tau\in[L(t-1),Lt]$. Under our policy, no demand for product i is served during this lead time. Therefore

$$\tilde{B}_i^{(L)}(t)\geqslant\tilde{D}_i^{(L)}(t)=\hat{D}_i^{(L)}(t)+\sqrt{L}\mu_i. \tag{56}$$

Let $j(i)$ be any j such that

$$a_{j(i)i}^{-1}[\tilde{Q}_{j(i)}^{(L)}(t)]^+\equiv\max_{j:a_{ji}>0}\{a_{ji}^{-1}[\tilde{Q}_j^{(L)}(t)]^+\}\geqslant\tilde{B}_i^{(L)*}(t)\geqslant 0, \tag{57}$$

where the first inequality holds because otherwise $\tilde{\mathbf{B}}^{(L)*}(t)$ is not optimal for (44). From (26),

$$\begin{aligned}\tilde{Q}_{j(i)}^{(L)}(t)&=\mathbf{A}_{j(i)}\cdot\tilde{\mathbf{B}}^{(L)}(t)-\tilde{\mathbf{I}}_{j(i)}^{(L)}(t)=\mathbf{A}_{j(i)}\cdot\tilde{\mathbf{D}}^{(L)}(t)-y_{j(i)}^{(L)\gamma}/\sqrt{L}\\&=\mathbf{A}_{j(i)}\cdot\hat{\mathbf{D}}^{(L)}(t)-\hat{y}_{j(i)}^{(L)\gamma}.\end{aligned} \tag{58}$$

Define

$$\bar{\beta}=\bar{a}/\underline{a}\quad\text{and}\quad\bar{y}_{\min}^{(L)}=\min_{i,j}\{-|\hat{y}_j^{(L)\gamma}|/a_{ji}\}.$$

Then $\bar{\beta}\geqslant 1$ and $\bar{y}_{\min}^{(L)}$ is finite by Lemma 4. From (57) and (58),

$$\begin{aligned}\tilde{B}_i^{(L)*}(t)&\leqslant\max_{j:a_{ji}>0}\frac{[\tilde{Q}_j^{(L)}(t)]^+}{a_{ji}}\leqslant\max_{j:a_{ji}>0}\frac{[\mathbf{A}_j\cdot\hat{\mathbf{D}}^{(L)}(t)-\hat{y}_j^{(L)\gamma}]^+}{a_{ji}}\\&\leqslant\bar{\beta}\sum_{l=1}^m|\hat{D}_l^{(L)}(t)|-\bar{y}_{\min}^{(L)}.\end{aligned}$$

Apply the above inequality and using (56),

$$\begin{aligned}&E\big[(\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t))\mathbf{1}(t_i^{(L)}<t-1)\big]\\&\quad\leqslant E\bigg[\bigg(\bar{\beta}\sum_{l=1}^m|\hat{D}_l^{(L)}(t)|-\bar{y}_{\min}^{(L)}+|\hat{D}_i^{(L)}(t)|-\sqrt{L}\mu_i\bigg)^+\bigg]\\&\quad\leqslant E\bigg[\bigg((\bar{\beta}+1)\sum_{l=1}^m|\hat{D}_l^{(L)}(t)|-(\bar{y}_{\min}^{(L)}+\sqrt{L}\mu_i)\bigg)^+\bigg]\\&\quad\leqslant(\bar{\beta}+1)\sum_{l=1}^m E\bigg[\bigg(|\hat{D}_l^{(L)}(t)|-\frac{\bar{y}_{\min}^{(L)}+\sqrt{L}\mu_i}{m(\bar{\beta}+1)}\bigg)^+\bigg]\\&\quad\leqslant\frac{(\bar{\beta}+1)^2m}{\bar{y}_{\min}^{(L)}+\sqrt{L}\mu_i}\sum_{l=1}^m\sigma_{ll}^2,\end{aligned} \tag{59}$$

where the last inequality comes from Chebyshev's inequality.

If $t_i^{(L)}\geqslant t-1$, we first establish that

$$\begin{aligned}&\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)}(t)\\&\quad\leqslant(\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)*}(t_i^{(L)}))-(\tilde{B}_i^{(L)}(t)-\tilde{B}_i^{(L)}(t_i^{(L)}))\\&\qquad+|\tilde{B}_i^{(L)*}(t_i^{(L)})-\tilde{B}_i^{(L)*}(t_i^{(L)-})|\end{aligned} \tag{60}$$

by considering the following two scenarios. In scenario one, demand for product i is served at time $t_i^{(L)}$. This implies that the product's preallocation backlog exceeds its target. Following (22) and the definition of $t_i^{(L)}$,

$$\tilde{B}_i^{(L)}(t_i^{(L)})=\tilde{B}_i^{(L)*}(t_i^{(L)}),$$

so (60) holds. In scenario two, demand for product i is not served at time $t_i^{(L)}$, so

$$\tilde{B}_i^{(L)}(t_i^{(L)})=\tilde{B}_i^{(L)}(t_i^{(L)-})+\tilde{d}_i^{(L)}(t_i^{(L)}).$$

By the definition of $t_i^{(L)}$ and because $\tilde{d}_i^{(L)}(t_i^{(L)})\geqslant 0$,

$$\tilde{B}_i^{(L)*}(t_i^{(L)-})<\tilde{B}_i^{(L)}(t_i^{(L)-})\leqslant\tilde{B}_i^{(L)}(t_i^{(L)}),$$

which also leads to (60).

We now consider the right-hand side of (60), starting from the last term. From Lemma 2 and (46), there exists some constant θ_1 such that

$$E\Big[\sup_{t-1\leqslant\tau\leqslant t}\{|\tilde{B}_i^{(L)*}(\tau)-\tilde{B}_i^{(L)*}(\tau^-)|\}\Big]\leqslant\theta_1L^{-\delta/(2(2+\delta))}. \tag{61}$$

For the first two terms, applying (45) with $t_1=t_i^{(L)}$ and $t_2=t$,

$$\begin{aligned}&\tilde{B}_i^{(L)*}(t)-\tilde{B}_i^{(L)*}(t_i^{(L)})\\&\quad\leqslant g\sum_{l=1}^m\big|\hat{D}_l^{(L)}(t_i^{(L)},t)-\hat{D}_l^{(L)}(t_i^{(L)}-1,t-1)\big|.\end{aligned} \tag{62}$$

Under our inventory policy, no product i demand is served during $(t_i^{(L)},t]$, so

$$\begin{aligned}&\tilde{B}_i^{(L)}(t)-\tilde{B}_i^{(L)}(t_i^{(L)})\\&\quad\geqslant\tilde{D}_i^{(L)}(t_i^{(L)},t)=\hat{D}_i^{(L)}(t_i^{(L)},t)+\sqrt{L}\mu_i(t-t_i^{(L)}).\end{aligned} \tag{63}$$

Let $\tau = t - t_i^{(L)}$. Then $0 \leqslant \tau \leqslant 1$, and

$$\sqrt{L}\mu_i(t - t_i^{(L)}) = \sqrt{L}\mu_i\tau, \quad \hat{D}_i^{(L)}(t_i^{(L)}, t) \stackrel{d}{=} \hat{D}_i^{(L)}(1, 1+\tau),$$

and for all $l = 1, \ldots, m$,

$$\hat{D}_l^{(L)}(t_i^{(L)}, t) - \hat{D}_l^{(L)}(t_i^{(L)} - 1, t-1) \stackrel{d}{=} \hat{D}_l^{(L)}(1, 1+\tau) - \hat{D}_l^{(L)}(0, \tau).$$

Thus applying (62), (63), and above equivalences in distribution,

$$E\big[(\tilde{B}_i^{(L)*}(t) - \tilde{B}_i^{(L)*}(t_i^{(L)})) - (\tilde{B}_i^{(L)}(t) - \tilde{B}_i^{(L)}(t_i^{(L)}))\big] \leqslant E\Big[\sup_{0\leqslant\tau\leqslant1} \mathscr{V}^{(L)+}(\tau)\Big], \tag{64}$$

where

$$\mathscr{V}^{(L)}(\tau) \equiv g\sum_{l=1}^m \big|\hat{D}_l^{(L)}(1, 1+\tau) - \hat{D}_l^{(L)}(0, \tau)\big| + \big|\hat{D}_i^{(L)}(1, 1+\tau)\big| - \sqrt{L}\mu_i\tau.$$

Let $\kappa = \mu_i/(2mg+1)$ and define

$$u_l^{(L)}(\tau) = |\hat{D}_l^{(L)}(0, \tau)| - \sqrt{L}\tau\kappa \quad \text{and}$$
$$v_l^{(L)}(\tau) = |\hat{D}_l^{(L)}(1, 1+\tau)| - \sqrt{L}\tau\kappa, \quad 1 \leqslant l \leqslant m.$$

Then

$$\mathscr{V}^{(L)}(\tau) \leqslant g\sum_{l=1}^m u_l^{(L)}(\tau) + g\sum_{l=1}^m v_l^{(L)}(\tau) + v_i^{(L)}(\tau),$$

and thus

$$E\Big[\sup_{0\leqslant\tau\leqslant1} \mathscr{V}^{(L)+}(\tau)\Big] \leqslant g\sum_{l=1}^m E\Big[\sup_{0\leqslant\tau\leqslant1} u_l^{(L)+}(\tau)\Big] + (g+1)\sum_{l=1}^m E\Big[\sup_{0\leqslant\tau\leqslant1} v_l^{(L)+}(\tau)\Big]. \tag{65}$$

Since $0 \leqslant \tau \leqslant 1$,

$$\begin{aligned} E\Big[\sup_{0\leqslant\tau\leqslant1} u_l^{(L)+}(\tau)\Big] &= E\Big[\sup_{0\leqslant\tau\leqslant1} \big(|\hat{D}_l^{(L)+}(0, \tau)| - \sqrt{L}\tau\kappa\big)^+\Big] \\ &\leqslant E\Big[\sup_{0\leqslant\tau\leqslant L^{-1/4}} |\hat{D}_l^{(L)}(0, \tau)|\Big] + E\Big[\sup_{L^{-1/4}\leqslant\tau\leqslant1} \big(|\hat{D}_l^{(L)}(0, \tau)| - \sqrt{L}\tau\kappa\big)^+\Big] \\ &\leqslant (1 + \sigma_{ll}^2)L^{-1/8} + \frac{\sigma_{ll}^2}{\kappa}L^{-1/4}, \quad 1 \leqslant l \leqslant m, \end{aligned} \tag{66}$$

where the last inequality is a direct application of Lemma 3. Since $u_l^{(L)}(\tau) \stackrel{d}{=} v_l^{(L)}(\tau)$ $(1 \leqslant l \leqslant m,\ 0 \leqslant \tau \leqslant 1)$, we conclude from (60), (61), (64), (65), and (66) that

$$E\big[(\tilde{B}_i^{(L)*}(t) - \tilde{B}_i^{(L)}(t))\mathbf{1}(t_i^{(L)} \geqslant t-1)\big] \leqslant \theta_1 L^{-\delta/(2(2+\delta))} + \theta_2 L^{-1/8} + \theta_3 L^{-1/4},$$

where θ_1, θ_2 and θ_3 are constants. This inequality and (59) prove (54), and thus the theorem.

5. Conclusion

We have developed a family of inventory control policies, applicable to both continuous and discrete review, for ATO systems with a general BOM and identical lead times. We proved that our approach is asymptotically optimal on the diffusion scale, and thus the percentage difference of the long-run average expected inventory cost from its optimal value converges to zero as the lead time increases. Since inventory costs of ATO systems increase with the lead time, our policies are most advantageous in parameter regions where optimizing inventory control is most desirable.

Observe that the lower bound in Theorem 1 is proved in Doğru et al. (2010) by showing that $\underline{C}$ is a lower bound on the expected total inventory cost at every single point of time. Here in §4 we have proved that under our inventory policies, the expected inventory cost uniformly converges to $\underline{C}$ over an infinite time horizon. Although it may thus seem that our approach can be applied to prove asymptotic optimality of our policies in a setting with discounting, this does not follow immediately. An extension in this direction would clearly be of interest.

To put our work into a broader context, we present complementary numerical illustrations on the performance of our policies, a discussion on challenging issues in extending our approach to systems with nonidentical lead times, and a brief reflection on the general use of asymptotic analysis for inventory control problems.

5.1. Numerical Studies

It is easy to see from their definitions that our base-stock replenishment policy and allocation principle are generally implementable regardless of system scale. In addition to the above proof of asymptotic optimality, we can strengthen the support for these policies by showing that they also perform well outside the asymptotic regime. This task has been carried out by extensive numerical evaluations in Doğru et al. (2010, 2014), which address the aforementioned representative cases, W and M systems, respectively. Results show that our approach holds a commanding lead over alternative policies.

Below we present more numerical results from experiments on the M system in Doğru et al. (2014). The purpose of the latter is to thoroughly evaluate and "stress-test" our approach. Thus the main benchmark used for comparison there is a priority policy (LSP), which is a more efficient implementation of the NHB principle than the FRFS in the literature (Lu et al. 2010). Here our goal is to illustrate that our development in this paper represents major progress over common approaches in the literature. Thus we focus on comparisons with common policies in published papers, FIFO with commitment in Lu and Song (2005) and FRFS in Lu et al. (2010). Although both are used in Doğru et al. (2014) for comparison, neither is a focus of discussion there.

Following Doğru et al. (2014), we refer to our approach as SP for its dependence on stochastic programs, and

Figure 2. Histograms of differences between optimality gaps.

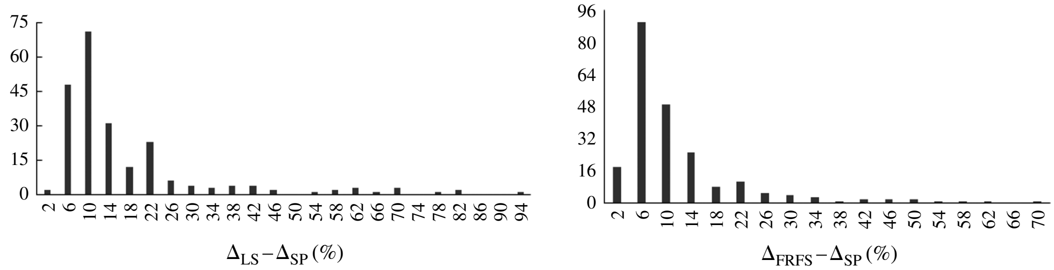

the FIFO policy in Lu and Song (2005) as LS. LS follows a base-stock replenishment policy with base-stock levels optimized for FIFO allocation (which includes component commitment). In contrast, FRFS serves demands FIFO without component commitment. In Lu et al. (2010), FRFS is applied with a base-stock replenishment policy, but how to determine an appropriate base-stock level is not addressed there. We consider FRFS to be a variation of LS and continue to use the method in Lu and Song (2005) to set base-stock levels. As we have shown in §3.2, in the M system, our allocation approach specializes to either a priority or a reservation policy. We use (18) with $\gamma = 0$ to set base-stock levels. The comparison is carried out for 224 sample cases outside the asymptotic regime. The design of these cases are given in Doğru et al. (2014, §4.2).

Performance of different policies are evaluated by the percentage difference of the resulting average cost in (9) from its lower bound in (17)

$$\Delta_p = 100\frac{\mathscr{C}^p - C^*}{C^*}, \quad p = \text{LS, FRFS, SP},$$

where C^* and C^{LS} are calculated directly and C^{FRFS} and C^{SP} are estimated by simulations. Figure 2 presents histograms of $\Delta_{LS} - \Delta_{\text{SP}}$ and $\Delta_{\text{FRFS}} - \Delta_{\text{SP}}$, which are the cost differences between our approach and the two alternatives, scaled by the lower bound. All bins have the same size of 4% and their midpoints are labeled on the horizontal axis. The midpoint of the tallest (third) bar in the left histogram is 10%, indicating that in the most frequent occurrences, SP cuts the optimality gap of LS by 8%–12%. The right histogram shows that $[8\%, 12\%)$ is also the second most likely range of reductions of the optimality gap of SP over FRFS. All bins are in the positive range, indicating that SP performs better in all of the 224 sample cases. In extreme cases, SP can reduce the optimality gap by 90% over LS and 70% over FRFS.

To help with understanding, we pick two sample cases for a more detailed comparison between SP and FRFS, the better of the two alternatives. Case 1 is picked from the range of the likely cases and case 2 is an extreme case. Relevant parameters are given at the bottom of Figure 3. The figure breaks down the inventory costs into different categories, where hc_j is the average expected holding cost of component j $(j = 1, 2)$ and bc_i is the average backlog cost of demand i $(i = 0, 1, 2)$.

In case 1, base-stock levels are $(45, 46)$ under FRFS and $(43, 43)$ under SP, so Figure 3 shows significant savings in inventory holding cost by using SP. As a trade-off, backlogs may go up, but our allocation principle directs the increase to demands of less value. Here the backlog cost of demand 0 is 8.125, 30% higher than the sum of backlog costs of the other two demands $(3.75 + 2.5)$. As we discussed in §3.2.2, for this case, the principle specializes to a reservation policy that not only gives priority to demand 0 but also withholds components from other demands. As a consequence, Figure 3 shows the backlog cost of demand 0 decreases relative to FRFS, and as a result, the increase of the total backlog cost under SP is less than the inventory

Figure 3. Breakdown of inventory costs in two cases.

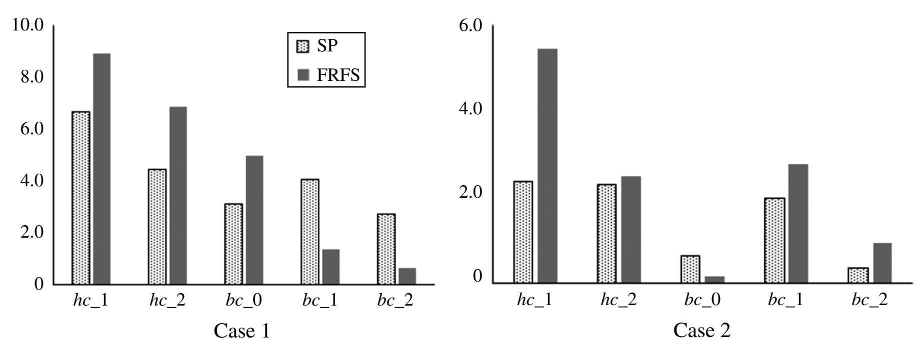

$(h_1 = 1.5, h_2 = 1.0, b_0 = 8.125, b_1 = 3.75, b_2 = 2.5)$ $(h_1 = 1.5, h_2 = 1.0, b_0 = 0.07, b_1 = 3.75, b_2 = 1.6)$

holding cost improvement, reducing the optimality gap by about 9% ($C^* = 19.1$, $C^{\mathrm{FRFS}} = 22.77$, and $C^{\mathrm{SP}} = 21.0$).

In case 2, base-stock levels are $(31, 22)$ under SP and $(41, 30)$ under FRFS. Despite much lower base-stock levels, which saves inventory holding costs, the total backlog cost also falls under SP, which is attributable to our allocation principle. The backlog cost of demand 0 in this case is trivial in comparison with other costs, so our allocation principle specializes to a priority policy that favors demands 1 and 2. As Figure 3 shows, because of this prioritization, demand 0 bears all the burden of reduced base-stock levels whereas backlog costs of demands 1 and 2 are actually reduced, leading to a fall in the total backlog cost. Correspondingly, the optimality gap of SP is 24%, much lower than that of FRFS (93%) ($C^* = 6.12$, $C^{\mathrm{FRFS}} = 11.79$, and $C^{\mathrm{SP}} = 7.59$).

5.2. Systems with Nonidentical Lead Times

Our allocation policy depends only on component shortage and inventory costs, so it can be directly combined with a base-stock replenishment policy as a complete solution for systems with general lead times. However the analysis in Rosling (1989) and Reiman and Wang (2012) suggest a base-stock policy is highly unlikely to be asymptotically optimal in this setting. Thus we need to explore new replenishment policies.

We conjecture that a similar SP-based approach can also lead to asymptotically optimal policies for general systems with nonidentical lead times. For these cases, Reiman and Wang (2012) has completed the first step of the four-step process introduced in §1. They show that the optimal solution of a particular $K+1$ stage SP is a lower bound on the average inventory cost of ATO systems with $K \geqslant 1$ different replenishment lead times. Although a replenishment policy, based on the solution of the $K+1$ stage SP is suggested in Reiman and Wang (2012), the analysis of this policy and the proof of its asymptotic optimality is challenging. It seems clear, based on the optimality of this replenishment policy in certain contexts (including that of Rosling 1989, where it reduces to his), that when lead times are not identical, a sensible replenishment policy should set inventory positions for components with shorter lead times based on the availability of components with longer lead times. Such a policy would in general not be of base-stock form, making the analysis much harder than the one in this paper, which relies on properties of base-stock policies. In particular, $\mathbf{Q}(t)$ has a simple expression in terms of the demand process under a base-stock policy (cf. (27)) that would no longer hold if the replenishment policy were not base stock. We leave the task of dealing with this significant issue to future work. (Note that we would also need to prove suitably generalized versions of Theorems 2 and 3, along with Lemmas 4 and 6.) On the other hand, our target-based allocation policy should be transferable to systems with nonidentical lead times.

Besides proving asymptotic optimality, it is important to make our approach computationally feasible for "industrial size" problems. Although two-stage SPs are sufficient for systems with identical lead times, the required number of stages increases with the number of distinct lead times, making the SP harder to solve. Our future research will address this issue to make the solution tractable.

5.3. Asymptotic Analysis for Inventory Problems

Inventory management has many difficult problems for which exactly optimal control policies are not known despite years of effort. Heuristics like FIFO allocation and its variations are developed out of consideration for the ability to perform policy evaluation with little justification from the optimality perspective. It helps to observe that in the related area of stochastic processing networks, asymptotic analysis serves as a powerful tool that has gained tremendous ground on many difficult problems. Plambeck and Ward (2006) demonstrate that the same success can be replicated in optimizing ATO production/inventory systems by applying big-step review policies and solving Brownian control problems. In this paper, we make dual use of asymptotic analysis to minimize the long-run average cost of ATO inventory systems. As a confirmatory tool, the analysis provides theoretical support to the existing use of base-stock policy for systems with identical lead times. As an exploratory mechanism, the analysis leads us to the allocation principle that spawns effective allocation policies.

The confirmatory role of asymptotic analysis applies to other inventory problems. For instance, this paper focuses on the backlog model. In cases when customers do not have patience to wait, one needs to consider a lost sales model. In some related ATO systems, such as the Markovian production/inventory model, similar control strategies apply to both cases (e.g., Benjaafar and ElHafsi 2006). However in our problem, especially in the long lead time asymptotic regime, the appropriate replenishment policies for the two models are qualitatively different. In the backlog model, because customers leave only after their demands are satisfied, if the inventory position is lower currently, more components will be needed one lead time later, so more components need to be ordered, which is achieved under a base-stock policy. In the lost sales model, a lower inventory position at the moment only means that more demands will not be satisfied during the next lead time. These unsatisfied demands will be gone when the newly ordered components arrive, so there is not much reason to link the order quantity with the current inventory position. This insight motivates an open-loop policy, which is to order a fixed amount of components for every fixed time interval (Reiman 2004). Despite its simplicity, the policy is found to outperform more sophisticated approaches (Zipkin 2008), and its effectiveness is rigorously proven by Goldberg et al. (2014), who show that in the long lead time asymptotic regime, the policy is asymptotically optimal without any scaling of the cost. Nevertheless, their analysis is on the single-product system, and how to address multiproduct systems with lost sales remains a challenge.

Asymptotic analysis can also play an exploratory role to untangle difficult issues of our problems in broader settings. For instance, Plambeck and Ward (2008) consider the minimization of the discounted cost in ATO systems with capacitated component production and fixed transportation cost. They show that in the high-volume asymptotic regime, it is optimal to use economic order quantity to determine the shipment size, subject to the capacity constraint at the production facility. It would be interesting to examine if similar policies/analysis would apply to minimizing the long-run average cost in ATO systems with fixed cost.

Our analysis and planned extension described in §5.2 are for deterministic lead times. It is not clear how to extend our approach to stochastic lead times, or even if it can be extended. Nonetheless, this is an interesting topic for further exploration.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/opre.2015.1372.

Acknowledgments

The authors thank Jong-Shi Pang for directing them to Han et al. (2012, Theorem 4.1). The authors also thank Mustafa Doğru for his help with numerical studies on the M system in §5.1. This paper is based upon work supported by the National Science Foundation [Grant CMMI 1363314].

References

Agrawal N, Cohen M (2001) Optimal material control in an assembly system with component commonality. *Naval Res. Logist. (NRL)* 48(5):409–429.

Akçay Y, Xu S (2004) Joint inventory replenishment and component allocation optimization in an assemble-to-order system. *Management Sci.* 50(1):99–116.

Aubin J, Frankowska H (2008) *Set-Valued Analysis* (Birkhäuser, Boston).

Benjaafar S, ElHafsi M (2006) Production and inventory control of a single product assemble-to-order system with multiple customer classes. *Management Sci.* 52(12):1896–1912.

Doğru M, Reiman M, Wang Q (2010) A stochastic programming based inventory policy for assemble-to-order systems with application to the W model. *Oper. Res.* 58(4):849–864.

Doğru M, Reiman M, Wang Q (2014) Stochastic programming based inventory control for the M system. Working paper, Alcatel-Lucent Bell Labs, Murray Hill.

Goldberg DA, Katz-Rogozhnikov DA, Lu Y, Sharma M, Squillante MS (2014) Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. Preprint, arXiv:1211.4063v2.

Han L, Camlibel MK, Pang J-S, Heemels WMH (2012) A unified numerical scheme for linear-quadratic optimal control problems with joint control and state constraints. *Optim. Methods and Software* 27(4–5):761–799.

Harrison J (1988) Brownian models of queueing networks with heterogeneous customer populations. Fleming W, Lions P-L, eds. *Stochastic Differential Systems, Stochastic Control Theory and Applications*, Vol. 10 (Springer-Verlag, New York), 147–186.

Harrison J (1996) The BIGSTEP approach to flow management in stochastic processing networks. Kelly FP, Ziedins I, Zachary S, eds. *Stochastic Networks: Theory and Applications*, Vol. 4 (Oxford University Press, Cambridge, UK), 147–186.

Harrison J, López M (1999) Heavy traffic resource pooling in parallel-server systems. *Queueing Systems* 33(4):339–368.

Harrison J, Van Mieghem J (1999) Multi-resource investment strategies: Operational hedging under demand uncertainty. *Eur. J. Oper. Res.* 113(1):17–29.

Harrison J, Wein L (1990) Scheduling networks of queues: Heavy traffic analysis of a two-station closed network. *Oper. Res.* 38(6):1052–1064.

Hausman W, Lee H, Zhang A (1998) Joint demand fulfillment probability in a multi-item inventory system with independent order-up-to policies. *Eur. J. Oper. Res.* 109(3):646–659.

Hoffman AJ (1952) On approximate solutions of systems of linear inequalities. *J. Res. National Bureau Standards* 49(4):263–265.

Karlin S, Scarf H (1958) Inventory models of the Arrow-Harris-Marschak type with time lag. *Stud. Math. Theory Inventory Production* 1:155–178.

Lu L, Song J, Zhang H (2014) Optimal and asymptotically optimal policies for assemble-to-order N- and W-Systems. Columbia Business School Research Paper No. 14-55.

Lu Y, Song J (2005) Order-based cost optimization in assemble-to-order systems. *Oper. Res.* 53(1):151–169.

Lu Y, Song J, Yao D (2003) Order fill rate, leadtime variability, and advance demand information in an assemble-to-order system. *Oper. Res.* 51(2):292–308.

Lu Y, Song J, Zhao Y (2010) No-holdback allocation rules for continuous-time assemble-to-order systems. *Oper. Res.* 58(3):691–705.

Nadar E, Akan M, Scheller-Wolf A (2014) Technical note—Optimal structural results for assemble-to-order generalized M-systems. *Oper. Res.* 62(3):571–579.

Plambeck E, Ward A (2006) Optimal control of a high-volume assemble-to-order system. *Math. Oper. Res.* 31(3):453–477.

Plambeck E, Ward A (2008) Optimal control of a high-volume assemble-to-order system with maximum leadtime quotation and expediting. *Queueing Systems* 60(1–2):1–69.

Reiman M, Wang Q (2012) A stochastic program based lower bound for assemble-to-order inventory systems. *Oper. Res. Lett.* 40(2):89–95.

Reiman MI (2004) A new and simple policy for the continuous review lost sales inventory model. Working paper, Alcatel-Lucent Bell Labs, Murray Hill.

Rosling K (1989) Optimal inventory policies for assembly systems under random demands. *Oper. Res.* 37(4):565–579.

Schrijver A (1998) *Theory of Linear and Integer Programming* (John Wiley & Sons, New York).

Song J, Zipkin P (2003) Supply chain operations: Assemble-to-order systems. Graves S, de Kok T, eds. *Handbooks in Operations Research and Management Science* (Elesvier, Amsterdam), 561–596.

van Jaarsveld W, Scheller-Wolf A (2015) Optimization of industrial-scale assemble-to-order sytems. *INFORMS J. Comput.* Forthcoming.

Zhang A (1997) Demand fulfillment rates in an assemble-to-order system with multiple products and dependent demands. *Production Oper. Management* 6(3):309–324.

Zipkin P (2008) On the structure of lost-sales inventory models. *Oper. Res.* 56(4):937–944.

Martin I. Reiman is a distinguished member of the technical staff in the Mathematics of Networks and Systems Research department of Alcatel-Lucent Bell Labs in Murray Hill, NJ. His research has focused on the analysis, optimization and control of various stochastic service systems, with an emphasis on fluid and diffusion limits. He is an INFORMS Fellow and has served as chair of the Applied Probability Society of INFORMS.

Qiong Wang is an associate professor in the Department of Industrial and Enterprise Systems Engineering at the University of Illinois at Urbana–Champaign. Before joining UIUC, he was a member of the technical staff at Alcatel-Lucent Bell Labs in Murray Hill, New Jersey. His research focuses on inventory control, revenue management, and pricing, as well as public policy analysis.