Asymptotic Optimality of Base-Stock Policies with Idle Times for Stochastic Economic Lot Scheduling Problems

Yue Hu
Decision, Risk, and Operations, Columbia Business School, yhu22@gsb.columbia.edu

Jing Dong
Decision, Risk, and Operations, Columbia Business School, jing.dong@gsb.columbia.edu

Ohad Perry
Department of Industrial Engineering and Management Sciences, Northwestern University, ohad.perry@northwestern.edu

We study the Stochastic Economic Lot Scheduling Problem (SELSP) which is concerned with the scheduling of production of multiple products by a common facility. Our goal is to characterize an optimal lot-sizing policy, with possible idle times inserted between productions, so as to minimize a long-run average holding, backlog, and setup cost. Due to its complexity, the optimal-control problem of SELSP is unsolved, with existing results mostly focusing on restricted classes of policies. In this paper, we take an asymptotic approach and prove that the classical base-stock policy with idle times is asymptotically optimal among a large class of admissible controls as the setup times grow large. Extensive numerical studies and computations of the optimal control in sufficiently simple cases (that can be solved using Markov decision processes) demonstrate the effectiveness and the robustness of the asymptotically optimal policy.

Key words: Multi-Product Inventory System, Setup Time, Idling, Base-Stock Policy, Optimal Control, Asymptotic Optimality

1. Introduction

The Stochastic Economic Lot Scheduling Problem (SELSP) considers the production of multiple products on a single machine with random demands, random production times, and random setup times. In this problem, the system manager aims to find a "good" production plan that specifies for each state of the system whether to continue the production of the current product, to switch to another product, or to idle the facility. The goal is to optimally balance inventory holding, order backlogging, and production setup costs. The SELSP has been widely applied to model and optimize large-scale manufacturing systems (see, for example, Gascon et al. (1994), Sox and Muckstadt (1997), Grasman et al. (2008)). The key features it captures are the setup time and cost incurred when switching from producing one type of products to another, the uncertainty in demand and production, and the heterogeneity of different products (including demands, production times, holding and backlog costs, etc).

From a theoretical perspective, the SELSP has long been identified as a challenging problem. A production plan for a multi-product manufacturing facility needs to specify two critical components:

a production sequence and a lot sizing policy. The production sequence specifies which product the facility turns to once it stops producing the current product, and can be either fixed or dynamic. The lot sizing policy decides how long the facility produces the current product, possibly with idle times inserted between the completion of the current production batch and the setup for the next product. Jointly optimizing the production sequence and the lot sizing policy seems out of reach: Hsu (1983) shows that, even a restricted version of the *deterministic* Economic Lot Scheduling Problem (ELSP) is NP-hard. Thus, most of the existing works fix the production sequence and study the lot sizing policy (see, for example, Federgruen and Katalan (1996a,b), Krieg and Kuhn (2004)), which is the focus of our paper as well.

An important class of lot sizing policies is the base-stock policies, under which the facility continues producing a product until its inventory level reaches a pre-specified base-stock level, with idle times inserted between the completion of the current production batch and the setup for the next product. It is argued in Federgruen and Katalan (1996b) that the base-stock policy is effective because (i) it is easy to implement and monitor, and (ii) it can be seen as a variant of the common cycle policy for deterministic ELSPs and the (s, S)-rule (i.e., periodic review order-up-to policy) for inventory replenishment systems, both of which are shown to be effective in coordinating replenishments across products. Although the simple and natural base-stock policy has been prevalently studied and applied, its merits have only been advocated qualitatively; in particular, there exists no optimality result for this policy; see, e.g., the survey in Winands et al. (2011). This fact is not surprising given the complexity of SELSP, which renders this problem intractable.

Thus, in this paper we take the prevalent approach in the queueing-network literature, and prove that the base-stock policy is *asymptotically optimal* among a large class of policies. Our asymptotic results are established under a large setup-times regime, and are therefore appropriate when the production times of an individual product are small relative to the setup times, as is often the case in practice. It is significant that, even if one restricts oneself to the base-stock policy, deriving the optimal control parameters can be computationally intensive, especially when optimization over the idle times is involved; see (Federgruen and Katalan 1996b). However, the *asymptotically optimal* control parameters can be derived relatively easily by solving a simple deterministic optimization problem which, for some important special cases, can be computed explicitly. Finally, we demonstrate through extensive numerical experiments that the policies derived based on our asymptotic analysis are effective.

Our analysis provides both theoretical insights and prescriptive solutions to the management of multi-product inventory systems with large setup times. Our main contributions can be summarized as follows. We leverage a novel theoretical framework developed in Hu et al. (2020a) to establish the asymptotic optimality of the base-stock policy with idle times in a large setup-times fluid regime. In particular, we first consider a deterministic relaxation for the stochastic network (namely, a fluid model), that is represented as a hybrid dynamical system (HDS). Employing HDS theory, we prove that the *common cycle approach* for the deterministic ESLP (Hanssmann 1962, Maxwell 1964), which optimizes the control over a single production cycle, is equivalent to optimizing the control over an arbitrary number of production cycles. Further, we prove that, for any initial condition, the fluid model converges to the optimal trajectory under our proposed optimal fluid control. Finally, we "translate" the optimal fluid control to a base-stock policy with idle times for the corresponding stochastic system, and prove that this policy is asymptotically optimal, in a large-setup-time asymptotic regime, among a large class of admissible controls. We mention that, to the best of our knowledge, all of the policies studied in the literature belong to our family of admissible controls; examples include various variations of base-stock policies and fixed-quantity policies.

1.1. Related Literature

In this section, we provide a brief review of the literature that is most related to our work. We refer to the survey papers Elmaghraby (1978), Allahverdi et al. (1999) for comprehensive reviews regarding ELSP, and to Sox et al. (1999), Winands et al. (2011) for reviews on SELSPs.

Deterministic Systems. The ELSPs are deterministic versions of the SELSPs, in which the production and arrival streams, and the setup times are continuous deterministic functions. The objective is to find a production schedule that minimizes the long-run average holding, backlog, and setup costs. Even in this deterministic setting, jointly optimizing the production sequence and the lot sizing is prohibitive. To the best of our knowledge, the optimal solution has only been obtained in a special case considered in Jones and Inman (1989). Since the general form of the problem is too hard to solve, most existing works take the approach of finding the optimum of a restricted version of the original problem. One of the most noteworthy confinement schemes is the common cycle (or rotation cycle) approach (Hanssmann 1962, Maxwell 1964). This method optimizes over periodic production schedules in which each item is produced exactly once in a production cycle. The varying lot size approach (Gallego and Roundy 1988, Gallego 1990, 1994, Zipkin 1991, Dobson 1992) generalizes the common cycle approach by allowing some products to be produced more than once in a cycle. Other works combine the optimization of the lot-sizing policy with dynamic programming formulations (Bomberger 1966, Elmaghraby 1978, Adelman and Barz 2014a,b). Linear costs are assumed in all the aforementioned works, and backlog is permitted only in Gallego and Roundy (1988), Gallego (1990, 1994).

Our work optimizes the lot sizing policy by first optimizing the system-level process over one production cycle using the common cycle approach, and then showing that the base-stock policy, with properly chosen parameters, can achieve that optimal one-cycle behavior in the long run. Our work contributes to the ELSP literature in three aspects: First, we consider general (non-negative, non-decreasing, and continuous) cost functions that grow at most at a polynomial rate (see Assumption 2); second, we prove that it is sufficient to optimize over one (instead of multiple) production cycle; third, our work draws rigorous connection between the ELSP and SELSP via fluid limits in a large setup-times asymptotic regime.

The SELSP with setup times. Most of the papers concerning the SELSP with setup costs or times deal with heuristics. Some well-known ones include Gallego (1990), which derives a target cyclic production schedule based on the deterministic ELSP with an added safety stock to hedge against demand uncertainty. Bourland and Yano (1994) assume a reorder point for each individual product; once the next product hits its reorder point, the facility immediately moves on to the next product, and the rest of the production quantity for the current product is made on overtime. Gascon et al. (1994) compare six different heuristics with the objective of minimizing the sum of setup and inventory costs. It is significant that all these heuristics optimize over a specific class of policies.

More recently, systems with large setup times have attracted great attention due to their practical relevance. Focusing on systems without inventories, Van der Mei (1999), Olsen (2001), Winands (2007, 2011) characterize the limiting system performance (i.e., steady-state waiting time of backlogged orders) in the large-setup-time limit under a particular control. Our work also considers the large-setup-time scaling, but we focus on proving asymptotic optimality of our control under a general class of admissible controls, as opposed to approximating performance measures under a specific policy. Lastly, similar large-time scaling has been considered for inventory control problems other than the SELSP. For example, Xin and Goldberg (2018) prove that a tailored base-surge policy is asymptotically optimal in dual-sourcing inventory systems, when the lead time difference between the regular and express suppliers grows to infinity.

Queueing-based approaches to the SELSP. There is a strong connection between the SELSP and the optimal control problem for polling systems (i.e., stochastic queueing networks where several queues are served by a single server). The switchover times in polling systems can be considered as the setup times in the SELSP, and the queue length can be equated to the backlog level. Though the SELSP is often considered more challenging than the control problem of polling systems due to the inclusion of inventory, possible idling behavior, and more complicated cost structure (i.e., the inventory level can be both positive or negative with different cost rates for the positive and negative parts), results for polling systems can shed light to SELSPs in certain settings.

Federgruen and Katalan (1996b) apply analysis of the queue length process under the exhaustive policy in Federgruen and Katalan (1994) to characterize the inventory level process under the base-stock policy. In addition, leveraging insights from the heavy-traffic approximation of polling systems in Coffman Jr et al. (1995, 1998), Markowitz et al. (2000) study two versions of the SELSP with linear cost rates, one with setup costs and the other with setup times, and solve a diffusion control problem for the two-product SELSPs. While the problem can be explicitly solved for the case with positive setup cost, only heuristics are obtained for the case with positive setup times. Markowitz and Wein (2001) further generalize the heavy-traffic approximation to other related inventory problems. Our work here also leverages results developed for polling systems, and in particular, the mathematical framework developed in our recent paper Hu et al. (2020a), although the setting and the control we consider here differ from those in this latter reference.

1.2. Notation

All the random variables and processes are defined on a single probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We denote $\mathbb{E}$ as the expectation operator. We let $\mathbb{R}$, $\mathbb{Z}$ and $\mathbb{N}$ denote the sets of real numbers, integers and strictly positive integers, respectively, $\mathbb{Z}_+ := \mathbb{N} \cup \{0\}$, and $\mathbb{R}_+ := [0, \infty)$. For $K \in \mathbb{N}$, we let $\mathbb{R}^K$ denote the space of K-dimensional vectors with real components, and denote these vectors with bold letters and numbers. We let D^K denote the space of right-continuous $\mathbb{R}^K$-valued functions (on arbitrary finite time intervals) with limits everywhere, endowed with the usual Skorokhod J_1 topology; see Chapter 11 of Whitt (2002). Let $D := D^1$. We use C^K (and $C := C^1$) to denote the subspace of D^K of continuous functions. It is well-known that the J_1 topology relativized to C^K coincides with the uniform topology on C^K, which is induced by the norm $||x||_t := \sup_{0 \le u \le t} \|x(u)\|$, where $||x||$ denotes the usual Euclidean norm of $x \in \mathbb{R}^K$.

1.3. Organization

The rest of the paper is organized as follows. In Section 2 we introduce the model, state the problem formulation, and preview the main results. In Section 3 we studies a deterministic relaxation (a fluid model) to the SELSP, which is characterized as a hybrid dynamical system (HDS). This deterministic relaxation gives rise to a fluid control problem (FCP). We show that the base-stock policy is optimal for the FCP. In Section 4 we translate the fluid-optimal base-stock policy for the stochastic system, and prove that it is asymptotically optimal as the setup times increase without bound. Sections 5 and 6 are devoted to numerical experiments. In Section 5, we validate the effectiveness of the fluid-translated policy for the stochastic system. In Section 6, we compare the performance of the proposed base-stock policy to that of the "exact" optimal policy derived based on the corresponding Markov decision process (MDP). We conclude in Section 7. All the proofs appear in the appendix.

2. Model Description, Problem Statement, and Main Results

We consider a make-to-stock inventory system with K products numbered $1,...,K$. Let $\mathcal{K} := \{1,...,K\}$. Demands for product $k \in \mathcal{K}$ arrive to the system according to a Poisson process with rate $\lambda_k > 0$, independently of everything else. When a demand arrives, it is immediately satisfied by one unit[1] if an inventory of the specific product exists. Otherwise, the demand is backlogged in an infinite buffer. Production times for individual units are independent; those of product k are identically distributed and denoted by a generic random variable G_k with mean $1/\mu_k < \infty$. The utilization rate for product k is $\rho_k := \lambda_k/\mu_k$, and for the system is $\rho := \sum_{k\in\mathcal{K}} \rho_k$. We assume $\rho < 1$, which is the necessary condition for the system to be stable (Fricker and Jaibi 1994) in the sense that there exists a control under which the system has a steady state. (However, $\rho < 1$ does not guarantee that the system is stable; this depends also on the control that is exercised.)

A single facility rotates between the products in a fixed order of $(1,...,K)$, switching back to product 1 once it is done producing product K. (The products are numbered in the order at which they are produced.) When switching from the production of product $k-1$ to the production of product k, a setup time is incurred. We let S_k denote a generic random variable representing the setup time incurred when switching from product $k-1$ to k, and let $s_k < \infty$ denote its mean, where we define $1-1 := K$ and $K+1 := 1$. We denote the total expected setup time during a cycle by $s := \sum_{k\in\mathcal{K}} s_k$, and assume that $s > 0$.

After setting up for product k, the facility produces this product according to a lot sizing policy. The lot sizing policy determines how many units should be produced in the current production run, and is one of the two control levers in our study. Once the production of product k is terminated, a

[1] Throughout the paper, we refer to each product type as "product," and an individual unit of a product type as "unit."

deterministic, fixed idle time $W_k \in \mathbb{R}_+$ is inserted before the facility sets up for producing product $k+1$. Let $\mathbf{W} := (W_k, k \in \mathcal{K})$. Whether and for how long to idle is the other part of the control. As we shall see below, the insertion of idle times may be beneficial, because it reduces the frequency of switchings in the long run, and thus reduces long-run average setup costs. Figure 1 illustrates the model configuration for a three-product facility.

Figure 1 A three-product facility

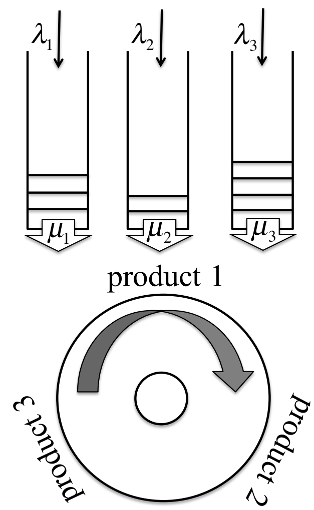

We refer to the start of the production of product k (after the setup time) as the *polling* epoch of product k, and the termination of production of this product as the *departure* epoch (before the potential idling time). A production *cycle* is the time elapsed between two consecutive polling epochs of product 1.

Without loss of generality, we assume that the facility is at the polling epoch of product 1 at time 0. For $m \geq 1$, we denote the length of the mth production cycle by $T^{(m)}$, so that $U^{(m)} := \sum_{i=1}^{m} T^{(i)}$ (with $U^{(0)} := 0$) is the beginning of the $(m+1)$st cycle. Let $A_k^{(m)}$ and $D_k^{(m)}$, $k \in \mathcal{K}$, denote the polling and departure epochs of product k during the mth cycle. Then, $V_k^{(m)} := D_k^{(m)} - A_k^{(m)}$ is the *busy time* of product k, namely, the length of the production run of product k in the mth cycle. We do not rule out the possibility of having instantaneous busy times, so that it is possible to have $V_k^{(m)} = 0$ for some $k \in \mathcal{K}$ and $m \geq 1$, in which case the polling epoch is equal to the departure epoch of product k in this cycle.

We let $X_k(t)$ denote the inventory level of product k at time $t \geq 0$, and $X(t) := (X_k(t), k \in \mathcal{K})$. Then $\max(0, X_k(t))$ is the available inventory, and $\max(-X_k(t), 0)$ is the number of backlogged orders of product k at time t. We use $Z(t)$ to denote the location of the facility at time t, i.e., whether the facility is producing, idling, or setting up at time t. More specifically, $Z(t) = k^S$ means that the facility is setting up for product k, $Z(t)=k$ means that the facility is actively producing product k, and $Z(t)=k^I$ means that the facility is idling after terminating the production of product k at time t. Then $Z(t)\in\mathcal{Z}:=\{k^S,k,k^I:k\in\mathcal{K}\}$.

2.1. Admissible Policies

The control of the system is a specification of the lot sizing and idling behavior. To define the family of admissible controls, which we denote by Π, let $\mathcal{F}_t:=\{\mathcal{F}_t:t\geq 0\}$ be the σ-algebra generated by X. A control is said to be non-anticipative if its decisions are made based on $\mathcal{F}_t$ (and not on any future information). It is said to be (discrete) Markovian if, conditional on the inventory level at the polling epoch of a product, the number of units to produce is independent of the history up to that instant. In particular, given $X(A_k^{(m)})$, $k\in\mathcal{K}$, $m\geq 1$, the lot size for product k in the mth cycle is conditionally independent of $\mathcal{F}_{A_k^{(m)}}$.

Definition 1 (admissible control) *Each element $\pi\in\Pi$ specifies a lot sizing policy and an idle time vector $\mathbf{W}\in\mathbb{R}_+^K$. The idle time vector is deterministic and remains fixed over time (i.e., independent of the evolution of the inventory level). In contrast, the lot sizing policy can be state dependent, and satisfies the following conditions:*

(i) The policy is non-anticipative.

(ii) The policy is (discrete) Markovian.

Recall that for $m\geq 0$, $U^{(m)}$ is the beginning of the $(m+1)$st production cycle, namely, the polling epoch of product 1 in the m-th cycle. Let

$$\tilde{X}(m):=X(U^{(m)}),\quad \text{for } m\geq 0.$$

The following Lemma follows from Proposition 1 in Fricker and Jaibi (1994).

Lemma 1 *Under any admissible policy $\pi\in\Pi$, the process $\{\tilde{X}(m):m\geq 0\}$ is an aperiodic time-homogeneous DTMC.*

For the SELSP (formally defined below in (2.1)), we can restrict attention to admissible controls that are stable, as defined below.

Definition 2 *An admissible lot sizing policy $\pi\in\Pi$ is stable if the DTMC $\tilde{X}$ is absorbed in a positive recurrent class.*

The family of admissible controls Π is large, and includes (to the best of our knowledge) all of the policies that were considered in the literature; here we bring four examples. Consider K-dimensional vectors $\mathbf{B}=(B_1,...,B_K)$ and $\mathbf{M}=(M_1,...,M_K)$, where B_k is the base-stock level, and M_k is the production quantity of product k.

Base-stock policies: when the facility polls product $k\in\mathcal{K}$, it continues producing the product until the target inventory level B_k is reached. Product k is skipped if the inventory level at its polling epoch is smaller than the target level.

Gated base-stock policies: when the facility turns to product $k\in\mathcal{K}$, it produces a batch of size that is equal to the difference between the base-stock level and the inventory level at the polling epoch. The product is skipped if its inventory level at the polling epoch is smaller than the base-stock level.

Fixed-quantity policies: At each visit of product k, the facility produces a batch of size of M_k.

Quantity-limited base-stock policies: At each polling of product $k\in\mathcal{K}$, the facility continues producing the product until either the target inventory level B_k is reached, or a number of M_k of units has been produced, whichever occurs first. The product is skipped if its starting inventory level is no less than the target base-stock level.

2.2. The Cost Structure

The optimal-control problem we consider aims to minimize the long-run average total cost incurred by the system by exercising an admissible control. To describe the objective function, let

$h_k:\mathbb{R}_+\to\mathbb{R}_+$ denote the inventory holding-cost function for product k per unit time;

$p_k:\mathbb{R}_+\to\mathbb{R}_+$ denote the backlog cost function for product k per unit time;

$r_k\in\mathbb{R}_+$ denote the fixed setup cost incurred per setup for product k.

For $k\in\mathcal{K}$ and $t\geq 0$, and for $\Gamma_k(t)$ denoting the number of setups for product k by time t, the objective function we consider is the following long-run average cost

$$C_\pi:=\lim_{t\to\infty}\frac{1}{t}\sum_{k\in\mathcal{K}}\left(r_k\Gamma_{\pi,k}(t)+\int_0^t\left(h_k\left(X_{\pi,k}(s)^+\right)+p_k\left(X_{\pi,k}(s)^-\right)\right)ds\right),\tag{2.1}$$

where we use the subscript π to mark the dependence of the system dynamics on the control. Note that for the class of stable controls (Definition 2), the inventory level process is a regenerative process, and thus the limit in (2.1) exists w.p.1. In what follows, we shall drop the subscript π if there is no ambiguity regarding which control is considered.

In addition, we define functions $\psi_k:\mathbb{R}\to\mathbb{R}_+$ and $\psi:\mathbb{R}^K\to\mathbb{R}_+$, where for $x\in\mathbb{R}^K$,

$$\psi_k(x_k):=h_k\left(x_k^+\right)+p_k\left(x_k^-\right),\quad \text{and}\quad \psi(x):=\sum_{k\in\mathcal{K}}\psi_k(x_k),\quad k\in\mathcal{K}.$$

2.3. Large-Setup-Time Asymptotic

Since we require an admissible control to be discrete-Markov (Definition 1), we could in principle employ discrete-time Markov decision processes (MDP) embedded at the polling epochs to optimize the control. However, due to the complex dynamics, the transition probabilities and costs for the MDPs are prohibitively hard to characterize, so that, even a two-product system is hard to analyze—a problem that is exacerbated in large networks. See Section 6. Thus, instead of optimizing the control for the system, we search for an *asymptotically optimal* control, namely, a control that achieves the best possible asymptotic outcome. To this end, we consider a fluid-limit regime obtained by letting the setup times grow without bound; this limiting approximation is efficient for systems in which the setup times are substantially larger than the per-product production times.

For out asymptotic analysis, we consider a sequence of facilities indexed by $n \geq 1$, and mark the primitives and processes that scale with n with a superscript. We assume that the demand rate and production rate for product $k \in \mathcal{K}$ do not scale with n. Thus, we keep the notations λ_k and μ_k for the nth system. We let the setup times grow without bound, and assume that the setup time for product k in the nth system, S_k^n, satisfies

$$\bar{S}_k^n := S_k^n / n \Rightarrow s_k \quad \text{as } n \to \infty, \quad k \in \mathcal{K},$$

where we use "$\Rightarrow$" to denote convergence in distribution. Furthermore, the mean setup time satisfies $\mathbb{E}\left[\bar{S}_k^n\right] = s_k$, $k \in \mathcal{K}$, $n \geq 1$.

For $n \geq 1$, let the fluid-scaled inventory and the setup counting processes be defined as follows:

$$\bar{X}^n(t) := \left(X_k^n(nt)/n,\ k \in \mathcal{K}\right), \quad \text{and} \quad \bar{\Gamma}^n(t) := \Gamma^n(nt) \quad t \geq 0.$$

Note that because the setup times are $\Theta_p(n)$ in the nth system, there are $\Theta_p(1)$ number of setups in $O(n)$ amount of time.[1] We refer to the above scaling as *the large-setup-time scaling* (also referred to as the fluid scaling).

Then under the fluid scaling, the objective function for the nth system is given by

$$\bar{C}_{\pi^n}^n := \lim_{t \to \infty} \frac{1}{t} \left(\sum_{k \in \mathcal{K}} r_k \bar{\Gamma}_{\pi^n,k}^n(t) + \int_0^t \psi\left(\bar{X}_{\pi^n}^n(s)\right) ds \right), \tag{2.2}$$

where π^n is the control employed in system n, assumed to be admissible (and stable).

We next define the notion of asymptotic optimality.

[1] Given a sequence of random variables $\{X^n : n \geq 1\}$ and a sequence of non-negative real numbers $\{a^n : n \geq 1\}$, we write $X^n = \Theta_p(a^n)$ if $||X^n||/a^n$ is stochastically bounded, i.e., for any $\epsilon > 0$, there exist finite $M, N \in \mathbb{N}$ such that $\mathbb{P}\left(||X^n||/a^n > M\right) < \epsilon$ for all $n \geq N$.

Definition 3 *A sequence of admissible policies* $\{\tilde{\pi}_*^n : n \geq 1\}$ *is said to be asymptotically optimal if, for any other sequence of admissible policies* $\{\pi^n : n \geq 1\}$*, it holds that*

$$\limsup_{n\to\infty} \bar{C}^n_{\tilde{\pi}^n_*} \leq \liminf_{n\to\infty} \bar{C}^n_{\pi^n} \quad w.p.1.$$

2.4. Main Results

Since we consider general cost functions, we impose the following assumptions on the production and setup time distributions. Recall that G_k and S_k denote the production time and setup time corresponding to product k, respectively.

Assumption 1 *For all* $k \in \mathcal{K}$,

(i) $\mathbb{E}\left[e^{tG_k}\right]$ *for all* $t \in (-\epsilon, \epsilon)$*, for some* $\epsilon > 0$.

(ii) $\mathbb{E}\left[e^{tS_k^n}\right] < \infty$ *for all* $t \in \mathbb{R}_+$ *and* $n \geq 1$*. In addition,* $\mathbb{E}\left[(\bar{S}_k^n)^\ell\right] \to s_k^\ell$ *as* $n \to \infty$ *for any* $\ell \geq 1$.

Assumption 1 requires that the production times and setup times have light-tailed distributions; such distributions include the uniform, Erlang (including exponential), and Weibull distribution with shape parameter larger than or equal to 1. Moreover, note that under the large-setup-time scaling, we have $\bar{S}_k^n \Rightarrow s_k$ as $n \to \infty$ for all $k \in \mathcal{K}$. Then the convergence of means in Assumption 1 is equivalent, or implies uniform integrability condition for the setup times.

For $n \geq 1$, we use π^n_{BSI} to denote the control that employs the base-stock policy with base-stock level $\mathbf{B}^n$ and idle time vector $\mathbf{W}^n$, for some vectors $\mathbf{B}^n \in \mathbb{R}^K$ and $\mathbf{W}^n \in \mathbb{R}^K_+$. (The subscript "BSI" is mnemonic for *base-stock idling*, in order to differentiate it from base-stock policies without idling.) We determine the parameters of the BSI policy by considering the following deterministic optimization problem:

$$\begin{aligned}
\min_{w\in\mathbb{R}_+,\, \mathbf{b}\in\mathbb{R}^K,\, x_e\in C^K} \quad & \frac{1}{\tau}\left(\sum_{k\in\mathcal{K}} r_k + \int_0^\tau \psi\left(x_e(s)\right) ds\right) \\
s.t. \quad & x_{e,k}(0) = b_k - (\mu_k - \lambda_k)\rho_k \tau, \quad k \in \mathcal{K} \\
& \dot{x}_{e,k}(t) = \mu_k - \lambda_k \ \text{ for } t \in (0, \rho_k\tau), \quad k \in \mathcal{K} \\
& \dot{x}_{e,k}(t) = -\lambda_k \ \text{ for } t \in (\rho_k\tau, \tau), \quad k \in \mathcal{K} \\
& \tau = (s+w)/(1-\rho).
\end{aligned} \tag{2.3}$$

In particular, (2.3) can be understood as a restricted version of the ELSP, which optimizes the state trajectory over one production cycle (where the inventory levels at the beginning and end are held equal) for the deterministic inventory system under the BSI policy with base-stock levels

$\mathbf{b}$ and total idle time w in each cycle. The existence of an optimal solution to problem (2.3) is established under Assumption 2 below. Let $\bar{h}_k : \mathbb{R}_+ \to \mathbb{R}_+$ and $\bar{p}_k : \mathbb{R}_+ \to \mathbb{R}_+$ be defined via

$$\bar{h}_k(y) := \frac{1}{y}\int_0^y h_k(a_k(t))dt, \quad \text{and} \quad \bar{p}_k(y) := \frac{1}{y}\int_0^y p_k(a_k(t))dt,$$

where

$$a_k(t) := \begin{cases} (\mu_k - \lambda_k)t, & \text{if } 0 \le t \le y\rho_k \\ (\mu_k - \lambda_k)y\rho_k - \lambda_k(t - y\rho_k) & \text{if } y\rho_k < t \le y \end{cases}$$

Assumption 2 *For each $k \in \mathcal{K}$, the cost functions h_k and p_k are non-decreasing and continuous. In addition, the functions $\bar{h}_k$ and $\bar{p}_k$ satisfy $\bar{h}_k(y) \to \infty$ and $\bar{p}_k(y) \to \infty$ as $y \to \infty$. Furthermore, $h_k(x) = O(x^p)$ and $p_k(x) = O(x^p)$, for some $p \in \mathbb{N}$.*[2]

We mention that the third condition in the assumption, regarding the (sub)-polynomial growth rate of h_k and p_k, is not needed for problem (2.3), but is required in order to establish a certain uniform integrability result for the stochastic inventory system (see the proof of Lemma 5). We further remark that Assumption 2 is not restrictive in practice, and allows for a wide range of cost functions. Examples for $h_k(x)$ and $p_k(x)$ include αx, x^p, $\log(x)$, etc., for $\alpha \in \mathbb{R}_+$ and $p \in \mathbb{N}$. Unlike Assumption 1 which can be relaxed in special settings (see, e.g., Lemma 6), Assumption 2 is assumed to hold throughout the paper and is thus omitted from the statements of the main results.

Let $(\mathbf{b}_*, w_*)$ denote the optimal solution to (2.3), and let $\mathbf{w}_* := (w_{*,k}, k \in \mathcal{K})$ be a vector in $\mathbb{R}_+^K$ such that $\sum_{k\in\mathcal{K}} w_{*,k} = w_*$. We construct a sequence of BSI policies denoted by $\{\pi^n_{\mathrm{BSI},*} : n \ge 1\}$ with the following parameters:

(i) The base-stock level for product k in the nth system is $B^n_{*,k} := \lceil nb_{*,k} \rceil = \min\{z \in \mathbb{Z} : z \ge nb_{*,k}\}$, and $\mathbf{B}^n_* := (B^n_{*,k}, k \in \mathcal{K})$.

(ii) The idle time following the production of product k in the nth system is $W^n_{*,k} := nw_{*,k}$, and $\mathbf{W}^n_* := (W^n_{*,k}, k \in \mathcal{K})$.

Our main result establishes that the sequence of BSI policies $\{\pi^n_{\mathrm{BSI},*} : n \ge 1\}$ with parameters $\{\mathbf{B}^n_* : n \ge 1\}$ and $\{\mathbf{W}^n_* : n \ge 1\}$ is asymptotically optimal.

Theorem 1 *Under Assumption 1 and the large-setup-time scaling, it holds that*

$$\limsup_{n\to\infty} \bar{C}^n_{\pi^n_{BSI,*}} \le \liminf_{n\to\infty} \bar{C}^n_{\pi^n} \quad w.p.1, \tag{2.4}$$

for any other sequence of admissible controls $\{\pi^n : n \ge 1\}$

[2] For $f : \mathbb{R}^k \to [0,\infty)$, $g : \mathbb{R}^k \to [0,\infty)$, and $a \in \mathbb{R}^k_+ \cup \{\infty\}$, we write $f(x) = O(g(x))$ as $x \to a$ if $\limsup_{x\to a} f(x)/g(x) < \infty$, and $f(x) = o(g(x))$ if $\lim_{x\to a} f(x)/g(x) = 0$, $k \in \mathbb{N}$.

3. The Economic Lot Scheduling Problem

In this section, we study the ELSP corresponding to our SELSP. We provide a novel HDS representation for the system's dynamics, and generalize the existing common cycle approach (Hanssmann 1962) by optimizing over multi-cycle PE. In the ELSP, the demand, production and setup occur continuously and deterministically at the same rates as those in the stochastic system. We thus refer to the ELSP model as the "fluid model" of the original (stochastic) system.

To describe the ELSP, let for each product $k \in \mathcal{K}$, λ_k be the demand rate, μ_k be the production rate, s_k and r_k be the deterministic setup time and setup cost prior to the production of product k, and h_k and p_k be the holding cost function and backlog cost function, respectively. We use lower-case letters to denote the fluid counterparts of the corresponding stochastic processes. In particular, $x(t)$ denotes the inventory level, $z(t)$ is the facility location at time t, and $a(t)$ denotes the most recent polling epoch prior to time t. In addition, we let $a_k^{(m)}$, $d_k^{(m)}$ and $v_k^{(m)}$ denote the polling epoch, the departure epoch, and the busy time (the length of the production run) of product k during the mth production cycle, $k \in \mathcal{K}$, $m \geq 1$. We use $\tau^{(m)}$ to denote the length of the mth cycle, and $u^{(m-1)}$ to denote the beginning of the mth cycle.

For the fluid model in (C.4) to be well-defined, we must specify, for each cycle m, the idle times $(w_k^{(m)}, k \in \mathcal{K})$ and the production lengths $(v_k^{(m)}, k \in \mathcal{K})$, from which the values of $(a_k^{(m)}, d_k^{(m)}, k \in \mathcal{K})$ can be inferred. In turn, the values of $(w_k^{(m)}, k \in \mathcal{K})$ and $(v_k^{(m)}, k \in \mathcal{K})$ are determined by the idling and lot sizing policy.

To introduce the HDS representation of the fluid model, we note that its dynamics depend on the state of $x(t)$ the position of the facility $z(t)$, and potentially (depending on the lot sizing policy) the inventory level at the most recent polling epoch $x(a(t))$. Thus, we consider the process

$$\hat{x}(t) := (x(t), x(a(t)), z(t)) \quad t \geq 0.$$

Since the dynamics of $\hat{x}(t)$ are determined simultaneously by those of the continuous process $x(t)$, and the discontinuous process $z(t)$, having values in a finite state space, $\hat{x}$ is an HDS on $\mathbb{R}^K \times \mathbb{R}^K \times \mathcal{Z}$. Its state equations have the form

$$\begin{aligned} \dot{x}(t) &= f(z(t)) \\ z(t) &= g(x(t), x(a(t-)), z(t-)), \\ a(t) &= h(x(t), x(a(t-)), z(t-)), \end{aligned} \tag{3.1}$$

for functions $f : \mathcal{Z} \to \mathbb{R}^K$, $g : \mathbb{R}^K \times \mathbb{R}^K \times \mathcal{Z} \to \mathcal{Z}$, and $h : \mathbb{R}^K \times \mathbb{R}^K \times \mathcal{Z} \to \mathbb{R}_+$ described below.

First, the function f specifies the law of motion of $x(t)$ as a function of the location of the facility:

$$f_k(z(t)) = \begin{cases} -\lambda_k + \mu_k & \text{if } z(t) = k \\ -\lambda_k & \text{otherwise} \end{cases}, \quad k \in \mathcal{K}.$$

The functions g and h are determined by the control that is exercised, and can be described via K *lot sizing functions* $\phi_k : \mathbb{R}^K \to \mathbb{R}_+$, and K *idling functions* $\eta_k : \mathbb{R}^K \to \mathbb{R}_+$, such that

$$v_k^{(m)} = \phi_k(x(a_k^{(m)})), \quad w_k^{(m)} = \eta_k(x(a_k^{(m)})), \quad k \in \mathcal{K}, \quad m \geq 1. \tag{3.2}$$

The functions $(\phi_k, \eta_k, k \in \mathcal{K})$ then jointly define the functions g for the server location, and h for the most recent polling epoch. Since the exact forms of g and h are algebraically lengthy, we relegate their characterization to Appendix A.

We will henceforth refer to equations of the form (3.1) as HDS and to $\hat{x}$ as a solution to the HDS. We will sometimes also refer to the inventory level x alone as a "solution" to the HDS of a "fluid model," when a control is specified.

3.1. Optimizing the Control of the ESLP

Since we optimize the long-run average cost, we need to study the long-run behavior of the fluid model; we thus begin by defining some fundamental terms.

Definition 4 (periodic equilibrium) *A solution $\hat{x}_e$ to the HDS* (3.1) *is a periodic equilibrium (PE) for the HDS if there exists $\tau > 0$ such that $\hat{x}_e(t+\tau) = \hat{x}_e(t)$ for all $t \geq 0$. The smallest such τ is called the period.*

Note that a PE may contain more than one cycle of the facility. We refer to a PE whose period contains L cycles as an *L-cycle PE*.

For a given control, we would like to know whether the system eventually converges to a (desirable) PE. Since a PE is itself a full trajectory, convergence of any other trajectory to the PE will hold if it is possible to "align" the two trajectories in such a way that the two become uniformly close to each other as time increases.

Definition 5 (convergence to a PE) *A solution $\hat{x}$ to the HDS* (3.1) *is said to converge to a PE $\hat{x}_e$ if there exists $L \in \mathbb{N}$, such that $||x(u^{(mL)} + \cdot) - x_e(\cdot)||_t \to 0$ as $m \to \infty$, for all $t > 0$.*

Finally, a desirable property of a PE is for it to be a *global limit cycle*, namely, a PE to which all the solutions to the HDS converge. Clearly, if a global limit cycle exists, then it must also be the unique PE.

3.2. The Fluid Control Problem

We assume that each control u of the ESLP satisfies the following two conditions:

(i) There exists a unique solution $x_u^\phi := \{x_u^\phi(t) : t \geq 0\}$ to the HDS (3.1) for any initial condition $\phi \in \mathbb{R}^K$.

(ii) Any solution x_u^ϕ converges to a PE as $t \to \infty$.

We denote the set of all such controls by $\mathcal{U}$. Note that a solution x_u^ϕ to the HDS need not converge to a one-cycle PE. In other words, we allow x_u^ϕ to converge to an L-cycle PE for some $L \in \mathbb{N}$. Nevertheless, we show below that if is sufficient to restrict attention to a subclass of $\mathcal{U}$ of controls that guarantee that a one-cycle global limit cycle exists.

Let $\gamma_{u,k}^\phi(t)$ denote the total number of setups for product k by time t when the initial condition is ϕ and control u is applied. For $\phi \in \mathbb{R}^K$, let

$$C_u(\phi) := \lim_{t\to\infty} \frac{1}{t}\left(\sum_{k\in\mathcal{K}} r_k \gamma_{u,k}^\phi(t) + \int_0^t \psi\left(x_u^\phi(s)\right) ds\right).$$

The aim of the FCP is to identify a fluid-optimal control $u^* \in \mathcal{U}$ that satisfies

$$C_{u^*}(\phi) \leq C_u(\phi) \quad \text{for all } \phi \in \mathbb{R}^K, \quad u \in \mathcal{U}.$$

We next show that optimizing the long-run average cost can be reduced to optimizing over all PE. Let x_u^ϕ denote the unique solution to the HDS when control $u \in \mathcal{U}$ is employed and when the initial condition is ϕ. Let x_e^ϕ denote the PE to which x_u^ϕ converges, and let τ_e^ϕ denote its period which contains L_e^ϕ production cycles, for some $L_e^\phi \in \mathbb{N}$.

Lemma 2 *We have*

$$C_u(\phi) = \frac{1}{\tau_e^\phi}\left(\sum_{k\in\mathcal{K}} r_k L_e^\phi + \int_0^{\tau_e^\phi} \psi\left(x_e^\phi(s)\right) ds\right).$$

Let $\underline{c} := \inf_{u\in\mathcal{U}} \inf_{\phi\in\mathbb{R}^K} C_u(\phi)$. Due to Lemma 2, we solve the FCP by searching for a control $u_* \in \mathcal{U}$ under which there exists a global limit cycle x_* (with period τ_* containing L_* production cycles) whose long-run average cost is $\underline{c}$, namely,

$$\underline{c} = \frac{1}{\tau_*}\left(\sum_{k\in\mathcal{K}} r_k L_* + \int_0^{\tau_*} \psi\left(x_*(s)\right) ds\right). \tag{3.3}$$

By Lemma 2 and Lemma 3 below, (3.3) holds if and only if

$$\frac{1}{\tau_*}\left(\sum_{k\in\mathcal{K}} r_k L_* + \int_0^{\tau_*} \psi\left(x_*(s)\right) ds\right) \leq \frac{1}{\tau_e}\left(\sum_{k\in\mathcal{K}} r_k L_e + \int_0^{\tau_e} \psi\left(x_e(s)\right) ds\right),$$

for any other PE x_e (under some control for some initial condition), whose period is τ_e and contains L_e production cycles.

3.3. The Optimal Base-Stock Idling Policy for the Fluid Model

To solve the FCP, we start by identifying closed curves in $\mathbb{R}^K$ that are possible solutions of the HDS, namely, they can be obtained under some control for some initial condition. We refer to such a closed curve x_e as a *PE-candidate* and treat it as a mapping from $[0,\tau_e]$ to $\mathbb{R}^K$, satisfying $x_e(\tau_e)=x_e(0)$. We then optimize over all possible PE-candidates in order to find the optimal PE-candidate x_* with the lowest time-average cost. Finally, we design an optimal control $u_*\in\mathcal{U}$ under which the optimal PE-candidate is a global limit cycle for the HDS.

3.3.1. Computing an Optimal PE-Candidate Formally, an optimal PE-candidate solves the fluid optimization problem

$$\begin{aligned} \min_{L\in\mathbb{N},\ \tau_L\in\mathbb{R}_+,\ x_e\in C^K} \quad & \frac{1}{\tau_L}\left(\sum_{k\in\mathcal{K}} r_k L+\int_0^{\tau_L}\psi\left(x_e(s)\right)ds\right) \\ s.t. \quad & x_e \text{ is an L-cycle PE-candidate with period } \tau_L. \end{aligned} \tag{3.4}$$

We note that the fluid optimization problem 3.4 is equivalent to the common cycle approach for the ELSP (Hanssmann 1962, Maxwell 1964, Elmaghraby 1978) if (i) backlog is not permitted, (ii) the holding cost function h_k is linear for all $k\in\mathcal{K}$, and (iii) L is prespecified at the value of 1 as opposed to being a control variable to optimize. While one may expect that an optimal solution to (3.4) outperforms the optimal one-cycle PE-candidate, we shall see below that setting L to 1 is indeed optimal for (3.4).

Solving problem (3.4) is prohibitive, as it requires solving for all $L\geq 1$. However, the following proposition shows that one can focus on the case $L=1$.

Proposition 1 *There exists an optimal solution to* (3.4) *with* $L=1$*.*

In light of Proposition 1, an optimal solution to the following problem also solves (3.4).

$$\begin{aligned} \min_{\tau\in\mathbb{R}_+,\ x_e\in C^K} \quad & \frac{1}{\tau}\left(\sum_{k\in\mathcal{K}} r_k+\int_0^{\tau}\psi\left(x_e(s)\right)ds\right) \\ s.t. \quad & x_e \text{ is a one-cycle PE-candidate with period } \tau. \end{aligned} \tag{3.5}$$

We can parametrize any one-cycle PE-candidate x_e with period τ, by letting b_k denote the *produce-up-to* level of product k in its production run, namely,

$$b_k := \max\left\{x_{e,k}(t): 0\leq t\leq\tau\right\},\quad k\in\mathcal{K}.$$

For a given idle-time vector $\mathbf{w}$, let $w:=\sum_{k\in\mathcal{K}} w_k$. It follows from basic flow balance equations that the period τ satisfies

$$\tau=(s+w)/(1-\rho). \tag{3.6}$$

To see this, note that, since the inventory level at the beginning and the end of a period is identical, each demand arriving during this time interval gets produced eventually. Hence, the facility must be producing a fraction ρ of the time, and setting up or idling for a fraction $(1-\rho)$ of the time. Since the total setup plus idle time over the production cycles is $(s+w)$, it holds that $\tau(1-\rho)=(s+w)$, which gives (3.6). as illustrated in Figure 2. Next, it can be easily observed (see, e.g., Figure 2) that the inventory level $x_{e,k}$ of product k is fully determined by τ and b_k, and thus by w and b_k due to (3.6). In addition, the one-cycle fluid optimization problem (3.5) depends on $\mathbf{w}$ only through the sum of its elements, w. Let $\mathbf{b} := (b_k, k \in \mathcal{K})$. Problem (3.5) can be equivalently written as optimizing over w and $\mathbf{b}$, which gives rise to the fluid optimization problem (2.3), which is the deterministic optimization problem that we introduce in Section 2.4 (preview of main results) to solve for the parameters of the asymptotically optimal BSI policy.

Figure 2 **Inventory level of product k in a one-cycle PE-candidate**

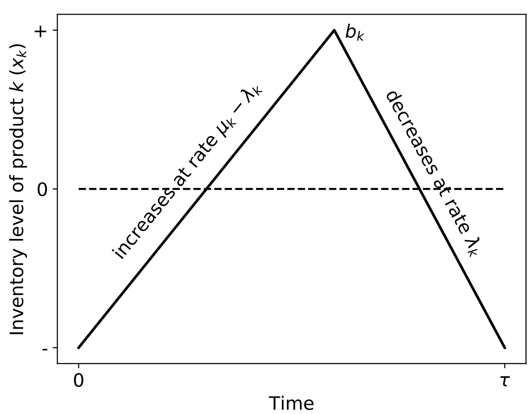

For a given cost function, problem (2.3) can be solved analytically or numerically with little effort. We next provide two examples where the optimal solution to (2.3) can be characterized in closed form.

Example 1 (linear holding and backlog costs, no setup cost) If $r_k = 0$ for all $k \in \mathcal{K}$, and if p_k and h_k are linear, i.e., $h_k(x) = h_k x$ and $p_k(x) = p_k x$ for some constants $h_k \in \mathbb{R}_+$ and $p_k \in \mathbb{R}_+$, then the optimal solution to the fluid optimization problem (2.3) has $w_* = 0$, and so the period is $\tau_* = s/(1-\rho)$. Furthermore, the optimal inventory level of product k increases from $-\frac{h_k}{p_k+h_k}(1-\rho_k)\lambda_k\tau_*$ to $\frac{p_k}{p_k+h_k}(1-\rho_k)\lambda_k\tau_*$ each time product $k \in \mathcal{K}$ is polled.

The next example shows that in the presence of setup costs, it may be beneficial to idle the facility in order to decrease the frequency of switching and reduce the long-run average setup cost.

Example 2 (symmetric two-product system) Consider a two-product symmetric system with $\lambda := \lambda_1 = \lambda_2$, $\mu := \mu_1 = \mu_2$, $s := s_1 = s_2$, $r := r_1 = r_2$, and linear cost functions with rates $c := h_1 = h_2 = p_1 = p_2$. Let

$$T_{sym} := \frac{cs^2\lambda\mu(\mu-\lambda)}{4(-2\lambda+\mu)^2}.$$

Then the optimal solution to the fluid optimization problem (2.3) is as follows.

(i) If $r > T_{sym}$, then

$$w_* = \frac{cs\lambda(\mu-\lambda) + 2(2\lambda-\mu)\sqrt{cr\lambda(\mu-\lambda)/\mu}}{c\lambda(\lambda-\mu)},$$

and the optimal inventory level of product k increases from $-\frac{\sqrt{cr\lambda(\mu-\lambda)/\mu}}{c}$ to $\frac{\sqrt{cr\lambda(\mu-\lambda)/\mu}}{c}$ each time product k is polled, $k = 1, 2$.

(ii) If $r \le T_{sym}$, then $w_* = 0$, and the optimal inventory level of product k increases from $-\frac{s\lambda(\mu-\lambda)}{2\mu-4\lambda}$ to $\frac{s\lambda(\mu-\lambda)}{2\mu-4\lambda}$ each time product k is polled, $k = 1, 2$.

3.3.2. Designing a BSI Policy that Achieves the Optimal PE Let $(\mathbf{b}_*, w_*)$ denote an optimal solution to the one-cycle fluid optimization problem (2.3), and c_* denote the optimal objective value. Observe that it necessarily holds that $c_* \le \underline{c}$, for $\underline{c}$ defined in (3.3). We next show that there exists some BSI policy with properly selected parameters, under which x_* is a global limit cycle for the HDS, implying that the reverse is also true, namely, that $c_* \ge \underline{c}$, so that $c_* = \underline{c}$.

Lemma 3 (global stability of BSI) *Let x_e be a one-cycle PE-candidate with produce-up-to levels $\mathbf{b}_e$ and total idle time w_e. Then x_e is a global limit cycle for the HDS* (3.1) *under the BSI policy with base-stock levels $\mathbf{b}_e$ and any idle time vector $\mathbf{w}_e \in \mathbb{R}_+^K$ that satisfies $\sum_{k\in\mathcal{K}} w_{e,k} = w_e$.*

Lemma 3 applies to any PE-candidate x_e given its parameterization $(\mathbf{b}_e, w_e)$, and thus to an optimal solution of the one-cycle fluid optimization problem (2.3). Since the global limit cycle under the translated BSI policy has time-average cost $c_* = \underline{c}$, the optimality of the BSI policy for the FCP is straightforward, and is formalized in the next theorem.

Theorem 2 (fluid-optimal BSI) *Let $(\mathbf{b}_*, w_*)$ be an optimal solution to the one-cycle fluid optimization problem* (2.3)*. Then the BSI policy with base-stock levels $\mathbf{b}_*$ and any idle time vector $\mathbf{w}_*$ that satisfies $\sum_{k\in\mathcal{K}} w_{*,k} = w_*$ is optimal for the FCP.*

4. Asymptotic Optimality

We first establish that c_* is a lower bound on the achievable costs, asymptotically.

Lemma 4 *For any sequence of admissible controls $\{\pi^n : n \ge 1\}$, it holds that $\liminf_{n\to\infty} \bar{C}^n_{\pi^n} \ge c_*$ w.p.1.*

Recall from Section 2.4 that π^n_{BSI} denotes the control for the nth system that employs the base-stock policy with some base-stock level $\mathbf{B}^n$ and idle time vector $\mathbf{W}^n$, $n \geq 1$. We let $\{\pi^n_{BSI,*} : n \geq 1\}$ denote the sequence of controls with parameters determined by the fluid-optimal BSI policy as follows $B^n_{*,k} := \lceil nb_{*,k} \rceil$ and $W^n_{*,k} := nw_{*,k}$ for $k \in \mathcal{K}$, $n \geq 1$.

Lemma 5 *Under Assumption 1, it holds that* $\lim_{n\to\infty} \bar{C}^n_{\pi^n_{BSI,*}} = c_*$.

Lemmas 4 and 5 imply that $\{\pi^n_{BSI,*} : n \geq 1\}$ is asymptotically optimal, from which Theorem 1 follows.

4.1. Linear Holding and Backlog Costs

The distributional assumption on the production and setup times needed for Lemma 5 largely depends on the cost structure. For the special case in which the holding and backlog costs are linear, we can relax Assumption 1 and impose conditions only on the second moments of the production and setup time distributions.

Assumption 3 *For all* $k \in \mathcal{K}$,

(i) $\mathbb{E}\left[G_k^2\right] < \infty$.

(ii) $\mathbb{E}\left[(S^n_k)^2\right] < \infty$ *for all* $n \geq 1$ *and, in addition,* $\mathbb{E}\left[(\bar{S}^n_k)^2\right] \to s_k^2$ *as* $n \to \infty$.

Lemma 6 *Assume that the holding and backlog cost functions* h_k *and* p_k *are linear for all* $k \in \mathcal{K}$, *and that Assumption 3 holds. Then*

$$\lim_{n\to\infty} \bar{C}^n_{\pi^n_{BSI,*}} = c_*.$$

4.2. The Proposed Policy with Stochastic Refinement

There are multiple ways to translate the optimal solution to the FCP into a sequence of controls for the stochastic systems. All of them achieve the optimal performance in the limit, but their pre-limit performances can differ. In this section, we discuss a refined policy for convex holding and backlog costs, in which, for each $n \geq 1$, we fix the idle time vector $\mathbf{W}^n_* = n\mathbf{w}_*$, and solve for the exact optimal base-stock level given the idle times. While this policy achieves the same limit as $\{\pi^n_{BSI,*} : n \geq 1\}$, it improves the performance of the stochastic system.

Given vectors of base-stock levels $\mathbf{B}^n$ and idle times $\mathbf{W}^n$, we call

$$Y^n_k(t) := B^n_k - X^n_k(t), \quad k \in \mathcal{K}, \quad t \geq 0,$$

the *shortfall process* of product k; it is the difference between the prescribed base-stock level and the actual inventory level.

Let $Y^n := (Y_k^n, k \in \mathcal{K})$. The advantage of working with the shortfall process Y^n over working with X^n is that the evolution of Y^n does not depend on the base-stock level $\mathbf{B}^n$. To see this, assume without loss of generality, that the initial inventory level is below the base-stock level, i.e., that $X_k^n(0) \leq B_k^n$, for all $k \in \mathcal{K}$. It is then easy to see that the evolution of Y^n is the same as that of the queue process (with the same initial condition) in a polling system under the exhaustive policy: Y_k^n is driven up by the arrival process of demands ("customers"), and down by the production process of product k ("services"). Each time the facility polls product k, it continues producing product k until Y_k^n reaches 0. By construction, Y^n stays non-negative at all time.

In what follows, we use $Y_k^n(\infty)$ to denote the steady-state shortfall level for product k in the nth system, which is well-defined under any stable control, $k \in \mathcal{K}$, $n \geq 1$. It is immediate that the distribution of $Y_k^n(\infty)$ also does not depend on the base-stock level $\mathbf{B}^n$. Indeed, $Y_k^n(\infty)$ is completely determined by the following three primitives of the model plus the prescribed idle times: $(\boldsymbol{\lambda}, \mathbf{G}, \mathbf{S}^n, \mathbf{W}^n)$, where $\boldsymbol{\lambda} := (\lambda_k, k \in \mathcal{K})$, $\mathbf{G} := (G_k, k \in \mathcal{K})$, and $\mathbf{S}^n := (S_k^n, k \in \mathcal{K})$.

The following proposition follows directly from Proposition 1 in Federgruen and Katalan (1996b).

Proposition 2 *Assume that the holding and backlog costs h_k and p_k are convex for each $k \in \mathcal{K}$. For the nth stochastic system, $n \geq 1$, consider all base-stock policies with the given idle time vector $\mathbf{W}_*^n$. Let $(Y_k^n(\infty), k \in \mathcal{K})$ denote the steady-state shortfall level. For product $k \in \mathcal{K}$, the optimal base-stock level $B_{*,k}^{'n}$ is obtained by determining the unique minimum of the single-variable convex function*

$$\psi_k^n(x) = \mathbb{E}\left[h_k\left([x - Y_k^n(\infty)]^+\right) + p_k\left([Y_k^n(\infty) - x]^+\right)\right], \quad k \in \mathcal{K}. \tag{4.1}$$

By Proposition 2, the optimal base-stock level $B_{*,k}^{'n}$ is the optimal solution to a newsvendor problem with random demand $Y_k^n(\infty)$, overage cost rate h_k, and underage cost rate p_k, $k \in \mathcal{K}$. Note that characterizing $\mathbf{B}_*^{'n} := (B_{*,k}^{'n}, k \in \mathcal{K})$ requires knowing the distribution of the steady-state shortfall level $Y^n(\infty)$. An efficient algorithm to approximate the distribution of $Y^n(\infty)$ is provided in Federgruen and Katalan (1996b), and thus $\mathbf{B}_*^{'n}$ can be computed numerically. Alternatively, the distribution of $Y^n(\infty)$ can be estimated via simulation.

Recall that under Assumption 1 (alternatively, Assumption 3 for linear costs), Theorem 1 establishes that the sequence of fluid-translated BSI policies $\{\pi_{BSI,*}^n : n \geq 1\}$ is asymptotically optimal. Let $\{\pi_{BSI,*}^{'n} : n \geq 1\}$ denote the sequence of refined BSI policies with the exact optimal base-stock levels $\{\mathbf{B}_*^{'n} : n \geq 1\}$ for the given vectors of idle times $\{\mathbf{W}_*^n : n \geq 1\}$. Since for each $n \geq 1$, the refined policy $\pi_{BSI,*}^{'n}$ is no worse than $\pi_{BSI,*}^n$, the sequence $\{\pi_{BSI,*}^{'n} : n \geq 1\}$ must be asymptotically optimal as well, which yields the following corollary.

Corollary 1 *Assume that the holding and backlog costs h_k and p_k are convex for each $k \in \mathcal{K}$. Under Assumption 1 (alternatively, Assumption 3 for linear holding and backlog cost functions p_k and h_k, $k \in \mathcal{K}$), the sequence of controls $\{\pi'^{n}_{BSI,*} : n \geq 1\}$ with stochastically refined base-stock levels is asymptotically optimal; in particular $\bar{C}'^{n}_{\pi^n_{BSI,*}} \to c_*$ as $n \to \infty$.*

Example 3 (linear holding and backlog costs, no setup cost) Assume that $r_k = 0$, and that p_k and h_k are linear, for all $k \in \mathcal{K}$. It was shown in Example 1 that the optimal solution to the fluid optimization problem (2.3) has total idle time $w_* = 0$, cycle length $\tau_* = s/(1-\rho)$, and produce-up-to level $b_{*,k} = p_k(p_k + h_k)^{-1}(1-\rho_k)\lambda_k \tau_*$, $k \in \mathcal{K}$. The direct translation of the fluid-optimal solution yields

$$B^n_{*,k} = \lceil n b_{*,k} \rceil \quad \text{and} \quad W^n_{*,k} = n w_{*,k} = 0, \quad k \in \mathcal{K}, \quad n \geq 1.$$

For the refined policy, we impose the same idle time vector $\mathbf{W}^n_* = \mathbf{0}$ for the nth system, $n \geq 1$. It follows from Proposition 2 that, given $\mathbf{W}^n_*$, the exact optimal base-stock level for the nth system is given by

$$B'^{n}_{*,k} = F^{-1}_{Y^n_k(\infty)}\left(\frac{p_k}{p_k + h_k}\right), \quad k \in \mathcal{K}, \quad n \geq 1, \tag{4.2}$$

where $F^{-1}_{Y^n_k(\infty)}$ is the inverse of the cumulative distribution function (cdf) of the steady-state shortfall level $Y^n_k(\infty)$. By Theorem 3.10 and equation (26) in Winands (2011), the steady-state shortfall level converges weakly to a uniform random variable under the fluid scaling, namely,

$$\bar{Y}^n_k(\infty) \;\Rightarrow\; U\left[0, (1-\rho_k)\lambda_k \tau_*\right] \quad \text{as } n \to \infty, \quad k \in \mathcal{K}. \tag{4.3}$$

Thus, by (4.2) and (4.3), the sequence of fluid-scaled optimal base-stock levels for product k converges to the inverse of the cdf of $U\left[0, (1-\rho_k)\lambda_k \tau_*\right]$ evaluated at the ratio $p_k/(p_k + h_k)$, i.e.,

$$\bar{B}'^{n}_{*,k} \;\to\; \frac{p_k}{p_k + h_k}(1-\rho_k)\lambda_k \tau_* = b_{*,k} \quad \text{as } n \to \infty, \quad k \in \mathcal{K}, \tag{4.4}$$

where the equality follows from the solution to the FCP in Example 1. We note that the convergence of $\bar{B}'^{n}_{*,k}$ to $b_{*,k}$ in (4.4) is expected, because $\bar{B}^n_{*,k} \to b_{*,k}$ as $n \to \infty$ by construction, and both $\mathbf{B}'^{n}_*$ and $\mathbf{B}^n_*$ are asymptotically optimal for the given idle time vector $\mathbf{W}^n_*$.

5. A Numerical Study of the Fluid-Translated BSI Policy

Based on our asymptotic optimality result, we propose the following heuristic to develop good BSI policies:

Step 1. Solve the one-cycle fluid optimization problem (2.3) analytically or numerically, and obtain an optimal produce-up-to level $\mathbf{b}_*$ and cumulated idle time w_*.

Step 2. Construct a waiting time vector $\mathbf{W}_* \in \mathbb{R}_+^K$, with $\sum_{k\in\mathcal{K}} W_{*,k} = w_*$. **Step 3.** Take the base-stock levels to be $\mathbf{B}_* = \lceil \mathbf{b}_* \rceil$. Alternatively, for the given idle time vector $\mathbf{W}_*$, compute the exact optimal base-stock level $\mathbf{B}'_*$ as in Proposition 2, either numerically (e.g., via the algorithm in Federgruen and Katalan (1996b)), or via simulation. We refer to the former base-stock levels as *fluid-translated* and the latter as *stochastically-refined*.

In Step 2 of the heuristic, there are multiple choices for the idle time vector. In particular, how the sum w_* is allocated among the K elements of the fluid idle time vector $\mathbf{w}_*$ is immaterial for the HDS, because the optimal PE-candidate x_* is a global limit cycle under any such BSI policy, so that all these policies are asymptotically optimal. However, different divisions of w_* among the elements of $\mathbf{W}_*$ may lead to different performances in the stochastic system. Similarly, in Step 3, we propose two alternative ways to construct the base-stock levels. While they are both asymptotically optimal, the refined base-stock level $\mathbf{B}'_*$ leads to better performance in the stochastic system for a given $\mathbf{W}_*$. In what follows, we compare the performance of different translations of the fluid optimal policy using simulation.

We examine twelve systems with different setup times and traffic intensities. These systems are grouped into three parameter sets as summarized in Table 1. Each system has three products, exponential inter-demand and production times, and deterministic setup times. We assume polynomial holding cost functions, where each p_k, $k \in \mathcal{K}$, is of the form $p_k(x) = (ax+b)^c$ for some coefficient $a \in \mathbb{R}$, constant $b \in \mathbb{R}$, and exponent $c \in \mathbb{R}$. The same structure is assumed for the backlog costs. Within each set, we increase the setup times from 1 to 4. In addition, we increase the nominal traffic intensity ρ from 0.7 to 0.9 at an increment of 0.1 across the three sets. For the systems listed in Table 1, we numerically solve the fluid optimization problem (2.3); the solutions can be found in Table 2. We observe that for the optimal fluid solutions, the total idle time tends to decrease and the base-stock levels tend to increase as the setup times and nominal traffic intensity grow. Note that in Table 2, several instances have the same optimal base-stock levels and objective values, e.g., systems 1–4 and systems 5 and 6. We formalize the intuition behind this phenomenon in Lemma 16 in Appendix E.1.

5.1. Translation of the Idle Times

For each system listed in Table 1, we assess four different divisions of the total idle time w_*: 1) $\mathbf{W}_* = (w_*/3, w_*/3, w_*/3)$, 2) $\mathbf{W}_* = (w_*, 0, 0)$, 3) $\mathbf{W}_* = (0, w_*, 0)$, and 4) $\mathbf{W}_* = (0, 0, w_*)$. We simulate the system under the BSI policy with base-stock levels $\mathbf{B}_* = \lceil \mathbf{b}_* \rceil$ and the four idle time vectors, respectively. Table 3 summarizes the long-run average costs of the stochastic system under

Table 1 System parameters

	System										p_1			p_2			p_3			h_1			h_2			h_3					
		λ_1	λ_2	λ_3	μ_1	μ_2	μ_3	s_1	s_2	s_3	a	b	c	a	b	c	a	b	c	a	b	c	a	b	c	a	b	c	r_1	r_2	r_3
Set 1 ($\rho = 0.7$)	**1**	**0.23**	**0.23**	**0.23**	1	1	1	**1**	**1**	**1**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	2	0.23	0.23	0.23	1	1	1	**2**	**2**	**2**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	3	0.23	0.23	0.23	1	1	1	**3**	**3**	**3**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	4	0.23	0.23	0.23	1	1	1	**4**	**4**	**4**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
Set 2 ($\rho = 0.8$)	**5**	**0.27**	**0.27**	**0.27**	1	1	1	**1**	**1**	**1**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	6	0.27	0.27	0.27	1	1	1	**2**	**2**	**2**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	7	0.27	0.27	0.27	1	1	1	**3**	**3**	**3**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	8	0.27	0.27	0.27	1	1	1	**4**	**4**	**4**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
Set 3 ($\rho = 0.9$)	**9**	**0.3**	**0.3**	**0.3**	1	1	1	**1**	**1**	**1**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	10	0.3	0.3	0.3	1	1	1	**2**	**2**	**2**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	11	0.3	0.3	0.3	1	1	1	**3**	**3**	**3**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000
	12	0.3	0.3	0.3	1	1	1	**4**	**4**	**4**	1	1	2	2	2	2	3	3	2	3	1	2	2	2	2	1	3	2	1000	2000	1000

Table 2 Optimal fluid solutions

	System	Total idle time	Base-stock level			Cost
			b_1	b_2	b_3	
Set 1 ($\rho = 0.7$)	**1**	10.4	2.5	3.8	5.7	181.1
	2	7.4	2.5	3.8	5.7	181.1
	3	4.4	2.5	3.8	5.7	181.1
	4	1.4	2.5	3.8	5.7	181.1
Set 2 ($\rho = 0.8$)	**5**	4.7	2.7	4.0	5.9	191.3
	6	1.7	2.7	4.0	5.9	191.3
	7	0.0	3.1	4.7	6.7	194.3
	8	0.0	4.3	6.2	8.7	219.1
Set 3 ($\rho = 0.9$)	**9**	0.9	2.7	4.1	6.0	197.6
	10	0.0	4.3	6.3	8.8	224.9
	11	0.0	6.6	9.5	12.8	322.0
	12	0.0	9.0	12.6	16.8	464.5

the four policies. We also report the maximum performance difference, defined as the percentage gap between the best and worst costs, among the four long-run average costs in the "Max gap" column. We observe that the performances among different assignments of the total idle time w_* are very similar. The gap between the best and the worst is less than 5% in all systems tested. (For systems with $w_* = 0$ in the optimal fluid solution (systems 10–12), there is no difference among the four divisions because $\mathbf{W}_* = \mathbf{0}$.)

5.2. Translation of the Base-Stock Levels

In this section, we study the performance of the simple fluid-translated BSI policy with base-stock level $\mathbf{B}_* = \lceil \mathbf{b}_* \rceil$ and idle time $\mathbf{W}_* = (w_*/3, w_*/3, w_*/3)$ in the stochastic system. We consider a benchmark BSI policy where we optimize both the base-stock levels and the total idle time for the stochastic system. In particular, for a given total idle time $w \in \mathbb{R}_+$, we consider the idle time vector $\mathbf{W} = (w/3, w/3, w/3)$, and use simulation to find the optimal base-stock levels corresponding to $\mathbf{W}$ (Proposition 2). The optimal base-stock levels and total idle time are calculated by enumerating

Table 3 **Idle time assignments for the fluid-translated BSI policy**

Fluid-Translated BSI										
	System	**Total idle time**	**Base-stock level**			**Cost**				**Max gap**
			b_1	b_2	b_3	$(w_*/3, w_*/3, w_*/3)$	$(w_*, 0, 0)$	$(0, w_*, 0)$	$(0, 0, w_*)$	
Set 1 ($\rho = 0.7$)	**1**	10.4	3	4	6	285.4	281.8	285.0	286.2	**1.6%**
	2	7.4	3	4	6	285.4	288.6	281.3	289.8	**3.0%**
	3	4.4	3	4	6	285.4	287.1	281.7	282.4	**1.9%**
	4	1.4	3	4	6	285.4	282.4	280.4	287.4	**2.5%**
Set 2 ($\rho = 0.8$)	**5**	4.7	3	4	6	360.3	361.4	366.1	372.1	**3.3%**
	6	1.7	3	4	6	360.3	363.4	371.2	354.3	**4.8%**
	7	0.0	3	5	7	381.9	381.9	381.9	381.9	**0.0%**
	8	0.0	4	6	9	408.8	408.8	408.8	408.8	**0.0%**
Set 3 ($\rho = 0.9$)	**9**	0.9	3	4	6	573.0	581.4	585.7	600.0	**4.7%**
	10	0.0	4	6	9	721.1	721.1	721.1	721.1	**0.0%**
	11	0.0	7	9	13	792.3	792.3	792.3	792.3	**0.0%**
	12	0.0	9	13	17	1041.2	1041.2	1041.2	1041.2	**0.0%**

the values of w (with the corresponding optimal base-stock levels) to minimize the long-run average cost. We refer to this benchmark policy as the optimal BSI policy.

Table 4 summarizes the parameters of the optimal BSI policy, including the total idle time and the base-stock levels, and the cost under this policy. We then compare the simulated long-run average cost under the fluid-translated BSI policy with parameters $(\mathbf{B}_*, \mathbf{W}_*)$ to that under the optimal BSI policy, and summarize the optimality gap in the last column of Tables 4. We observe that the optimality gap is small in most cases. The gap is less than 6% for systems 1–8, 11, and 12. However, when the system is critically loaded, i.e., $\rho = 0.9$, and when the setup times are small, i.e., $s_i = 1$ or 2, the optimality gap can be larger than 15%. The low accuracy in these examples follows from the fact that the fluid model is not an appropriate approximation because the system is in heavy traffic when ρ is large and the setup times are relatively small; a diffusion approximation may be more appropriate in such cases; see, e.g., Markowitz et al. (2000).

6. Exact MDP Solutions

In this section, we compare the fluid-translated BSI policies $(\mathbf{B}_*^n, \mathbf{W}_*^n)$ to the optimal polices derived via MDP. We consider two MDP formulations for the SELSP. The first MDP generates policies that are within our set of admission controls, while the second MDP may generate policies that are not admissible according to Definition 1. In the first MDP (termed "MDP Idle-Fix"), we impose the fluid-translated idle time vector $\mathbf{W}_*$, and solve the MDP for the optimal lot sizing policy. In the second MDP (termed "MDP General"), we relax the assumption on the static idle time vector, and allow state-dependent idling behavior. In particular, when to start idling and how long to idle before switching to (setting up for) the next product may depend on the inventory level at the

Table 4 Exponentially distributed production times

Simulated optimal BSI							Opt gap of fluid-translated BSI
	System	Total idle time	Base-stock level b_1	b_2	b_3	Cost $(w_*/3, w_*/3, w_*/3)$	
Set 1 ($\rho = 0.7$)	**1**	10.1	3	4	6	272.4	**4.8%**
	2	7.1	3	4	6	272.4	**4.8%**
	3	4.1	3	4	6	272.4	**4.8%**
	4	1.1	3	4	6	272.4	**4.8%**
Set 2 ($\rho = 0.8$)	**5**	4.0	3	5	7	342.0	**5.3%**
	6	1.0	3	5	7	342.0	**5.3%**
	7	0.4	4	6	9	366.4	**4.2%**
	8	0.0	5	7	11	407.5	**0.3%**
Set 3 ($\rho = 0.9$)	**9**	0.3	4	6	9	494.6	**15.9%**
	10	0.5	6	9	14	626.4	**15.1%**
	11	0.0	8	12	17	748.6	**5.8%**
	12	0.0	10	15	21	1004.0	**3.7%**

decision epoch. We provide the detailed MDP formulations in Appendix E.2. We mention at the outset that using MDP is computationally intensive, and is therefore not appropriate for systems with multiple products. Thus, our goal is to demonstrate the effectiveness of our simple asymptotic analyses.

To keep the problem tractable, we consider a two-product system with exponential inter-demand and production times, deterministic setup times, and linear costs. The parameters for the problem are summarized in Table 5. In each set of parameters, we examine a system of size $n = 1, 5, 10, 15$, where the mean setup time for the nth system is $s_k^n = s_k n$, and the setup cost is $r_k^n = r_k n^2$, for the "base" mean setup time s_k and setup cost r_k in the system with $n = 1$, $k \in \mathcal{K}$. In the case of linear costs, inflating the setup costs by n^2 is consistent with the large-setup-time scaling introduced in Section 2.3. To see this, note that the fluid-scaled long-run average cost for the nth system can be expressed as

$$\begin{aligned}
\bar{C}^n_{\pi^n} &= \frac{1}{n} \lim_{t\to\infty} \sum_{k\in\mathcal{K}} \left(\frac{1}{t} r_k n^2 \Gamma^n_{\pi^n,k}(t) + \frac{1}{t} \int_0^t \left(h_k\left(X^n_{\pi^n,k}(s)^+\right) + p_k\left(X^n_{\pi^n,k}(s)^-\right)\right) ds \right) \\
&= \frac{1}{n} \lim_{t\to\infty} \sum_{k\in\mathcal{K}} \left(\frac{1}{nt} r_k n^2 \Gamma^n_{\pi^n,k}(nt) + \frac{1}{nt} \int_0^{nt} \left(h_k\left(X^n_{\pi^n,k}(s)^+\right) + p_k\left(X^n_{\pi^n,k}(s)^-\right)\right) ds \right) \\
&= \lim_{t\to\infty} \sum_{k\in\mathcal{K}} \left(\frac{1}{t} r_k \Gamma^n_{\pi^n,k}(nt) + \frac{1}{n^2 t} \int_0^{nt} \left(h_k\left(X^n_{\pi^n,k}(s)^+\right) + p_k\left(X^n_{\pi^n,k}(s)^-\right)\right) ds \right) \\
&= \lim_{t\to\infty} \frac{1}{t} \sum_{k\in\mathcal{K}} \left(r_k \bar{\Gamma}^n_{\pi^n,k}(t) + \int_0^t \left(h_k\left(\bar{X}^n_{\pi^n,k}(s)^+\right) + p_k\left(\bar{X}^n_{\pi^n,k}(s)^-\right)\right) ds \right),
\end{aligned}$$

which recovers (2.2).

For each system, we simulate the long-run average cost under 1) the fluid-translated BSI policy, 2) the MDP Idle-Fix policy, and 3) the MDP General policy. Table 6 lists the parameters for the fluid-translated BSI policy, costs under the three policies, and the optimality gaps of the fluid-translated BSI policy relative to the two MDP-based policies. We observe that the optimality gap of the fluid-translated BSI policy relative to the the MDP Idle-Fix policy is smaller than the gap relative to the MDP General policy. This is well expected as the MDP Idle-Fix policies are more restricted than the MDP General policies. Note that for the four systems in Set 2, the costs under the MDP Idle-Fix and MDP General policies are indistinguishable; this is because the facility idles for an almost negligible amount of time under the MDP General policy for this set of parameters. We also observe that in all cases, the optimality gap decreases the n increases.

Table 5 **MDP parameters**

	Scale	Parameters											
	n	λ_1	λ_2	μ_1	μ_2	s_1	s_2	h_1	h_2	p_1	p_2	r_1	r_2
Set 1	1	0.5	0.6	3	3	0.2	0.4	1	2	3	6	1	1
	5	0.5	0.6	3	3	1	2	1	2	3	6	25	25
	10	0.5	0.6	3	3	2	4	1	2	3	6	100	100
	15	0.5	0.6	3	3	3	6	1	2	3	6	225	225
Set 2	1	0.6	0.5	3	3	0.1	0.3	1.3	3	4.6	4.2	0	0
	5	0.6	0.5	3	3	0.5	1.5	1.3	3	4.6	4.2	0	0
	10	0.6	0.5	3	3	1	3	1.3	3	4.6	4.2	0	0
	15	0.6	0.5	3	3	1.5	4.5	1.3	3	4.6	4.2	0	0
Set 3	1	2	2	8	8.5	0.1	0.3	1	2	4	8	0.5	1
	5	2	2	8	8.5	0.5	1.5	1	2	4	8	12.5	25
	10	2	2	8	8.5	1	3	1	2	4	8	50	100
	15	2	2	8	8.5	1.5	4.5	1	2	4	8	112.5	225

Table 6 **MDP costs**

	Scale	Fluid-translated policy				MDP Idle-Fix		MDP General	
	n	b_1	b_2	Total idle time	Cost	Cost	Optimality gap	Cost	Optimality gap
Set 1	1	0.6	0.7	0.6	5.0	4.9	1.0%	3.6	38.0%
	5	3.1	3.5	3.2	14.1	14.0	0.8%	11.8	19.3%
	10	6.2	7.1	6.5	24.8	24.6	0.7%	22.4	10.8%
	15	9.2	10.6	9.7	35.1	35.0	0.3%	32.8	6.9%
Set 2	1	0.2	0.2	0.0	3.8	3.1	22.9%	3.1	22.9%
	5	1.2	0.8	0.0	5.6	5.2	7.5%	5.2	7.5%
	10	2.4	1.5	0.0	8.0	7.7	4.5%	7.7	4.5%
	15	3.5	2.3	0.0	9.7	9.7	0.0%	9.7	0.0%
Set 3	1	1.1	1.1	0.1	9.5	7.7	22.7%	6.6	44.1%
	5	5.4	5.5	0.3	22.7	22.3	1.9%	20.7	9.6%
	10	10.9	11.1	0.7	40.0	39.6	1.1%	37.8	5.7%
	15	16.3	16.6	1.0	56.7	56.2	1.0%	55.1	2.8%

In what follows, we examine the solutions to MDP Idle-Fix and MDP General. The goal is to gain more insights into the structure of the optimal policies.

6.1. MDP Idle-Fix

The state space for MDP Idle-Fix is $(X_1, X_2, Z) \in \mathbb{Z} \times \mathbb{Z} \times \{1, 2\}$, where $Z = i$ means that the facility is at product i. Note that at each decision epoch, if $Z = 1$, the facility can choose to keep producing the same product, or to idle and switch to the next product. If the facility chooses to idle and switch, it does not produce any product for $w_*/2 + s_2$ amount of time, after which the location process Z enters state 2.

Figure 3 illustrates the MDP Idle-Fix policies for the four systems in Set 1. (Policies in the other sets follow similar structure.) To facilitate comparison, we add the fluid-translated BSI policy to each figure, with the black solid line marking the states (i.e., base-stock levels) at which the facility switches from producing the current product to idling and setting up for the next one. We observe that while the MDP Idle-Fix policy for $n = 1$ is quite different from the BSI policy, the two policies are very close to each other for $n \geq 5$. In particular, we observe convergence of the MDP policy to the BSI policy as n increases.

Figure 3 **MDP Idle-Fix policies**

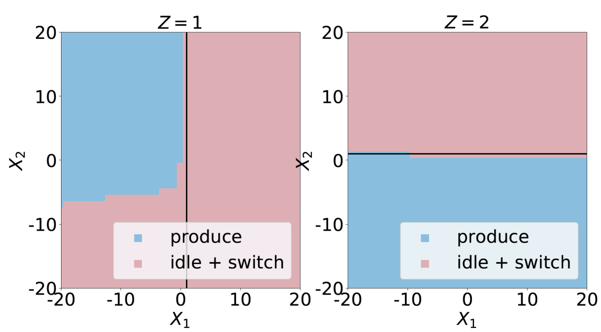

(a) $n = 1$

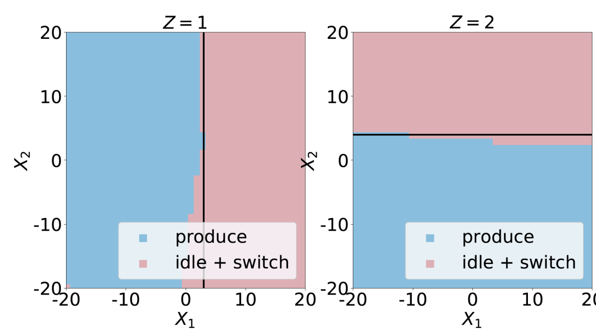

(b) $n = 5$

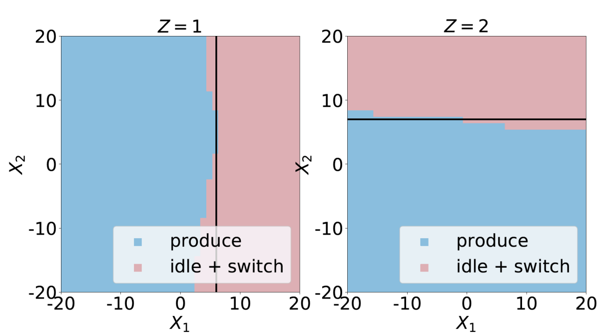

(c) $n = 10$

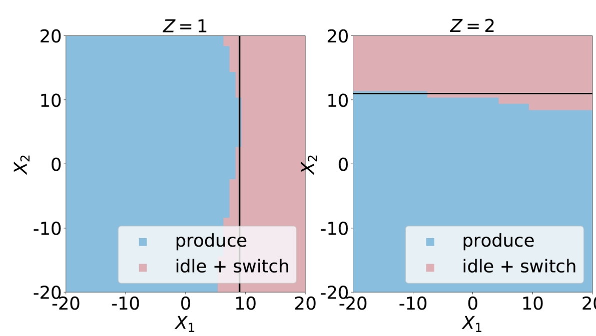

(d) $n = 15$

6.2. MDP General

Different from MDP Idle-Fix where the idle-time vector is pre-specified, MDP General solves for both the optimal lot scheduling policy and the state-dependent idling policy. In this case, the MDP has expanded state space $(X_1, X_2, Z) \in \mathbb{Z} \times \mathbb{Z} \times \{1, 2, 1^I, 2^I\}$. Here, $Z = 1$ indicates that the facility is at product 1, and can choose to keep producing the same product, start idling, or set up for product 2 for a total setup time s_2. If the facility chooses to idle from state $Z = 1$, the location process Z enters state $Z = 1^I$ immediately. When $Z = 1^I$, the actions available are either to keep idling (then Z stays at state 1^I) or to set up for product 2 (then Z changes to state 2 after s_2 amount of time). Similar dynamic holds for state $Z = 2$ and $Z = 2^I$.

Figure 4 illustrates the MDP General policies for $n = 1, 5, 10, 15$. In each figure with $Z = i$, $i = 1, 2$, the black solid line marks the states (i.e., base-stock levels) at which the fluid-translated BSI policy switches from producing the current product to idling and setting up for the next product. We observe that the two policies are substantially different even for large values of n, e.g., $n = 10, 15$.

One reason for the disparity between the fluid-translated and MDP General policies is that there can be multiple controls, all leading to the same asymptotically optimal PE. To see this, we translate the MDP General policy for the HDS by applying the MDP policy corresponding to state $(\lceil x_1 n \rfloor, \lceil x_2 n \rfloor, z)$ if the HDS is at state (x_1, x_2, z). Figure 5 plots the PE under the translated MDP General policy for $n = 1, 10, 20$. We see that as n increases, the PE under the translated MDP General policy converges to the optimal fluid PE-candidate (marked in blue). This explains the diminishing optimality gap of the BSI policy relative to the MDP General policy in Table 6.

Lastly, we comment that policies with state-dependent idling behavior are outside the set of admissible controls in our theoretical analysis. Nevertheless, under the MDP General policy, the system dynamic is still likely to converge to a stable HDS as the setup times increase. Thus, the fluid-translated BSI policy may be asymptotically optimal among a large class of controls, i.e., any sequence of controls under which the limiting fluid process converges to a PE.

7. Summary and Future Research

In this paper, we prove that the BSI policy is asymptotically optimal for the SELSP as the setup times increase to infinity if either of the two conditions holds: (i) The holding and backlog costs grow at most at a polynomial rate (of any degree), the production time distributions possess finite m.g.f.'s in a neighborhood of zero, and the setup time distributions possess finite m.g.f.'s on the positive real line; (ii) the holding and backlog costs are linear, and the production and setup time distributions possess finite variances. Our analysis leverages a novel theoretical framework developed in Hu et al. (2020a) to derive asymptotically optimal control for non-stationary stochastic networks. Our results provide a rigorous theoretical justification for employing the BSI policy, which has been widely studied in the literature and implemented in practice. Moreover, the deterministic

Figure 4 **MDP General policies**

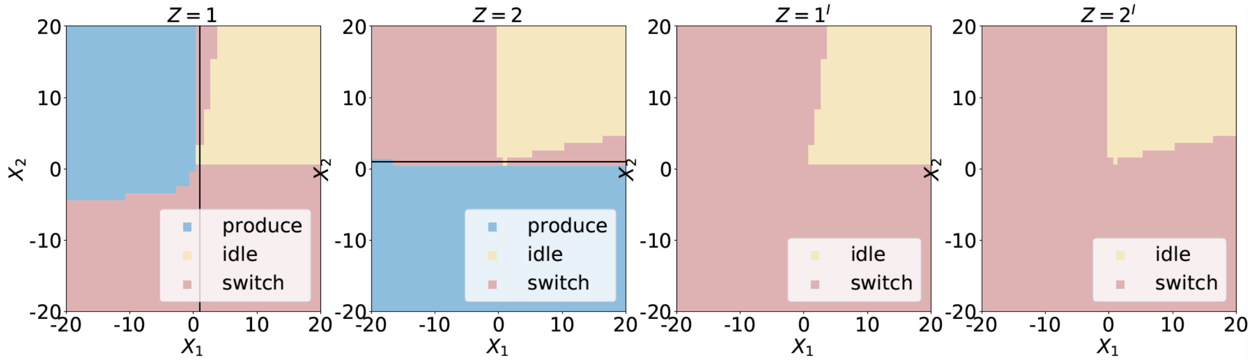

(a) $n=1$

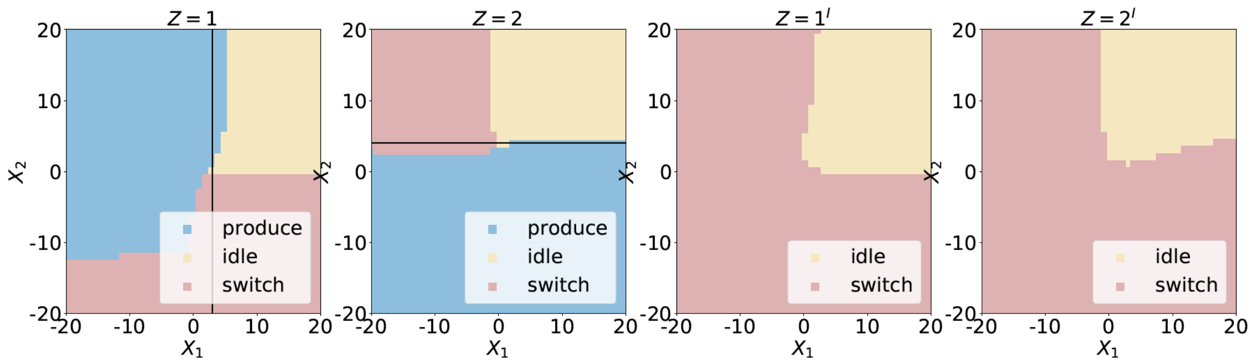

(b) $n=5$

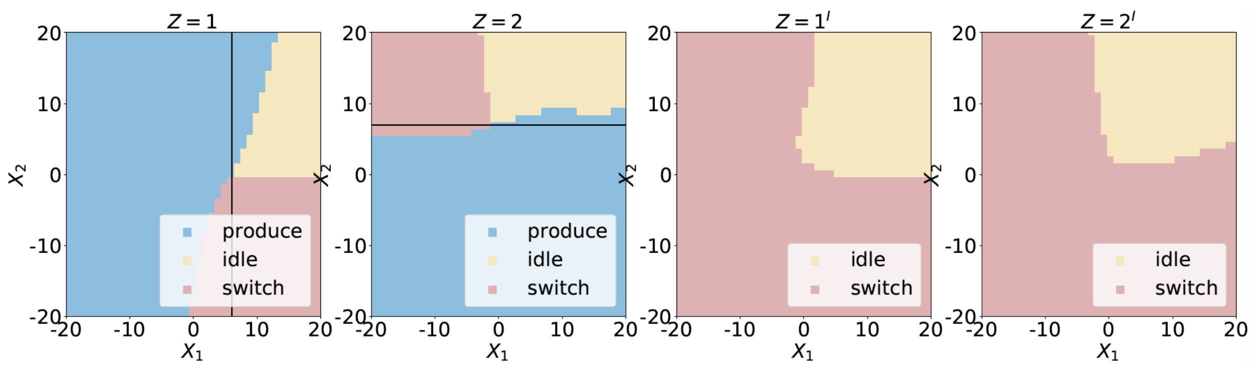

(c) $n=10$

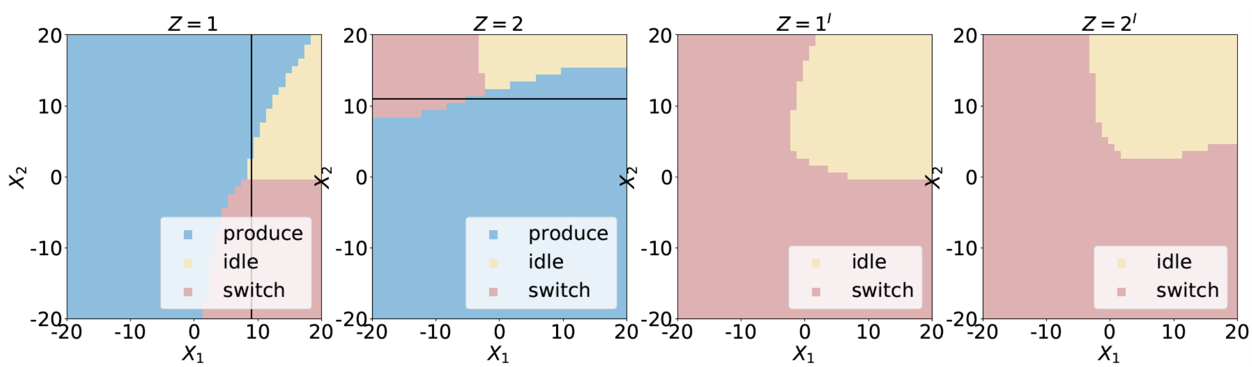

(d) $n=15$

Figure 5 **PE under the translated MDP General policy for the HDS**

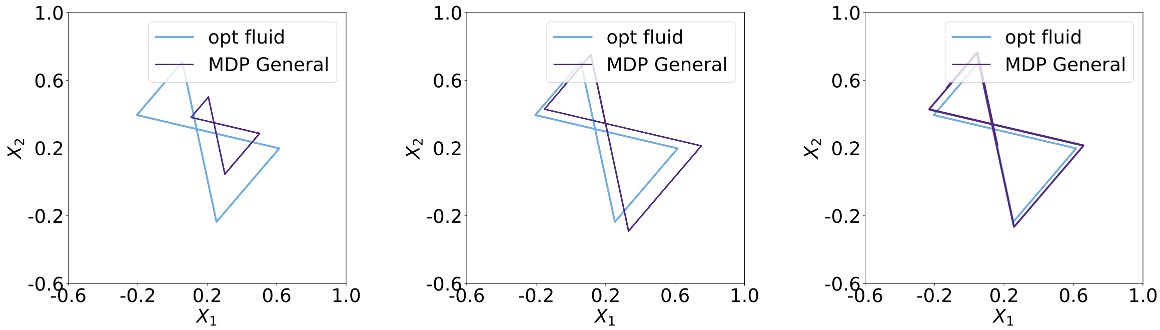

(a) $n = 1$ (b) $n = 10$ (c) $n = 20$

optimization problem that solves for the parameters of the asymptotically optimal BSI policy extends beyond the conventional common cycle approach for the ESLP. Numerical experiments suggest that a simple BSI policy derived from a deterministic optimization problem can achieve robust, and near optimal performance, even when the setup times are moderate.

References

Adelman, D., C. Barz. 2014a. A price-directed heuristic for the economic lot scheduling problem. *IIE Transactions* **46**(12) 1343–1356.

Adelman, D., C. Barz. 2014b. A unifying approximate dynamic programming model for the economic lot scheduling problem. *Mathematics of Operations Research* **39**(2) 374–402.

Allahverdi, A., J. N. Gupta, T. Aldowaisan. 1999. A review of scheduling research involving setup considerations. *Omega* **27**(2) 219–239.

Bomberger, E. E. 1966. A dynamic programming approach to a lot size scheduling problem. *Management science* **12**(11) 778–784.

Bourland, K. E., C. A. Yano. 1994. The strategic use of capacity slack in the economic lot scheduling problem with random demand. *Management Science* **40**(12) 1690–1704.

Coffman Jr, E., A. Puhalskii, M. Reiman. 1995. Polling systems with zero switchover times: a heavy-traffic averaging principle. *The Annals of Applied Probability* 681–719.

Coffman Jr, E. G., A. A. Puhalskii, M. I. Reiman. 1998. Polling systems in heavy traffic: A bessel process limit. *Mathematics of Operations Research* **23**(2) 257–304.

Das, T. K., A. Gosavi, S. Mahadevan, N. Marchalleck. 1999. Solving semi-markov decision problems using average reward reinforcement learning. *Management Science* **45**(4) 560–574.

Dobson, G. 1992. The cyclic lot scheduling problem with sequence-dependent setups. *Operations Research* **40**(4) 736–749.

Elmaghraby, S. E. 1978. The economic lot scheduling problem (elsp): review and extensions. *Management Science* **24**(6) 587–598.

Federgruen, A., Z. Katalan. 1994. Approximating queue size and waiting time distributions in general polling systems. *Queueing Systems* **18**(3-4) 353–386.

Federgruen, A., Z. Katalan. 1996a. Customer waiting-time distributions under base-stock policies in single-facility multi-item production systems. *Naval Research Logistics (NRL)* **43**(4) 533–548.

Federgruen, A., Z. Katalan. 1996b. The stochastic economic lot scheduling problem: cyclical base-stock policies with idle times. *Management Science* **42**(6) 783–796.

Feoktistova, V., A. Matveev, E. Lefeber, J. Rooda. 2012. Designs of optimal switching feedback decentralized control policies for fluid queueing networks. *Mathematics of Control, Signals, and Systems* **24**(4) 477–503.

Fricker, C., M. R. Jaibi. 1994. Monotonicity and stability of periodic polling models. *Queueing systems* **15**(1-4) 211–238.

Gallego, G. 1990. Scheduling the production of several items with random demands in a single facility. *Management Science* **36**(12) 1579–1592.

Gallego, G. 1994. When is a base stock policy optimal in recovering disrupted cyclic schedules? *Naval Research Logistics (NRL)* **41**(3) 317–333.

Gallego, G., R. Roundy. 1988. The extended economic lot scheduling problem. Tech. rep., Cornell University Operations Research and Industrial Engineering.

Gascon, A., R. C. Leachman, P. Lefrançois. 1994. Multi-item, single-machine scheduling problem with stochastic demands: a comparison of heuristics. *THE INTERNATIONAL JOURNAL OF PRODUCTION RESEARCH* **32**(3) 583–596.

Grasman, S. E., T. L. Olsen, J. R. Birge. 2008. Setting basestock levels in multi-product systems with setups and random yield. *IIE Transactions* **40**(12) 1158–1170.

Hanssmann, F. 1962. *Operations research in production and inventory control*. Wiley.

Hsu, W.-L. 1983. On the general feasibility test of scheduling lot sizes for several products on one machine. *Management science* **29**(1) 93–105.

Hu, Y., J. Dong, O. Perry. 2020a. Asymptotic optimality of the binomial-exhaustive policy for polling systems with large switchover times. URL `http://www.columbia.edu/~yh2987/Files/Polling.pdf`. (accessed on 2021-05-28).

Hu, Y., J. Dong, O. Perry. 2020b. On the existence and asymptotic approximation of moments for polling systems under the binomial-exhaustive policy. URL `http://www.columbia.edu/~yh2987/Files/Moments.pdf`. (accessed on 2021-05-28).

Jones, P. C., R. R. Inman. 1989. When is the economic lot scheduling problem easy? *IIE transactions* **21**(1) 11–20.

Krieg, G. N., H. Kuhn. 2004. Analysis of multi-product kanban systems with state-dependent setups and lost sales. *Annals of Operations Research* **125**(1-4) 141–166.

Markowitz, D. M., M. I. Reiman, L. M. Wein. 2000. The stochastic economic lot scheduling problem: heavy traffic analysis of dynamic cyclic policies. *Operations Research* **48**(1) 136–154.

Markowitz, D. M., L. M. Wein. 2001. Heavy traffic analysis of dynamic cyclic policies: a unified treatment of the single machine scheduling problem. *Operations Research* **49**(2) 246–270.

Maxwell, W. L. 1964. The scheduling of economic lot sizes. *Naval Research Logistics Quarterly* **11**(2) 89–124.

Olsen, T. L. 2001. Limit theorems for polling models with increasing setups. *Probability in the Engineering and Informational Sciences* **15**(1) 35–55.

Sox, C. R., P. L. Jackson, A. Bowman, J. A. Muckstadt. 1999. A review of the stochastic lot scheduling problem. *International Journal of Production Economics* **62**(3) 181–200.

Sox, C. R., J. A. Muckstadt. 1997. Optimization-based planning for the stochastic lot-scheduling problem. *IIE transactions* **29**(5) 349–357.

Tijms, H. C. 2003. *A first course in stochastic models*. John Wiley and sons.

Van der Mei, R. 1999. Delay in polling systems with large switch-over times. *Journal of applied probability* **36**(1) 232–243.

Whitt, W. 2002. *Stochastic-process limits: an introduction to stochastic-process limits and their application to queues*. Springer Science & Business Media.

Winands, E. M. 2007. On polling systems with large setups. *Operations research letters* **35**(5) 584–590.

Winands, E. M. 2011. Branching-type polling systems with large setups. *OR Spectrum* **33**(1) 77–97.

Winands, E. M., I. J. Adan, G.-J. van Houtum. 2011. The stochastic economic lot scheduling problem: A survey. *European Journal of Operational Research* **210**(1) 1–9.

Xin, L., D. A. Goldberg. 2018. Asymptotic optimality of tailored base-surge policies in dual-sourcing inventory systems. *Management Science* **64**(1) 437–452.

Zipkin, P. H. 1991. Computing optimal lot sizes in the economic lot scheduling problem. *Operations Research* **39**(1) 56–63.

Appendix A: Full HDS Representation

In Section 3, we characterize the fluid model as an HDS described by

$$\dot{x}(t) = f(z(t)) \quad z(t) = g(x(t), x(a(t-)), z(t-)), \quad a(t) = h(x(t), x(a(t-)), z(t-)),$$

where $f : \mathcal{Z} \to \mathbb{R}^K$, $g : \mathbb{R}^K \times \mathbb{R}^K \times \mathcal{Z} \to \mathcal{Z}$, and $h : \mathbb{R}^K \times \mathbb{R}^K \times \mathcal{Z} \to \mathbb{R}_+$. In particular, the function f regulates the drift of the inventory level process, g determines the server location, and h specifies the most recent polling epoch. We now characterize f, g, and h one by one, utilizing the lot sizing and idling functions $(\phi_k, \eta_k, k \in \mathcal{K})$ specified by the control. Noticeably, the functional forms of f, g, and h allow zero setup times (subject to strictly positive total setup time in a cycle, i.e., $s > 0$), zero idle times, and zero productions.

(I) It is easy to see (and is mentioned in Section 3) that the function f is given by

$$f_k(x(t), z(t)) = \begin{cases} -\lambda_k + \mu_k & \text{if } z(t) = k \\ -\lambda_k & \text{otherwise} \end{cases}, \quad k \in \mathcal{K}.$$

(II) The function g characterizes the location of the server as follows:

(i) If $z(t-) = k$ and $x_k(t) = x_k(a(t-)) + (\mu_k - \lambda_k)\phi_k\left(x(a(t-))\right)$, define

$$j_s := \min\{j > k : s_{(j \bmod K)} > 0\}$$

$$j_v := \min\{j > k : \phi_{(j \bmod K)}\left(x(a(t-))\right) > 0\}$$

$$j_w := \min\{j \geq k : \eta_{(j \bmod K)}\left(x(a(t-))\right) > 0\}.$$

(a) If $j_s = \min\{j_s, j_v, j_w\}$, then $g(x(t), x(a(t-)), z(t-)) = (j_s \bmod K)^S$.

(b) If $j_v = \min\{j_s, j_v, j_w\}$ and $j_v < j_s$, then $g(x(t), x(a(t-)), z(t-)) = (j_v \bmod K)$.

(c) If $j_w = \min\{j_s, j_v, j_w\}$, $j_w < j_s$ and $j_w < j_v$, then $g(x(t), x(a(t-)), z(t-)) = (j_w \bmod K)^I$.

(ii) If $z(t-) = k^I$ and $x_k(t) = x_k(a(t-)) + (\mu_k - \lambda_k)\phi_k\left(x(a(t-))\right) - \lambda_k \eta_k\left(x(a(t-))\right)$, define

$$j_s := \min\{j > k : s_{(j \bmod K)} > 0\}$$

$$j_v := \min\{j > k : \phi_{(j \bmod K)}\left(x(a(t-))\right) > 0\}$$

$$j_w := \min\{j > k : \eta_{(j \bmod K)}\left(x(a(t-))\right) > 0\}.$$

(a) If $j_s = \min\{j_s, j_v, j_w\}$, then $g(x(t), x(a(t-)), z(t-)) = (j_s \bmod K)^S$.

(b) If $j_v = \min\{j_s, j_v, j_w\}$ and $j_v < j_s$, then $g(x(t), x(a(t-)), z(t-)) = (j_v \bmod K)$.

(c) If $j_w = \min\{j_s, j_v, j_w\}$, $j_w < j_s$ and $j_w < j_v$, then $g(x(t), x(a(t-)), z(t-)) = (j_w \bmod K)^I$.

(iii) If $z(t-) = k^S$ and $x_k(t) = x_k(a(t-)) - \lambda_k\left(\phi_{k-1}\left(x(a(t-))\right) + \eta_{k-1}\left(x(a(t-))\right) + s_k\right)$, define

$$j_s := \min\{j > k : s_{(j \bmod K)} > 0\}$$

$$j_v := \min\{j \geq k : \phi_{(j \bmod K)}\left(x(a(t-))\right) > 0\}$$

$$j_w := \min\{j \geq k : \eta_{(j \bmod K)}\left(x(a(t-))\right) > 0\}.$$

(a) If $j_s = \min\{j_s, j_v, j_w\}$, then $g(x(t), x(a(t-)), z(t-)) = (j_s \bmod K)^S$.

(b) If $j_v = \min\{j_s, j_v, j_w\}$ and $j_v < j_s$, then $g(x(t), x(a(t-)), z(t-)) = (j_v \bmod K)$.

(c) If $j_w = \min\{j_s, j_v, j_w\}$, $j_w < j_s$ and $j_w < j_v$, then $g(x(t), x(a(t-)), z(t-)) = (j_w \bmod K)^I$.

(iv) Otherwise, $g(x(t), x(a(t-)), z(t-)) = z(t-)$.

(III) The function h updates the most recent polling epoch according to

$$h(x(t), x(a(t-)), z(t-)) = \\ \begin{cases} t & \text{if } z(t-) = k^S \text{ and } x_k(t) = x_k(a(t-)) - \lambda_k\left(\phi_{k-1}\left(x(a(t-))\right) + \eta_{k-1}\left(x(a(t-))\right) + s_k\right) \\ a(t-) & \text{otherwise.} \end{cases}$$

Appendix B: Proofs of Results in Section 3

B.1. Proof of Lemma 2

Proof: Let $v^{(m-1)}$ denote the beginning epoch of the $((m-1)L_e^\phi + 1)$th production cycle, namely, $v^{(m-1)} := u^{((m-1)L_e^\phi)}$, $m \geq 1$. Define $\tilde{T}^{(m)} := v^{(m)} - v^{(m-1)}$. By construction, $\tilde{T}^{(m)}$ contains exactly L_e^ϕ production cycles. Since x_e^ϕ is a PE, for any $\epsilon > 0$, there exists some $N_\epsilon \in \mathbb{N}$ such that for all $m \geq N_\epsilon$,

$$\|x_u^\phi(v^{(m-1)} + \cdot) - x_e^\phi(\cdot)\|_t < \epsilon \quad \text{for all } t > 0, \quad |\tilde{T}^{(m)} - \tau_e^\phi| < \epsilon, \tag{B.1}$$

and

$$\left| \int_{v^{(m-1)}}^{v^{(m)}} x_u^\phi(s) ds - \int_0^{\tau_e^\phi} x_e^\phi(s) ds \right| < \epsilon.$$

Since the PE-candidate x_e^ϕ is bounded, (B.1) implies that x_u^ϕ is also bounded. Due to the continuity of ψ_k, x_e^ϕ and x_u^ϕ, the compositions $\psi_k \circ x_e^\phi$ and $\psi_k \circ x_u^\phi$ are uniformly continuous over any compact time interval, $k \in \mathcal{K}$. It follows that there exists $M_\epsilon \geq N_\epsilon$ such that for all $m \geq M_\epsilon$,

$$\left| \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s)) ds - \int_0^{\tau_e^\phi} \psi(x_e^\phi(s)) ds \right| < \epsilon.$$

For any time $t \geq 0$, define $M(t) := \max\{m \geq 1 : v^{(m)} \leq t\}$. In addition, let $r := L_e^\phi \sum_{k=1}^K r_k$. Then, the time-average cost of x_u^ϕ can be written as

$$\begin{aligned} &\frac{1}{t}\left(M(t)r + \sum_{\ell=1}^{L_e^\phi} \sum_{k=1}^K r_k \mathbf{1}_{\{a_k^{((M(t)-1)L_e^\phi+\ell)} \leq t\}} + \int_0^t \psi(x_u^\phi(s)) ds \right) \\ &= \frac{M(t)}{t} r + \frac{1}{t} \sum_{m=1}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s)) ds + \frac{1}{t} \left(\sum_{\ell=1}^{L_e^\phi} \sum_{k=1}^K r_k \mathbf{1}_{\{a_k^{((M(t)-1)L_e^\phi+\ell)} \leq t\}} + \int_{v^{(M(t))}}^t \psi(x_u^\phi(s)) ds \right). \end{aligned} \tag{B.2}$$

In what follows, we analyze each of the three terms on the right-hand side of (B.2). We shall consider t sufficiently large, so that $M(t) > M_\epsilon$.

(I) For the first term on the right-hand side of (B.2), observe that

$$\begin{aligned} \frac{t}{M(t)} &= \frac{1}{M(t)} \left(\sum_{m=1}^{M_\epsilon - 1} \tilde{T}^{(m)} + \sum_{m=M_\epsilon}^{M(t)} \tilde{T}^{(m)} + \left(t - v^{(M(t))}\right) \right) \\ &= \frac{1}{M(t)} \sum_{m=M_\epsilon}^{M(t)} \tilde{T}^{(m)} + o(1) \\ &\geq \frac{M(t) - M_\epsilon}{M(t)} (\tau_e^\phi - \epsilon) + o(1) \to \tau_e^\phi - \epsilon \quad \text{as } t \to \infty. \end{aligned} \tag{B.3}$$

Then (B.3) implies that

$$\liminf_{t \to \infty} \frac{t}{M(t)} \geq \tau_e^\phi - \epsilon.$$

Similarly, we can show that

$$\limsup_{t \to \infty} \frac{t}{M(t)} \leq \tau_e^\phi + \epsilon.$$

(II) The second term on the right-hand side of (B.2) satisfies

$$\frac{1}{t} \sum_{m=1}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s)) ds = \frac{1}{t} \sum_{m=1}^{M_\epsilon - 1} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s)) ds + \frac{1}{t} \sum_{m=M_\epsilon}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s)) ds.$$

For fixed $\epsilon > 0$, M_ϵ is fixed, so the first term on the right-hand side of the equality converges to 0 as $t \to \infty$. Applying (B.3) for the second term gives that

$$\begin{aligned}
\frac{1}{t}\sum_{m=M_\epsilon}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s))ds &= \frac{M(t)}{t}\frac{1}{M(t)}\sum_{m=M_\epsilon}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s))ds \\
&\le \frac{M(t)}{t}\frac{1}{M(t)}\sum_{m=M_\epsilon}^{M(t)} \left(\int_0^{\tau_e^\phi} \psi(x_e^\phi(s))ds + \epsilon\right) \\
&\le \left(\frac{1}{\tau_e^\phi - \epsilon} + o(1)\right)\left(\int_0^{\tau_e^\phi} \psi(x_e^\phi(s))ds + \epsilon\right) \\
&\to \frac{1}{\tau_e^\phi - \epsilon}\left(\int_0^{\tau_e^\phi} \psi(x_e^\phi(s))ds + \epsilon\right) \qquad \text{as } t \to \infty.
\end{aligned} \tag{B.4}$$

It follows from (B.4) that

$$\limsup_{t\to\infty} \frac{1}{t}\sum_{m=1}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s))ds \le \frac{1}{\tau_e^\phi - \epsilon}\left(\int_0^{\tau_e^\phi} \psi(x_e^\phi(s))ds + \epsilon\right).$$

Similarly, we can show that

$$\liminf_{t\to\infty} \frac{1}{t}\sum_{m=1}^{M(t)} \int_{v^{(m-1)}}^{v^{(m)}} \psi(x_u^\phi(s))ds \ge \frac{1}{\tau_e^\phi + \epsilon}\left(\int_0^{\tau_e^\phi} \psi(x_e^\phi(s))ds - \epsilon\right).$$

(III) For the third term on the right-hand side of (B.2), note that by definition of $M(t)$,

$$\sum_{\ell=1}^{L_e^\phi}\sum_{k=1}^{K} r_k \mathbf{1}_{\{a_k^{((M(t)-1)L_e^\phi+\ell)} \le t\}} \le r.$$

In addition, we have that $0 \le t - v^{(M(t))} \le \tilde{T}^{(M(t)+1)}$, and $\tilde{T}^{(M(t)+1)}$ is bounded due to (B.1). Thus,

$$\lim_{t\to\infty} \frac{1}{t}\left(\sum_{\ell=1}^{L_e^\phi}\sum_{k=1}^{K} r_k \mathbf{1}_{\{a_k^{((M(t)-1)L_e+\ell)} \le t\}} + \int_{v^{(M(t))}}^{t} \psi(x_u^\phi(s))ds\right) = 0.$$

Combining **(I)**, **(II)**, and **(III)**, together with the fact that ϵ is arbitrary, we have

$$\frac{1}{t}\left(M(t)r + \sum_{\ell=1}^{L_e^\phi}\sum_{k=1}^{K} r_k \mathbf{1}_{\{a_k^{((M(t)-1)L_e^\phi+\ell)} \le t\}} + \int_0^t \psi(x_u^\phi(s))ds\right) \to \frac{r}{\tau_e^\phi} + \frac{1}{\tau_e^\phi}\int_0^{\tau_e^\phi} \psi(x_e^\phi(s))ds \qquad \text{as } t \to \infty,$$

where the limit on right-hand side is equal to the cost of the PE-candidate x_e^ϕ over one period. Q.E.D.

B.2. Proof of Proposition 1

Proof: We carry out the proof in two steps: First, we consider piecewise linear holding and backlog costs that satisfy Assumption 2, and prove Proposition 1 in this context. In the second step, we generalize the result to general (i.e., not necessarily piecewise linear) cost functions that satisfy Assumption 2 by approximating bounded continuous functions over compact intervals with piecewise linear functions to arbitrary precision.

Step 1. Piecewise linear costs. For each $k \in \mathcal{K}$, we assume that the holding cost function $h_k : \mathbb{R}_+ \to \mathbb{R}_+$ is piecewise linear with $N_{h_k} \in \mathbb{N}$ irregular points. Namely, there exists some constant $h_k^{(0)} \in \mathbb{R}_+$, coefficients

$h_k^{(1)}, h_k^{(2)}, ..., h_k^{(N_{h_k})}$ in $\mathbb{R}_+$, thresholds $0 = \alpha_k^{(1)} < \alpha_k^{(2)} \cdots < \alpha_k^{(N_{h_k})}$ in $\mathbb{R}_+$, and $\alpha_k^{(N_{h_k}+1)} = \infty$, such that for any $x \in \mathbb{R}_+$,

$$h_k(x) = \begin{cases} h_k^{(1)}(x - \alpha_k^{(1)}) + h_k^{(0)} & \text{if } \alpha_k^{(1)} \le x < \alpha_k^{(2)} \\ h_k^{(2)}(x - \alpha_k^{(2)}) + h_k^{(1)}(\alpha_k^{(2)} - \alpha_k^{(1)}) + h_k^{(0)} & \text{if } \alpha_k^{(2)} \le x < \alpha_k^{(3)} \\ \vdots & \\ h_k^{(N_{h_k})}(x - \alpha_k^{(N_{h_k})}) + \sum_{i=1}^{N_k - 1} h_k^{(i)}(\alpha_k^{(i+1)} - \alpha_k^{(i)}) + h_k^{(0)} & \text{if } \alpha_k^{(N_{h_k})} \le x < \alpha_k^{(N_{h_k}+1)}. \end{cases}$$

Similarly, we assume that the backlog cost function $p_k : \mathbb{R}_+ \to \mathbb{R}_+$ is piecewise linear with $N_{p_k} \in \mathbb{N}$ irregular points, and is associated with constant $p_k^{(0)} \in \mathbb{R}_+$, coefficients $p_k^{(1)}, p_k^{(2)}, ..., p_k^{(N_{p_k})}$ in $\mathbb{R}_+$, thresholds $0 = \beta_k^{(1)} < \beta_k^{(2)} \cdots < \beta_k^{(N_{p_k})}$ in $\mathbb{R}_+$, and $\beta_k^{(N_{p_k}+1)} = \infty$.

Let x_e be an L-cycle PE-candidate whose period is strictly larger than one production cycle. We will show that based on x_e, we can construct a one-cycle PE-candidate y_e whose time-average cost is no worse than that of x_e. To do this, we note that the objective function of the fluid optimization problem (3.4) consists of three parts: the time-average setup, holding, and backlog costs. In particular, it takes the form

$$\begin{aligned} &\frac{1}{\tau_L} \sum_{k \in \mathcal{K}} \left(r_k L + \int_0^{\tau_L} \left(h_k \left(x_{e,k}(s)^+ \right) + p_k \left(x_{e,k}(s)^- \right) \right) ds \right) \\ &= \sum_{k \in \mathcal{K}} \left(\frac{1}{\tau_L} r_k L + \frac{1}{\tau_L} \int_0^{\tau_L} \left(h_k \left(x_{e,k}(s)^+ \right) ds \right) + \frac{1}{\tau_L} \int_0^{\tau_L} \left(p_k \left(x_{e,k}(s)^- \right) ds \right) \right). \end{aligned}$$

In what follows, we will construct y_e such that for each product $k \in \mathcal{K}$, $y_{e,k}$ **(I)** (weakly) reduces the the time-average holding cost, **(II)** (weakly) decreases the time-average backlog cost, and **(III)** maintains the same time-average setup cost.

(I) For $k \in \mathcal{K}$, recall that N_{h_k} is the total number of discontinuity points associated with the holding cost function h_k. The time-average holding cost of $x_{e,k}$ can be rewritten as

$$\begin{aligned} &\frac{1}{\tau_L} \int_0^{\tau_L} \left(h_k \left(x_{e,k}(s)^+ \right) \right) ds \\ &= \frac{1}{\tau_L} \int_0^{\tau_L} \Bigg[\left(h_k^{(1)}(x_{e,k}(s) - \alpha_k^{(1)}) + h_k^{(0)} \right) \mathbf{1}_{\{\alpha_k^{(1)} \le x_{e,k}(s) < \alpha_k^{(2)}\}} \\ &\quad + \left(h_k^{(2)}(x_{e,k}(s) - \alpha_k^{(2)}) + h_k^{(1)}(\alpha_k^{(2)} - \alpha_k^{(1)}) + h_k^{(0)} \right) \mathbf{1}_{\{\alpha_k^{(2)} \le x_{e,k}(s) < \alpha_k^{(3)}\}} ds + \cdots \\ &\quad + \left(h_k^{(N_k)}(x_{e,k}(s) - \alpha_k^{(N_k)}) + \sum_{\ell=1}^{N_k - 1} h_k^{(\ell)}(\alpha_k^{(\ell+1)} - \alpha_k^{(\ell)}) + h_k^{(0)} \right) \mathbf{1}_{\{\alpha_k^{(N_k)} \le x_{e,k}(s) < \alpha_k^{(N_k+1)}\}} \Bigg] ds \\ &= \frac{1}{\tau_L} \int_0^{\tau_L} \Bigg[h_k^{(1)} \left(\min\{x_{e,k}(s), \alpha_k^{(2)}\} - \alpha_k^{(1)} \right) \mathbf{1}_{\{x_{e,k}(s) \ge \alpha_k^{(1)}\}} \\ &\quad + h_k^{(2)} \left(\min\{x_{e,k}(s), \alpha_k^{(3)}\} - \alpha_k^{(2)} \right) \mathbf{1}_{\{x_{e,k}(s) \ge \alpha_k^{(2)}\}} + \cdots \\ &\quad + h_k^{(N_k)} \left(\min\{x_{e,k}(s), \alpha_k^{(N_k+1)}\} - \alpha_k^{(N_k)} \right) \mathbf{1}_{\{x_{e,k}(s) \ge \alpha_k^{(N_k)}\}} \Bigg] ds + h_k^{(0)} \\ &= \frac{1}{\tau_L} \sum_{\ell=1}^{N_k} h_k^{(\ell)} \int_0^{\tau_L} \left(\min\{x_{e,k}(s), \alpha_k^{(\ell+1)}\} - \alpha_k^{(\ell)} \right) \mathbf{1}_{\{x_{e,k}(s) \ge \alpha_k^{(\ell)}\}} ds + h_k^{(0)}. \end{aligned} \tag{B.5}$$

For each summand in (B.5), define

$$A^{(\ell)}(x_{e,k}) := \int_0^{\tau_L} \left(\min\{x_{e,k}(s), \alpha_k^{(\ell+1)}\} - \alpha_k^{(\ell)} \right) \mathbf{1}_{\{x_{e,k}(s) \ge \alpha_k^{(\ell)}\}} ds, \quad \ell = 1, ..., N_{h_k},$$

so that the time-average holding cost can be written as

$$\frac{1}{\tau_L}\int_0^{\tau_L}\left(h_k\left(x_{e,k}(s)^+\right)\right)ds = \frac{1}{\tau_L}\sum_{\ell=1}^{N_{h_k}} h_k^{(\ell)} A^{(\ell)}\left(x_{e,k}\right) + h_k^{(0)}. \tag{B.6}$$

We further partition the area $A^{(\ell)}(x_{e,k})$ into L sub-areas in each production cycle. In particular, let $a^{(\ell,m)}(x_{e,k})$ denote the area beneath the inventory level $x_{e,k}$, and between the horizontal lines at values $\alpha_k^{(\ell)}$ and $\alpha_k^{(\ell+1)}$ over the mth production cycle within the time $[0,\tau_L]$, namely,

$$a^{(\ell,m)}\left(x_{e,k}\right) := \int_{u_e^{(m-1)}}^{u_e^{(m)}}\left(\min\{x_{e,k}(s),\alpha_k^{(\ell+1)}\} - \alpha_k^{(\ell)}\right)\mathbf{1}_{\{x_{e,k}(s)\geq\alpha_k^{(\ell)}\}}ds, \quad m=1,...,L, \quad \ell=1,...,N_{h_k},$$

where recall that $u_e^{(m-1)}$ denotes the beginning epoch of the mth production cycle for x_e. Hence, following (B.6), the time-average holding cost is equivalent to

$$\frac{1}{\tau_L}\int_0^{\tau_L}\left(h_k\left(x_{e,k}(s)^+\right)\right)ds = \frac{1}{\tau_L}\sum_{\ell=1}^{N_{h_k}} h_k^{(\ell)}\sum_{m=1}^{L} a^{(\ell,m)}\left(x_{e,k}\right) + h_k^{(0)}. \tag{B.7}$$

Corresponding to $a^{(\ell,m)}\left(x_{e,k}\right)$, define $\tau^{(\ell,m)}\left(x_{e,k}\right)$ as the total time in the mth production cycle during which the trajectory of $x_{e,k}$ stays above the horizontal line at $\alpha_k^{(\ell)}$, namely,

$$\tau^{(\ell,m)}\left(x_{e,k}\right) := \int_{u_e^{(m-1)}}^{u_e^{(m)}}\mathbf{1}_{\{x_{e,k}(s)\geq\alpha_k^{(\ell)}\}}ds, \quad m=1,...,L, \quad \ell=1,...,N_{h_k}.$$

We make two important observations on the magnitude of $a^{(\ell,m)}\left(x_{e,k}\right)$ and $\tau^{(\ell,m)}(x_{e,k})$:

(i) Let $\tau_+(x_{e,k})$ denote the total time $x_{e,k}$ stays non-negative, namely,

$$\tau_+(x_{e,k}) := \int_0^{\tau_L}\mathbf{1}_{\{x_{e,k}(s)\geq 0\}}ds.$$

Define

$$M_{h_k}^{(\ell)} := \max\left\{\tau_+(x_{e,k}) - L\left(\frac{\alpha_k^{(\ell)}}{\mu_k-\lambda_k}+\frac{\alpha_k^{(\ell)}}{\lambda_k}\right),0\right\}, \quad \ell=1,...,N_{h_k}. \tag{B.8}$$

The total time $x_{e,k}$ stays above level $\alpha_k^{(\ell)}$ over the time $[0,\tau_L]$ is lower bounded by

$$\sum_{m=1}^{L}\tau^{(\ell,m)}(x_{e,k}) \geq M_{h_k}^{(\ell)}, \quad \ell=1,...,N_{h_k}. \tag{B.9}$$

To see (B.9), note that during the total $\tau_+(x_{e,k})$ time units when $x_{e,k}$ is non-negative, $x_{e,k}$ increases at rate $\mu_k-\lambda_k$, decreases at rate λ_k, and switches slope from $\mu_k-\lambda_k$ to $-\lambda_k$ at most L times. Thus, for the non-negative part of $x_{e,k}$, the maximum amount of time the trajectory stays below level $\alpha_k^{(\ell)}$ is given by $L\left(\frac{\alpha_k^{(\ell)}}{\mu_k-\lambda_k}+\frac{\alpha_k^{(\ell)}}{\lambda_k}\right)$. In other words, the non-negative part of $x_{e,k}$ stays above level $\alpha_k^{(\ell)}$ for at least $\tau_+(x_{e,k}) - L\left(\frac{\alpha_k^{(\ell)}}{\mu_k-\lambda_k}+\frac{\alpha_k^{(\ell)}}{\lambda_k}\right)$ time units. Reflecting this lower bound at 0 gives the expression of $M_{h_k}^{(\ell)}$, and (B.9) follows.

(ii) The area $a^{(\ell,m)}\left(x_{e,k}\right)$ is lower bounded by

$$\begin{aligned}a^{(\ell,m)}\left(x_{e,k}\right) \geq{}& \frac{1}{2}\left(2\tau^{(\ell,m)}(x_{e,k}) - \frac{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}}{\lambda_k} - \frac{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}}{\mu_k-\lambda_k}\right)\left(\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}\right)\\ &\mathbf{1}_{\{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}<(1-\rho_k)\lambda_k\tau^{(\ell,m)}(x_{e,k})\}}\\ &+\frac{1}{2}(1-\rho_k)\lambda_k\left(\tau^{(\ell,m)}(x_{e,k})\right)^2\mathbf{1}_{\{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}\geq(1-\rho_k)\lambda_k\tau^{(\ell,m)}(x_{e,k})\}}, \quad m=1,...,L, \quad \ell=1,...,N_{h_k}.\end{aligned}$$

This, together with (B.9), gives that

$$\sum_{m=1}^{L} a^{(\ell,m)}\left(x_{e,k}\right) \geq L\Bigg[\frac{1}{2}\left(2M_{h_k}^{(\ell)}/L-\frac{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}}{\lambda_k}-\frac{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}}{\mu_k-\lambda_k}\right)\left(\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}\right)$$
$$\mathbf{1}_{\{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}<(1-\rho_k)\lambda_k M_{h_k}^{(\ell)}/L\}} \tag{B.10}$$
$$+\frac{1}{2}(1-\rho_k)\lambda_k\left(M_{h_k}^{(\ell)}/L\right)^2\mathbf{1}_{\{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}\geq(1-\rho_k)\lambda_k M_{h_k}^{(\ell)}/L\}}\Bigg],\quad \ell=1,...,N_{h_k}.$$

Based on $x_{e,k}$, we construct a corresponding L-cycle closed loop $y_{e,k}$ whose non-negative part satisfies the following conditions:

(a) The trajectory of $y_{e,k}$ stays non-negative for a total of $\tau_+(x_{e,k})$ time units over the time interval $[0,\tau_L]$. Namely, $y_{e,k}$ and $x_{e,k}$ spend the same amount of time above zero.

(b) The non-negative part of $y_{e,k}$ is identical (symmetric) over each production cycle, and forms a triangular closed loop that increases at rate $\mu_k-\lambda_k$, decreases at rate λ_k, and has a base of length $\tau_+(x_{e,k})/L$ (equivalently, $M_{h_k}^{(1)}/L$).

(The negative part of $y_{e,k}$ is immaterial for our current consideration and will be specified in **(II)**).

Note that steps (a) and (b) imply that in each production cycle, the part of $y_{e,k}$ that stays above level $\alpha_k^{(\ell)}$ is again a triangular closed loop that increases at rate $\mu_k-\lambda_k$, decreases at rate λ_k, and has a base whose length is given by

$$\tau^{(\ell,m)}(y_{e,k})=M_{h_k}^{(\ell)}/L \quad \text{for } M_{h_k}^{(\ell)} \text{ in (B.8)},\quad m=1,...,L,\quad \ell=1,...,N_{h_k}.$$

Thus, the total area beneath the trajectory of $y_{e,k}$, and between the horizontal lines at values $\alpha_k^{(\ell)}$ and $\alpha_k^{(\ell+1)}$ over the time $[0,\tau_L]$ satisfies

$$\sum_{m=1}^{L} a^{(\ell,m)}\left(y_{e,k}\right) = L\Bigg[\frac{1}{2}\left(2M_{h_k}^{(\ell)}/L-\frac{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}}{\lambda_k}-\frac{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}}{\mu_k-\lambda_k}\right)\left(\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}\right)$$
$$\mathbf{1}_{\{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}<(1-\rho_k)\lambda_k M_{h_k}^{(\ell)}/L\}}$$
$$+\frac{1}{2}(1-\rho_k)\lambda_k\left(M_{h_k}^{(\ell)}/L\right)^2\mathbf{1}_{\{\alpha_k^{(\ell+1)}-\alpha_k^{(\ell)}\geq(1-\rho_k)\lambda_k M_{h_k}^{(\ell)}/L\}}\Bigg] \tag{B.11}$$
$$\leq \sum_{m=1}^{L} a^{(\ell,m)}\left(x_{e,k}\right) \qquad \text{by (B.10)},\quad \ell=1,...,N_{h_k}.$$

Following (B.7) and (B.11), we get that

$$\begin{aligned}\frac{1}{\tau_L}\int_0^{\tau_L}\left(h_k\left(y_{e,k}(s)^+\right)\right)ds &= \frac{1}{\tau_L}\sum_{\ell=1}^{N_{h_k}} h_k^{(\ell)}\sum_{m=1}^{L} a^{(\ell,m)}\left(y_{e,k}\right)+h_k^{(0)}\\ &\leq \frac{1}{\tau_L}\sum_{\ell=1}^{N_{h_k}} h_k^{(\ell)}\sum_{m=1}^{L} a^{(\ell,m)}\left(x_{e,k}\right)+h_k^{(0)}\\ &= \frac{1}{\tau_L}\int_0^{\tau_L}\left(h_k\left(x_{e,k}(s)^+\right)\right)ds.\end{aligned}$$

(II) Following similar procedures as in **(I)**, we can construct the negative part of the L-cycle PE-candidate y_e. In particular, for $k\in\mathcal{K}$, let $\tau_-(x_{e,k})$ denote the total time $x_{e,k}$ stays negative, i.e.,

$$\tau_-(x_{e,k}) := \int_0^{\tau_L}\mathbf{1}_{\{x_{e,k}(s)<0\}}ds.$$

In addition, define

$$M_{p_k}^{(\ell)} := \max\left\{\tau_-(x_{e,k}) - L\left(\frac{\beta_k^{(\ell)}}{\mu_k - \lambda_k} + \frac{\beta_k^{(\ell)}}{\lambda_k}\right), 0\right\}, \quad \ell = 1, ..., N_{p_k}.$$

We construct the negative part of $y_{e,k}$ as follows:

(a) The trajectory of $y_{e,k}$ stays negative for a total of $\tau_-(x_{e,k})$ time units over the time interval $[0, \tau_L]$. In other words, $y_{e,k}$ and $x_{e,k}$ spend the same amount of time below zero.

(b) The negative part of $y_{e,k}$ is identical (symmetric) over each production cycle, and forms an upside-down triangular closed loop that decreases at rate λ_k, increases at rate $\mu_k - \lambda_k$, and has a base of length $\tau_-(x_{e,k})/L$ (equivalently, $M_{p_k}^{(1)}/L$).

It follows from the same analysis as in **(I)** that

$$\frac{1}{\tau_L}\int_0^{\tau_L} \left(p_k\left(y_{e,k}(s)^-\right)\right) ds \leq \frac{1}{\tau_L}\int_0^{\tau_L} \left(p_k\left(x_{e,k}(s)^-\right)\right) ds.$$

(III) For $k \in \mathcal{K}$, the constructed $y_{e,k}$ is symmetric across the L production cycles. Thus, $y_{e,k}$ is indeed a one-cycle PE-candidate with period τ_L/L. The time-average setup cost of $y_{e,k}$ is given by $r_k/(\tau_L/L) = r_k L/\tau_L$, which is equal to that of $x_{e,k}$.

By **(I)**, **(II)**, and **(III)**, given any L-cycle PE-candidate x_e with $L > 1$, we can construct a one-cycle PE-candidate y_e with a (weakly) lower time-average cost than that of x_e. Therefore, exactly one of the following two cases hold: (i) the fluid optimization problem (3.4) does not admit an optimal solution; (ii) the fluid optimization problem (3.4) admits an optimal solution with L = 1. We complete the proof by showing that there exists an optimal solution to the one-cycle fluid optimization problem (2.3), which implies that case (ii) indeed holds.

In the fluid optimization problem (2.3), the decisions variables w and $\mathbf{b}$ fully determine the closed trajectory x_e. Thus, the optimization problem can be considered as optimizing $f(w, \mathbf{b})$ over $w \in \mathbb{R}_+$ and $\mathbf{b} \in \mathbb{R}^K$, for some continuous function $f : \mathbb{R}_+ \times \mathbb{R}^K \to \mathbb{R}_+$. It is then easy to see from the formulation of the problem and Assumption 2 that $f(w, \mathbf{b})$ is coercive in $(w, \mathbf{b})$, namely, for every sequence $\{(w_n, \mathbf{b}_n) : n \geq 1\}$ with $||(w_n, \mathbf{b}_n)|| \to \infty$, it holds that $\lim_{n\to\infty} f(w_n, \mathbf{b}_n) = \infty$. Since the feasibility region $\mathbb{R}_+ \times \mathbb{R}^K$ is non-empty and closed, and since f is continuous and coercive over $\mathbb{R}_+ \times \mathbb{R}^K$, the set of optimal solutions to problem (2.3) is non-empty and compact.

Step 2. Generalization to general cost functions. Let $L \in \mathbb{N}$. For the fluid optimization problem (3.4), it is without loss of optimality to consider L-cycle PE-candidates whose trajectory touches zero at least once over the time $[0, \tau_L]$. To see this, suppose the inventory level of product k, $x_{e,k}$, in an L-cycle PE-candidate x_e is strictly above zero, for some $k \in \mathcal{K}$. We can then reduce the holding cost of $x_{e,k}$ (while the associated setup and backlog costs stay the same) by shifting its trajectory downwards until its lowest point touches zero. Note that this shift does not affect the inventory level processes of the other products, and thus does not violate the feasibility condition (flow balance and synchronization of the trajectories) of the L-cycle PE-candidate. The other case where the trajectory remains strictly below zero can be ruled out by similar arguments. Now, for an L-cycle PE-candidate whose trajectory touches zero at least once, the inventory level of product k is upper bounded by $(\lambda_k + \mu_k)\tau_L$, and lower bounded by $-(\lambda_k + \mu_k)\tau_L$, $k \in \mathcal{K}$. The cost function

ψ_k evaluated at any point of the PE is then bounded above by $\psi_k((\lambda_k+\mu_k)\tau_L)+\psi_k(-(\lambda_k+\mu_k)\tau_L)<\infty$. We can then find two piecewise linear functions p'_k and h'_k that jointly approximate ψ_k to arbitrary precision. This step is standard, and follows the same derivation as in the proof of Proposition 4.1 in Hu et al. (2020a). Q.E.D.

B.3. Proof of Lemma 3

Proof: We prove the lemma by establishing that the mapping of the inventory process from one production cycle to the next under the BSI policy with parameters $(\mathbf{b}_e,\mathbf{w}_e)$ is a contraction map.

Let $k\in\mathcal{K}$. Define the fluid shortfall level process for product k as $y_k(t):=b_k-x_k(t)$, with $y(t):=(y_k(t),k\in\mathcal{K})$, $t\geq 0$. We assume without loss of generality that $x_k(0)\leq b_k$ for all $k\in\mathcal{K}$, so that $y(t)\geq 0$ for all $t\geq 0$.

Under the BSI policy with parameters $(\mathbf{b}_e,\mathbf{w}_e)$, we consider the operator $\mathcal{T}_k:\mathbb{R}_+^K\to\mathbb{R}_+^K$, $k\in\mathcal{K}$, which maps the shortfall level at the polling epoch of product k to that at the polling epoch of product $k+1$. The time interval between these two time epochs consists of the production time of product k, the idle time following the production, and the setup time for product $k+1$. Given $y(t)=\ell\in\mathbb{R}_+^K$ at the polling epoch of product k, the production run lasts for $\ell_k/(\mu_k-\lambda_k)$ units of time, during which the shortfall level of product k decreases at rate $\mu_k-\lambda_k$, and of product j increases at rate λ_j, for $j\neq k$, $j\in\mathcal{K}$. Next, during the idle time $w_{e,k}$ and setup time s_{k+1}, the shortfall level of product j increases at rate λ_j for all $j\in\mathcal{K}$.

Based on the system dynamics, we can write

$$\mathcal{T}_k(\ell)=\mathcal{A}_k\ell+\mathcal{B}_k,$$

where

$$\mathcal{A}_k\ell:=\begin{bmatrix}\ell_1+\lambda_1\frac{\ell_k}{\mu_k-\lambda_k}\\ \vdots\\ \ell_{k-1}+\lambda_{k-1}\frac{\ell_k}{\mu_k-\lambda_k}\\ 0\\ \ell_{k+1}+\lambda_{k+1}\frac{\ell_k}{\mu_k-\lambda_k}\\ \vdots\\ \ell_K+\lambda_K\frac{\ell_k}{\mu_k-\lambda_k}\end{bmatrix},\quad \mathcal{B}_k:=\begin{bmatrix}\lambda_1(w_{e,k}+s_{k+1})\\ \vdots\\ \lambda_K(w_{e,k}+s_{k+1})\end{bmatrix},\quad k\in\mathcal{K}. \tag{B.12}$$

We also write $\mathcal{T}:=\mathcal{T}_K\circ\ldots\circ\mathcal{T}_1$, the composition of the operators over one production cycle, which maps the shortfall level from the beginning epoch of one production cycle to the next. It is easy to see that

$$\mathcal{T}(\ell)=\mathcal{A}\ell+\mathcal{B},\text{ where }\mathcal{A}:=\mathcal{A}_K\cdots\mathcal{A}_1\text{ and }\mathcal{B}:=\sum_{i=1}^{K-1}\left(\prod_{j=i+1}^{K}\mathcal{A}_j\right)\mathcal{B}_i+\mathcal{B}_K.$$

Note that the operator $\mathcal{T}_k$, $k\in\mathcal{K}$, is (i) monotone, i.e., if $x\leq y$, then $\mathcal{T}_k(x)\leq\mathcal{T}_k(y)$, (ii) affine, i.e., $\mathcal{A}_k\ell+\mathcal{B}_k$ for $\mathcal{A}_k,\mathcal{B}_k$ in (B.12), and (iii) dominated, i.e., $\mathcal{B}_k\geq 0$. The properties of monotonicity, affinity, and dominance are inherited by compositions and thus hold for $\mathcal{T}$. Moreover, the operator $\mathcal{T}$ is not only dominated by composition, but also strictly dominated, i.e., $\mathcal{B}>0$, due to the assumption that $s>0$. The following lemma, which is adapted from Theorem 1 in Feoktistova et al. (2012), establishes a contraction property for such operators.

Lemma 7 *Suppose that an iteration of a piecewise affine continuous monotone map $\mathcal{T}$ is strictly dominated and this map has a fixed point $\mathcal{T}(\ell_e)=\ell_e\in\mathbb{R}_+^K$. Then this fixed point is unique and attracts all trajectories of the iterated system $\ell^{(m+1)}=\mathcal{T}(\ell^{(m)})$ with $\ell^{(0)}\in\mathbb{R}_+^K$, namely, $\lim_{m\to\infty}\ell^{(m)}=\ell_e$.*

Let ℓ_e denote the shortfall level at the polling epoch of stage 1 in the one-cycle PE-candidate x_e. It is easy to see that ℓ_e is a fixed point of the operator $\mathcal{T}$. It follows from Lemma 7 that ℓ_e is the unique fixed point to which all trajectories of the shortfall level embedded at the polling epochs of stage 1 converge. This implies that x_e is a global limit cycle for the HDS under the BSI policy with parameters $(\mathbf{b}_e, \mathbf{w}_e)$. Q.E.D.

Remark 1 A result similar to that of Lemma 3 can be found in (Gallego 1990, Proposition 3.2). There, the author establishes that the BSI policy can "steer" the system back to the equilibrium (which is, in our setting, a PE) after the system undergoes disruptions that takes it away from its "normal" state of operations. By treating the system as an HDS, our proof of Lemma 3 is simpler than that in Gallego (1990).

B.4. Proof of Theorem 2

PROOF: The statement follows from Lemmas 2 and 3 by taking the PE to be x_e^*—the optimal PE-candidate to problem (2.3), and the parameters of the BSI policy to be $(\mathbf{b}_*, \mathbf{w}_*)$. Q.E.D.

Appendix C: Auxiliary Lemmas for the Proof of Asymptotic Optimality

In this section, we state several technical lemmas which are needed for the proof of the asymptotic optimality result. First, recall from Definition 2 that for any idle time vector $(W_k, k \in \mathcal{K})$ and stable lot sizing policy $\pi \in \Pi$, $\{\tilde{X}(m) : m \geq 0\}$, the inventory level process embedded at the polling epochs of product 1, is absorbed into a positive recurrent class. Lemma 8 quantifies the expected cycle length for the inventory level process in stationarity. Second, in Section 3 we use a fluid model to approximate the dynamics in the stochastic system. Lemmas 9–11 formally justify the fluid approximation by establishing it as a FWLLN limit of the sequence of stochastic systems as the setup times increase to infinity.

Recall that $s = \sum_{k \in \mathcal{K}} s_k$ is the expected total setup time and $w = \sum_{k \in \mathcal{K}} W_k$ is the total idle time in a cycle.

Lemma 8 *Under any stable control $\pi \in \Pi$, the stationary cycle length T has mean $\mathbb{E}[T] = (s + w)/(1 - \rho)$.*

PROOF: Let $X(0)$ be distributed according to the stationary distribution of the embedded DTMC, namely, $X(0) \stackrel{d}{=} \tilde{X}(\infty)$. The cycle length T consists of the total time the facility spends producing each product, the total idle time w, and the the total setup time in a cycle. Let $L_k(T)$ denote the number of units of product k produced over the time interval $[0, T]$, $k \in \mathcal{K}$. It holds that

$$X_k(0) - P_k(T) + L_k(T) = X_k(T), \quad k \in \mathcal{K}, \tag{C.1}$$

where

$$T = \sum_{\ell=1}^{K} \sum_{j=1}^{L_\ell(T)} G_\ell^{(j)} + w + \sum_{\ell=1}^{K} S_\ell,$$

where $G_\ell^{(j)}$'s are i.i.d. random variables distributed according to the production time of one unit of product $\ell \in \mathcal{K}$. Since $X_k(0)$ and $X_k(T)$ are both distributed according to the stationary distribution of the embedded DTMC, we have that $\mathbb{E}[X_k(0)] = \mathbb{E}[X_k(T)]$, $k \in \mathcal{K}$. Thus, taking expectation on both sides of (C.1) and using Walld's identity, we get that

$$\mathbb{E}[L_k(T)] = \lambda_k \left(\sum_{\ell=1}^{K} \frac{1}{\mu_\ell} \mathbb{E}[L_\ell(T)] + w + s \right), \quad k \in \mathcal{K}. \tag{C.2}$$

To see (C.2), note that

$$
\begin{aligned}
\mathbb{E}\left[\sum_{\ell=1}^{K}\sum_{j=1}^{L_\ell(T)} G_\ell^{(j)}\right] &= \sum_{\ell=1}^{K}\sum_{j=1}^{\infty}\mathbb{E}\left[G_\ell^{(j)} 1_{\{L_\ell(T)>j-1\}}\right] \\
&= \sum_{\ell=1}^{K}\sum_{j=1}^{\infty}\mathbb{E}\left[\mathbb{E}\left[G_\ell^{(j)} 1_{\{L_\ell(T)>j-1\}}\middle|\mathcal{F}_{G_\ell^{(j-1)}}\right]\right] \qquad \text{(C.3)} \\
&= \sum_{\ell=1}^{K}\mathbb{E}\left[G_\ell\right]\sum_{j=1}^{\infty}\mathbb{E}\left[1_{\{L_\ell(T)>j-1\}}\right] = \sum_{\ell=1}^{K}\frac{1}{\mu_\ell}\mathbb{E}\left[L_\ell(T)\right],
\end{aligned}
$$

where $\mathcal{F}_{G_\ell^{(j-1)}}$ denotes the (stopped) σ-field containing the history of process $\hat{X}$ up to the end of the production time $G_\ell^{(j-1)}$, and the event $\{L_\ell(T)=j-1\}\in\mathcal{F}_{G_\ell^{(j-1)}}$. Since the production times are i.i.d. and $G_\ell^{(j)}$ is independent of $\mathcal{F}_{G_\ell^{(j-1)}}$, the third equality in (C.3) follows. Then, applying the same lines of analysis to the other terms and taking the conditional expectation of the Poisson process in (C.1) gives the expectation in (C.2). The unique solution to (C.2) is

$$
\mathbb{E}\left[L_k(T)\right] = \lambda_k\frac{s+w}{1-\rho} \quad \text{and} \quad \mathbb{E}\left[T\right] = \frac{s+w}{1-\rho}, \quad k\in\mathcal{K},
$$

and the statement follows. Q.E.D.

Let $A_k^{(m),n}$ and $D_k^{(m),n}$ respectively denote the polling and departure epochs associated with product k in the mth cycle in system n, $k\in\mathcal{K}$, $m\geq 1$, $n\geq 1$, and let $\bar{A}_k^{(m),n} := A_k^{(m),n}/n$ and $\bar{D}_k^{(m),n} := D_k^{(m),n}/n$ denote the fluid-scaled time epochs. Then $\bar{V}_k^{(m),n} := \bar{D}_k^{(m),n} - \bar{A}_k^{(m),n}$ is the corresponding fluid-scaled busy time the facility spends producing product k in the mth cycle in system n. Recall that $\bar{X}^n$ is the normalized inventory process in system n.

Lemma 9 (tightness) *If $\{\bar{X}^n(0):n\geq 1\}$ is tight in $\mathbb{R}^K$, then $\{\bar{X}^n:n\geq 1\}$ has a convergent subsequence in D^K. Further, the sample path of each subsequential limit of $\{\bar{X}^n:n\geq 1\}$ is a solution to an HDS differential equation of the form* (3.1).

Proof: It follows from inspection that the solution to the HDS differential equation (3.1) can be described inductively via its dynamics over each cycle:

$$
x_k(t) = x_k(u^{(m-1)}) - (t-u^{(m-1)})\lambda_k + \mu_k\int_{u^{(m-1)}}^{t} 1_{\left[a_k^{(m)},\, d_k^{(m)}\right)}(s)ds, \qquad \text{(C.4)}
$$

for $t\in[u^{(m-1)},u^{(m)})$, $k\in\mathcal{K}$ and $m\geq 1$, where $1_A(s)$ denote the indicator function of the set A, which is equal to 1 if $s\in A$, and equal to 0 otherwise.

Under the large-setup-time scaling, as $s>0$, on every interval $[0,t]$, $t>0$, there are finitely many fluid-scaled polling and departure epochs w.p.1. Thus, the polling epochs are sequentially compact in $\mathbb{R}_+^j$, where j is the number of polling epochs for all n large enough. If no such j exists, then there exists a subsequence indexed by n_k with such a j for all k large enough, and we focus on each such subsequence. The sequential compactness of the polling epochs implies that each of the subsequences has a further converging sub-subsequence.

Consider a (sub-)subsequence $\{\bar{X}^{n_j}:j\geq 1\}$ over which the initial condition $\{\bar{X}^{n_j}(0):j\geq 1\}$ and polling epochs $\{\bar{A}_k^{(m),n_j}, j\geq 1\}$ converge over any compact intervals for each $k\in\mathcal{K}$. Consider a further subsequence of $\{\bar{X}^{n_j}:j\geq 1\}$, e.g., $\{\bar{X}^{n_{j_\ell}}:\ell\geq 1\}$, over which the departure epochs $\{\bar{D}_k^{(m),n_j}, j\geq 1\}$ converge (over any compact intervals), $k \in \mathcal{K}$. The rest of the proof is standard. In particular, the dynamics of $\bar{X}^{n_{j_\ell}}$ are easily seen to be a continuous mapping of its primitive over time intervals between switching (i.e., polling and departure) epochs. The convergence of the switching epochs implies that the busy times $\{\bar{V}_k^{(m),n_j}, j \geq 1\}$ converge for each $k \in \mathcal{K}$. Moreover, since the setup times converge weakly under the large-setup-time scaling, it can be inferred that the sequence of fluid-scaled idle times, i.e., $\{\bar{W}_k^{(m),n_j}, j \geq 1\}$, converges in $\mathbb{R}_+$ for each $k \in \mathcal{K}$.

Let $\bar{X}$ be a subsequential limit point of $\{\bar{X}^n : n \geq 1\}$. It is easy to see that the sample paths of $\bar{X}$ are of the form

$$\bar{X}_k(t) = \bar{X}_k(0) - \lambda_k t + \mu_k \bar{\Phi}_k(t), \quad t \geq 0, \quad k \in \mathcal{K}, \tag{C.5}$$

where $\bar{\Phi}_k := \{\bar{\Phi}_k(t) : t \geq 0\}$ is a cumulative process of the form

$$\bar{\Phi}_k(t) = \int_0^t \phi_k(s) ds, \tag{C.6}$$

for a piecewise-constant function $\phi_k : \mathbb{R}_+ \to \{0, 1\}$. In particular, $\phi_k(t)$ is equal to 1 if the facility is producing product k at time t, and equal to 0 otherwise. (Note that ϕ_k is piecewise-constant because there are finitely many switching epochs over any compact time interval.) The statement follows from (C.4), (C.5), and (C.6). Q.E.D.

It is significant that the limits of converging subsequence of $\{\bar{X}^n : n \geq 1\}$ need not be deterministic. Lemma 9 can be strengthened to a FWLLN under an extra regularity condition on the double arrays of polling and departure epochs.

Lemma 10 (FWLLN) *Assume that $\bar{X}^n(0) \Rightarrow x(0)$ as $n \to \infty$, where $x(0)$ is a deterministic element of $\mathbb{R}^K$. If $\bar{A}_k^{(m),n} \Rightarrow a_k^{(m)}$ and $\bar{D}_k^{(m),n} \Rightarrow d_k^{(m)}$ for all $m \geq 1$ and $k \in \mathcal{K}$, where $a_k^{(m)}$ and $d_k^{(m)}$'s are deterministic elements of $\mathbb{R}_+$, then $\bar{X}^n \Rightarrow x$ in D^K as $n \to \infty$, where x is the unique solution to the HDS* (C.4) *with initial condition $x(0)$, polling epochs $\{a_k^{(m)} : m \geq 1\}$, and departure epochs $\{d_k^{(m)} : m \geq 1\}$, $k \in \mathcal{K}$.*

The proof of Lemma 10 follows closely that of Lemma 9 and is thus omitted.

It is important to note that if $\bar{X}^n(0) \Rightarrow x(0)$, convergence of $\bar{X}^n$ to the fluid model should hold under any "reasonable" policy, for example, under the sequence of BSI policies with parameters

$$\mathbf{B}^n := \lceil n\mathbf{b} \rceil \quad \text{and} \quad \mathbf{W}^n := n\mathbf{w}, \quad n \geq 1. \tag{C.7}$$

Lemma 11 (FWLLN under BSI) *For the sequence of stochastic systems under the sequence of BSI policies with base-stock levels $\{\mathbf{B}^n : n \geq 1\}$ and idle times $\{\mathbf{W}^n : n \geq 1\}$ in* (C.7)*, if $\bar{X}^n(0) \Rightarrow x(0)$ in $\mathbb{R}^K$, then $\bar{X}^n \Rightarrow x$ in D^K, where x the a fluid inventory level process under the BSI policy with base-stock levels* **b***, idle times* **w***, and initial condition $x(0)$.*

Proof: Let $n \geq 1$. Under the BSI policy with base-stock levels $\mathbf{B}^n$ and idle times $\mathbf{W}^n$, the shortfall level of product k in the nth system is

$$Y_k^n(t) = B_k^n - X_k^n(t), \quad k \in \mathcal{K}, \quad t \geq 0. \tag{C.8}$$

For $Y^n = (Y_k^n, k \in \mathcal{K})$, we have derived in Section 4.2 that Y^n can be equivalently considered as the queue length process under the exhaustive policy for polling systems. Moreover, the idle times and the setup times combined together can be considered as the switchover times for polling systems. The statement then follows from Corollary 5.1 in Hu et al. (2020a) by noting that the exhaustive policy for polling systems is equivalent to the binomial-exhaustive policy with parameters $(1, \mathbf{1})$. Q.E.D.

Appendix D: Proof of results in Section 4

We first introduce more notation. For the nth system, $n \geq 1$, let $U^{(m),n}$ denote the beginning epoch of the $(m+1)$st production cycle (equivalently, the polling epoch of product 1). We define the fluid-scaled discrete-time inventory level process embedded at the polling epochs of product 1 as

$$\tilde{X}^n(m) := \bar{X}^n(\bar{U}^{(m),n}), \quad m \geq 0, \quad n \geq 1.$$

Similarly, the embedded inventory level process for the fluid model under an admissible control is

$$\tilde{x}(m) := x(u^{(m)}), \quad m \geq 0.$$

D.1. Proof of Lemma 4

Proof: We conduct the proof in the following four steps: In **Step 1**, we examine the embedded DTMC $\{\tilde{X}^n(m) : m \geq 0\}$ under a sequence of admissible controls $\{\pi^n : n \geq 1\}$. We assume without loss of generality that the sequence of stationary DTMC $\{\tilde{X}^n(\infty) : n \geq 1\}$ is tight, and initialize system n according to $\tilde{X}^n(\infty)$, i.e., $\bar{X}^n(0) \stackrel{d}{=} \tilde{X}^n(\infty)$. We then restrict to a convergent subsequence of $\{\bar{X}^n : n \geq 1\}$, and let $\bar{X}$ denote its limit point. For $r \geq 0$, we construct a ball $\mathcal{B}_r$ in $\mathbb{R}^K$ with radius r, which both $\{\tilde{X}^n(m) : m \geq 0\}$ and the limit $\{\tilde{X}(m) : m \geq 1\})$ visit infinitely often, for sufficiently large $n \geq 1$. In **Step 2**, we translate the number of discrete transitions after which the DTMC $\{\tilde{X}^n(m) : m \geq 0\}$ (alternatively, $\{\tilde{X}(m) : m \geq 0\}$) returns to $\mathcal{B}_r$ for the continuous-time process $\bar{X}^n$ ($\bar{X}$), $n \geq 1$. Let $\bar{R}_r^n$ ($\bar{R}_r$) denote the first return time of $\bar{X}^n$ ($\bar{X}$) to the ball $\mathcal{B}_r$. We show that $\{\bar{R}_r^n : n \geq 1\}$ is uniformly integrable (UI), namely, $\bar{R}_r^n \Rightarrow \bar{R}_r$ and $\mathbb{E}\left[\bar{R}_r^n\right] \to \mathbb{E}\left[\bar{R}_r\right]$ as $n \to \infty$. In **Step 3**, motivated by the observation that the limit $\bar{X}$ returns to ball $\mathcal{B}_r$ infinitely often and r is arbitrary, we use a sample-path approach to formalize the idea that $\bar{X}$ "gets close to being a PE-candidate" as r decreases. In **Step 4**, combining the results in the first three steps, we show that the limiting cost under $\{\pi^n : n \geq 1\}$ is arbitrarily close to the cost of a (sample-path dependent) PE-candidate w.p.1, which is necessarily no lower than the optimal fluid cost c_*. Below, we elaborate on each step one by one.

Step 1. Stationary embedded DTMC. For asymptotic optimality, it is without loss of generality to restrict to admissible controls under which the sequence of embedded stationary DTMC's $\{\tilde{X}^n(\infty) : n \geq 1\}$ is tight with bounded limits in $\mathbb{R}^K$. Then with $\bar{X}^n(0) \stackrel{d}{=} \tilde{X}^n(\infty)$ for $n \geq 1$, the sequence $\{\bar{X}^n(0) : n \geq 1\}$ is tight. It follows from Lemma 9 that $\{\bar{X}^n : n \geq 1\}$ has a bounded convergent subsequence in D^K. Let $\bar{X}$ denote a subsequential limit of $\bar{X}^n$. Then the sample paths of $\bar{X}$ are HDS of the form (C.4). To simply notation, we shall henceforth use $\{\bar{X}^n : n \geq 1\}$ to denote the converging subsequence with limit point $\bar{X}$.

Let α^n denote the (stationary) distribution of $\{\tilde{X}^n(m) : m \geq 0\}$. Since each process in the pre-limit is stationary, the limit $\{\tilde{X}(m) : m \geq 0\}$ must be stationary with some stationary distribution α. For $r \geq 0$, let $\mathscr{B}(r)$ denote a ball in $\mathbb{R}^K$ with positive measure, namely, $\alpha(\mathscr{B}(r)) \in (0, 1]$. Let $B_r(0)$ denote the center of ball $\mathscr{B}(r)$. Note that we do not rule out the case where $r=0$, which is tantamount to $\mathscr{B}(r)$ being a point in $\mathbb{R}^K$ and the limiting distribution α has a point mass.

By the weak convergence of $\{\tilde{X}^n : n \geq 1\}$ to $\tilde{X}$, we have

$$\lim_{n\to\infty} \mathbb{P}\left(\tilde{X}^n(0) \in \mathscr{B}(r)\right) = \mathbb{P}\left(\tilde{X}(0) \in \mathscr{B}(r)\right), \tag{D.1}$$

so that

$$\lim_{n\to\infty} \alpha^n(\mathscr{B}(r)) = \alpha(\mathscr{B}(r)) > 0.$$

Note that (D.1) implies that $\alpha^n(\mathscr{B}(r)) > 0$ for sufficiently large n. For the purpose of proving the asymptotic lower bound, we shall assume without loss of generality that $\alpha^n(\mathscr{B}(r)) > 0$ for all $n \geq 1$. Then $\{\tilde{X}^n(m) : m \geq 0\}$ must return to $\mathscr{B}(r)$ infinitely often for all $n \geq 1$; the same holds for the limiting ball $\mathscr{B}(r)$ and the limiting process $\{\tilde{X}(m) : m \geq 0\}$.

Let

$$N_r^n := \inf\{m \geq 1 : \tilde{X}^n(m) \in \mathscr{B}(r)\} \quad \text{and} \quad N_r := \inf\{m \geq 1 : \tilde{X}(m) \in \mathscr{B}(r)\}.$$

Denote

$$\alpha_r^n(\cdot) := \mathbb{P}\left(\tilde{X}^n(0) \in \cdot \,|\, \tilde{X}^n(0) \in \mathscr{B}(r)\right) \quad \text{and} \quad \alpha_r(\cdot) := \mathbb{P}\left(\tilde{X}(0) \in \cdot \,|\, \tilde{X}(0) \in \mathscr{B}(r)\right).$$

In addition, we write $\mathbb{E}_{\alpha_r^n}$ as the expectation operator for which $\mathbb{P}(\tilde{X}^n(0) \in \cdot) = \alpha_r^n(\cdot)$, and $\mathbb{E}_{\alpha_r}$ as the expectation operator for which $\mathbb{P}(\tilde{X}(0) \in \cdot) = \alpha_r(\cdot)$. It follows from basic DTMC theory that

$$\lim_{n\to\infty} \mathbb{E}_{\alpha_r^n}[N_r^n] = \lim_{n\to\infty} \frac{1}{\alpha^n(\mathscr{B}(r))} = \frac{1}{\alpha(\mathscr{B}(r))} = \mathbb{E}_{\alpha_r}[N_r]. \tag{D.2}$$

Step 2. Uniformly integrable return time. To translate the number of discrete transitions of the embedded DTMC for the continuous-time process, we define

$$\bar{R}_r^n := \inf\{\bar{U}^{(m),n} > 0 : \bar{X}^n(\bar{U}^{(m),n}) \in \mathscr{B}(r)\}$$
$$\bar{R}_r := \inf\{\bar{U}^{(m)} > 0 : \bar{X}(\bar{U}^{(m)}) \in \mathscr{B}(r)\}.$$

By the weak convergence of $\bar{X}^n$ to $\bar{X}$, we have $\bar{R}_r^n \Rightarrow \bar{R}_r$; see, e.g., Theorem 13.6.4 in Whitt (2002). In addition, we can also establish the following UI result for $\{\bar{R}_r^n : n \geq 1\}$, i.e.,

$$\lim_{n\to\infty} \mathbb{E}_{\alpha_r^n}\left[\bar{R}_r^n\right] = \mathbb{E}_{\alpha_r}\left[\bar{R}_r\right]. \tag{D.3}$$

To see (D.3), let $L_k^n(t)$ denote the number of units of product k produced over the time interval $[0,t]$, and $\bar{L}_k^n(t) := L_k^n(nt)/n$. We also define $P_k^n(t)$ as a Poisson process with rate λ_k and $\bar{P}_k^n(t) = P_k^n(nt)/n$. It can be derived using similar lines of analysis as in the proof of Lemma 8 that

$$\bar{X}_k^n(0) - \bar{P}_k^n\left(\bar{R}_r^n\right) + \bar{L}_k^n(\bar{R}_r^n) = \bar{X}_k^n(\bar{R}_r^n), \quad k \in \mathcal{K}, \quad n \geq 1, \tag{D.4}$$

and

$$\bar{R}_r^n = \frac{1}{n}\sum_{\ell=1}^{K}\sum_{j=1}^{L_\ell^n(n\bar{R}_r^n)} G_\ell^{(j)} + N_r^n \bar{w}^n + \sum_{m=1}^{N_r^n}\sum_{\ell=1}^{K} \bar{S}_\ell^{(m),n},$$

where $G_\ell^{(j)}$'s are i.i.d. random variables distributed according to the production time of one unit of product $\ell \in \mathcal{K}$, and $\bar{w}^n := \sum_{k\in\mathcal{K}} \bar{W}_k^n$. Moreover, plugging $\mathbb{E}_{\alpha_r^n}\left[\bar{X}_k^n(0)\right] = \mathbb{E}_{\alpha_r^n}\left[\bar{X}_k^n(\bar{R}_r^n)\right]$ into (D.4) by taking expectation on both sides of the equation, and using Ward's identity, we get that for each $n \geq 1$ and $k \in \mathcal{K}$,

$$\mathbb{E}_{\alpha_r^n}\left[\bar{L}_k^n(\bar{R}_r^n)\right] = \lambda_k \left(\sum_{\ell=1}^K \frac{1}{\mu_\ell} \mathbb{E}_{\alpha_r^n}\left[\bar{L}_\ell^n(\bar{R}_r^n)\right] + \mathbb{E}_{\alpha_r^n}\left[N_r^n\right]\bar{w}^n + \mathbb{E}_{\alpha_r^n}\left[N_r^n\right]s \right). \tag{D.5}$$

The unique solution to (D.5) is

$$\mathbb{E}_{\alpha_r^n}\left[\bar{L}_k^n(\bar{R}_r^n)\right] = \lambda_k \frac{s+\bar{w}^n}{1-\rho}\mathbb{E}_{\alpha_r^n}\left[N_r^n\right] \quad \text{and} \quad \mathbb{E}_{\alpha_r^n}\left[\bar{R}_r^n\right] = \frac{s+\bar{w}^n}{1-\rho}\mathbb{E}_{\alpha_r^n}\left[N_r^n\right]. \tag{D.6}$$

Similar "flow equation" as in (D.4) holds for the subsequential limit process $\bar{X}$. In particular, by the dynamics of $\bar{X}$ in (C.5), we have

$$\bar{X}_k(\bar{R}_r) = \bar{X}_k(0) - \lambda_k \bar{R}_r + \mu_k \bar{\Phi}_k(\bar{R}_r), \quad k \in \mathcal{K},$$

where $\bar{R}_r = \sum_{\ell=1}^K \bar{\Phi}_\ell(\bar{R}_r) + N_r\bar{w} + N_r s$, and $\bar{w}$ is the limit point of $\{\bar{w}^n : n \geq 1\}$ in $\mathbb{R}_+$. Again, since both $\bar{X}_k(0)$ and $\bar{X}_k(\bar{R}_r)$ are distributed according to the stationary distribution α restricted to the ball $\mathscr{B}(r)$, it holds that $\mathbb{E}_{\alpha_r}\left[\bar{X}_k(0)\right] = \mathbb{E}_{\alpha_r}\left[\bar{X}_k(\bar{R}_r)\right]$, and that

$$\lambda_k \mathbb{E}_{\alpha_r}\left[\sum_{\ell=1}^K \bar{\Phi}_\ell(\bar{R}_r) + N_r\bar{w} + N_r s\right] = \mu_k \mathbb{E}_{\alpha_r}\left[\bar{\Phi}_k(\bar{R}_r)\right], \quad k \in \mathcal{K},$$

whose unique solution is

$$\mathbb{E}_{\alpha_r}\left[\bar{\Phi}_k(\bar{R}_r)\right] = \rho_k \frac{s+\bar{w}}{1-\rho}\mathbb{E}_{\alpha_r}\left[N_r\right] \quad \text{and} \quad \mathbb{E}_{\alpha_r}\left[\bar{R}_r\right] = \frac{s+\bar{w}}{1-\rho}\mathbb{E}_{\alpha_r}\left[N_r\right], \quad k \in \mathcal{K}. \tag{D.7}$$

Lastly, since $\bar{w}^n \to \bar{w}$ as $n \to \infty$, and $\mathbb{E}_{\alpha_r^n}\left[N_r^n\right] \to \mathbb{E}_{\alpha_r}\left[N_r\right]$ as $n \to \infty$ by (D.2), it follows from the second equalities in (D.6) and (D.7) that $\mathbb{E}_{\alpha_r^n}\left[\bar{R}_r^n\right] \to \mathbb{E}_{\alpha_r}\left[\bar{R}_r\right]$ as $n \to \infty$, and (D.3) follows.

Step 3. "Nearly-periodic" limit. We next use a sample path argument to quantify the difference between $\bar{X}$ and a N_r-cycle PE-candidate over one return time to $\mathscr{B}(r)$ (from time 0 to $\bar{R}_r$). For each sample point $\omega \in \Omega$, we use $N_r(\omega)$ to denote the number of cycles contained in the return time $\bar{R}_r(\omega)$ for $\bar{X}(\omega, \cdot)$, where the subscript ω marks the sample-path dependence. (Note that $N_r(\omega)$ and $\bar{R}_r(\omega)$ are deterministic elements of $\mathbb{N}$ and $\mathbb{R}_+$, respectively.) It follows from (3.6) that $\tau_{N_r(\omega)} = (s+\bar{w})N_r(\omega)/(1-\rho)$ is the corresponding cycle length for the $N_r(\omega)$-cycle PE-candidate.

Lemma 12 *There exist constants $d_1, d_2 > 0$ such that the following hold.*

(i) It holds w.p.1 that

$$\left|\bar{R}_r - \tau_{N_r}\right| \leq d_1 r. \tag{D.8}$$

(ii) For any sample point $\omega \in \Omega$ (except for those in a set of measure zero), there exists an $N_r(\omega)$-cycle PE-candidate x^ω such that

$$||\bar{X}(\omega,\cdot) - x^\omega||_{\bar{R}_r(\omega)\vee\tau_{N_r(\omega)}} \leq d_2 r, \tag{D.9}$$

where for $x, y \in \mathbb{R}$, $x \vee y := \min\{x, y\}$.

The proof of Lemma 12 is relegated to Appendix D.1.1.

For the following, we assume without loss of generality that $\bar{R}_r(\omega) \geq \tau_{N_r(\omega)}$. (The other case follows similar lines of analysis.) Note that (D.9) implies that

$$\max_{k\in\mathcal{K}} |\bar{X}_k(\omega,t) - x_k^\omega(t)| \leq d_2 r \quad \text{for all } t \in [0, \bar{R}_r(\omega)]. \tag{D.10}$$

Recall that B^0 is the center of $\mathscr{B}(r)$, and $x^\omega(0) = x^\omega(\tau_{N_r(\omega)})$. Since $||x^\omega(0) - B^0|| \leq r$, it holds that

$$\max_{k\in\mathcal{K}} x_k^\omega(\tau_{N_r(\omega)}) \leq \max_{k\in\mathcal{K}} B_k^0 + r.$$

Then (D.10) implies that

$$\max_{k\in\mathcal{K}} \bar{X}_k(\omega, \tau_{N_r(\omega)}) \leq \max_{k\in\mathcal{K}} B_k^0 + r + d_2 r,$$

and in turn, for all $t \in [\tau_{N_r(\omega)}, \bar{R}_r(\omega)]$,

$$\begin{aligned}
\max_{k\in\mathcal{K}} \bar{X}_k(\omega,t) &\leq \max_{k\in\mathcal{K}} \bar{X}_k(\omega, \tau_{N_r(\omega)}) + \max_{k\in\mathcal{K}}\{\lambda_k + \mu_k - \lambda_k\}(t - \tau_{N_r(\omega)}) \\
&\leq \max_{k\in\mathcal{K}} B_k^0 + r + d_2 r + \max_{k\in\mathcal{K}}\{\mu_k\}(\bar{R}_r(\omega) - \tau_{N_r(\omega)}) \\
&\leq \max_{k\in\mathcal{K}} B_k^0 + r(1 + d_2 + \max_{k\in\mathcal{K}}\{\mu_k\} d_1) \text{ by (D.8).}
\end{aligned} \tag{D.11}$$

Let $\epsilon > 0$. We select $r > 0$ that satisfies the following three conditions:

First, since $\bar{X}(\omega,\cdot)$ is bounded, and since ψ, x^ω and $\bar{X}(\omega,\cdot)$ are continuous, the composite functions $\psi \circ x^\omega$ and $\psi \circ \bar{X}(\omega,\cdot)$ are both uniformly continuous over any compact time interval. From (D.9), we have there exists $r > 0$ small enough such that

$$|\psi(\bar{X}(\omega,t)) - \psi(x^\omega(t))| < \epsilon/3, \quad \text{for all } t \in [0, \tau_{N_r(\omega)}]. \tag{D.12}$$

Second, (D.11) and the continuity of ψ_k, $k \in \mathcal{K}$, imply that r can be made sufficiently small such that for all $t \in [\tau_{N_r(\omega)}, \bar{R}_r(\omega)]$,

$$\begin{aligned}
\psi(\bar{X}(\omega,t)) &\leq K \max_{k\in\mathcal{K}} \psi_k(\bar{X}_k(\omega,t)) \\
&\leq K \max_{k\in\mathcal{K}} \psi_k(\max_{k\in\mathcal{K}} B_k^0 + r(1 + d_2 + \max_{k\in\mathcal{K}}\{\mu_k\} d_1)) \leq K \max_{k\in\mathcal{K}} \psi_k(\max_{k\in\mathcal{K}} B_k^0) + \epsilon.
\end{aligned}$$

Then by (D.8) and the fact that $\tau_{N_r(\omega)} \geq \tau_1$,

$$\frac{1}{\tau_{N_r(\omega)}} \int_{\tau_{N_r(\omega)}}^{\bar{R}_r(\omega)} \psi(\bar{X}(\omega,s))ds \leq \frac{1}{\tau_1}(K \max_{k\in\mathcal{K}} \psi_k(\max_{k\in\mathcal{K}} B_k^0) + \epsilon) d_1 r < \epsilon/3. \tag{D.13}$$

Third, we consider r such that

$$\frac{2r(1-\rho)}{s + \bar{w}} < \epsilon/3. \tag{D.14}$$

With our choice of r, the time-average cost of $\bar{X}(\omega,\cdot)$ over one return time, in comparison to that of the constructed PE-candidate x^ω, satisfies

$$
\begin{aligned}
&\left|\frac{1}{\bar{R}_r(\omega)}\left(rN_r(\omega)+\int_0^{\bar{R}_r(\omega)}\psi(\bar{X}(\omega,s))ds\right)-\frac{1}{\tau_{N_r(\omega)}}\left(rN_r(\omega)+\int_0^{\tau_{N_r(\omega)}}\psi(x^\omega(s))ds\right)\right|\\
=&\left|\left(\frac{1}{\bar{R}_r(\omega)}-\frac{1}{\tau_{N_r(\omega)}}\right)rN_r(\omega)+\frac{1}{\bar{R}_r(\omega)}\left(\int_0^{\tau_{N_r(\omega)}}\psi(\bar{X}(\omega,s))ds+\int_{\tau_{N_r(\omega)}}^{\bar{R}_r(\omega)}\psi(\bar{X}(\omega,s))ds\right)\right.\\
&\left.\quad-\frac{1}{\tau_{N_r(\omega)}}\int_0^{\tau_{N_r(\omega)}}\psi(x^\omega(s))ds\right|\\
\leq&2\frac{1}{\tau_{N_r(\omega)}}rN_r(\omega)+\frac{1}{\tau_{N_r(\omega)}}\int_0^{\tau_{N_r(\omega)}}|\psi(\bar{X}(\omega,s))-\psi(x^\omega(s))|ds+\frac{1}{\tau_{N_r(\omega)}}\int_{\tau_{N_r(\omega)}}^{\bar{R}_r(\omega)}\psi(\bar{X}(\omega,s))ds \qquad \text{(D.15)}\\
=&2\frac{1}{(s+\bar{w})N_r(\omega)/(1-\rho)}rN_r(\omega)+\frac{1}{\tau_{N_r(\omega)}}\int_0^{\tau_{N_r(\omega)}}|\psi(\bar{X}(\omega,s))-\psi(x^\omega(s))|ds\\
&+\frac{1}{\tau_{N_r(\omega)}}\int_{\tau_{N_r(\omega)}}^{\bar{R}_r(\omega)}\psi(\bar{X}(\omega,s))ds\\
\leq&\frac{2r(1-\rho)}{s+\bar{w}}+\frac{1}{\tau_{N_r(\omega)}}\int_0^{\tau_{N_r(\omega)}}|\psi(\bar{X}(\omega,s))-\psi(x^\omega(s))|ds+\frac{1}{\tau_{N_r(\omega)}}\int_{\tau_{N_r(\omega)}}^{\bar{R}_r(\omega)}\psi(\bar{X}(\omega,s))ds\\
<&\epsilon \text{ by (D.12)–(D.14).}
\end{aligned}
$$

Step 4. Lower bound for long-run average cost. Let c^ω denote the time-average total cost of x^ω over the return time, i.e.,

$$
c^w:=\frac{1}{\tau_{N_r(\omega)}}\left(rN_r(\omega)+\int_0^{\tau_{N_r(\omega)}}\psi(x^\omega(s))ds\right).
$$

By (D.15) and the fact that $c^w\geq c_*$, we have

$$
\frac{1}{\bar{R}_r(\omega)}\left(rN_r(\omega)+\int_0^{\bar{R}_r(\omega)}\psi(\bar{X}(\omega,s))ds\right)>c^\omega-\epsilon\geq c_*-\epsilon. \tag{D.16}
$$

Since (D.16) holds for all ω (except for those in a set of measure zero), it holds that

$$
\frac{1}{\bar{R}_r}\left(rN_r+\int_0^{\bar{R}_r}\psi(\bar{X}(s))ds\right)\geq c_*-\epsilon \quad \text{w.p.1.} \tag{D.17}
$$

Lastly, by (D.1) and the regenerative structure of the inventory level process, we have

$$
\begin{aligned}
\liminf_{n\to\infty}\bar{C}^n&=\liminf_{n\to\infty}\lim_{t\to\infty}\frac{1}{t}\left(r\bar{\Gamma}^n(t)+\int_0^t\psi\left(\bar{X}^n(s)\right)ds\right) \qquad \text{(D.18)}\\
&=\liminf_{n\to\infty}\frac{\mathbb{E}_{\alpha_r^n}\left[rN_r^n+\int_0^{\bar{R}_r^n}\psi\left(\bar{X}^n(s)\right)ds\right]}{\mathbb{E}_{\alpha_r^n}\left[\bar{R}_r^n\right]} \quad \text{w.p.1 by the renewal-reward theorem}\\
&\geq\frac{\liminf_{n\to\infty}\mathbb{E}_{\alpha_r^n}\left[rN_r^n+\int_0^{\bar{R}_r^n}\psi\left(\bar{X}^n(s)\right)ds\right]}{\limsup_{n\to\infty}\mathbb{E}_{\alpha_r^n}\left[\bar{R}_r^n\right]}\\
&=\frac{\liminf_{n\to\infty}\mathbb{E}_{\alpha_r^n}\left[rN_r^n+\int_0^{\bar{R}_r^n}\psi\left(\bar{X}^n(s)\right)ds\right]}{\mathbb{E}_{\alpha_r}\left[\bar{R}_r\right]} \quad \text{by (D.3)}\\
&\geq\frac{\mathbb{E}_{\alpha_r}\left[\liminf_{n\to\infty}\left(rN_r^n+\int_0^{\bar{R}_r^n}\psi\left(\bar{X}^n(s)\right)ds\right)\right]}{\mathbb{E}_{\alpha_r}\left[\bar{R}_r\right]} \quad \text{by Fatou's lemma}\\
&=\frac{\mathbb{E}_{\alpha_r}\left[\left(\frac{1}{\bar{R}_r}\left(rN_r+\int_0^{\bar{R}_r}\psi\left(\bar{X}(s)\right)ds\right)\right)\bar{R}_r\right]}{\mathbb{E}_{\alpha_r}\left[\bar{R}_r\right]}
\end{aligned}
$$

$$
\begin{aligned}
&\geq \frac{\mathbb{E}_{\alpha_r}\left[(c_* - \epsilon)\bar{R}_r\right]}{\mathbb{E}_{\alpha_r}\left[\bar{R}_r\right]} \quad \text{by (D.17)} \\
&= c_* - \epsilon.
\end{aligned}
$$

Note that in second equality in (D.18), for non-negative costs, the renewal-reward theorem holds regardless of whether $\mathbb{E}_{\alpha_r^n}\left[rN_r^n + \int_0^{\bar{R}_r^n} \psi\left(\bar{X}^n(s)\right) ds\right] < \infty$; see Theorem 2.2.1 and the corresponding remarks in (Tijms 2003, p.21). The statement of Lemma 4 follows because (D.18) holds for arbitrary $\epsilon > 0$. Q.E.D.

D.1.1. Proof of Lemma 12

Proof: **(i)** By (C.5) and the fact that $\|\bar{X}(\bar{R}_r) - \bar{X}(0)\| \leq 2r$, it holds that for each $k \in \mathcal{K}$,

$$-2r \;\leq\; -\mu_k \bar{\Phi}_k(\bar{R}_r) + \lambda_k \bar{R}_r \;\leq\; 2r, \tag{D.19}$$

so that

$$\sum_{k\in\mathcal{K}} (-2r/\mu_k + \rho_k \bar{R}_r) \;\leq\; \sum_{k\in\mathcal{K}} \bar{V}_k(\bar{R}_r) \;\leq\; \sum_{k\in\mathcal{K}} (2r/\mu_k + \rho_k \bar{R}_r). \tag{D.20}$$

Since $\bar{R}_r$ is the total length of the N_r cycles, it consists of sN_r units of time the facility spends setting up, $\bar{w}N_r$ units of time the facility idles, and the time the facility spends producing each product. Thus,

$$\bar{R}_r = sN_r + \bar{w}N_r + \sum_{k\in\mathcal{K}} \bar{\Phi}_k(\bar{R}_r).$$

It then follows from (D.20) that

$$sN_r + \bar{w}N_r + \sum_{k\in\mathcal{K}} \left(-2r/\mu_k + \rho_k \bar{R}_r\right) \;\leq\; \bar{R}_r \;\leq\; sN_r + \bar{w}N_r + \sum_{k\in\mathcal{K}} \left(2r/\mu_k + \rho_k \bar{R}_r\right),$$

so that

$$-\frac{2r}{1-\rho} \sum_{k\in\mathcal{K}} \frac{1}{\mu_k} \;\leq\; \bar{R}_r - \tau_{N_r} \;\leq\; \frac{2r}{1-\rho} \sum_{k\in\mathcal{K}} \frac{1}{\mu_k}. \tag{D.21}$$

Thus, letting $d_1 := \frac{2}{1-\rho} \sum_{k\in\mathcal{K}} \frac{1}{\mu_k}$ gives (D.8) in Lemma 12.

(ii) Let $\omega \in \Omega$. In what follows, all the random variables and processes are assumed to be realizations corresponding to that fixed ω in the sample space. To simplify notation, we drop ω from the notation, except for the constructed PE-candidate x^ω.

Consider the limiting process $\bar{X}$ over the time interval $[0, \bar{R}_r]$. For product $k \in \mathcal{K}$, let $\bar{V}_k^{(m)}$, $m \in 1, \ldots, N_r$, denote the busy time the facility spends producing product k in the mth cycle. By construction,

$$\bar{\Phi}_k(\bar{R}_r) = \sum_{m=1}^{N_r} \sum_{\ell=1}^{K} \bar{V}_\ell^{(m)}, \quad k \in \mathcal{K}.$$

It follows from (D.19) and (D.21) that for $k \in \mathcal{K}$,

$$-2r\left(\frac{1}{\mu_k} + \frac{\rho_k}{1-\rho} \sum_{k\in\mathcal{K}} \frac{1}{\mu_k}\right) \;\leq\; \bar{\Phi}_k(\bar{R}_r) - \rho_k \tau_{N_r} \;\leq\; 2r\left(\frac{1}{\mu_k} + \frac{\rho_k}{1-\rho} \sum_{k\in\mathcal{K}} \frac{1}{\mu_k}\right). \tag{D.22}$$

Let $\delta_k := \bar{\Phi}_k(\bar{R}_r) - \rho_k \tau_{N_r}$. (D.22) implies that

$$|\delta_k| \leq 2r\left(\frac{1}{\mu_k} + \frac{\rho_k}{1-\rho} \sum_{k\in\mathcal{K}} \frac{1}{\mu_k}\right), \quad k \in \mathcal{K}.$$

We next construct the corresponding N_r-cycle PE-candidate x^ω by taking $x^\omega(0) := \bar{X}(0)$ and then specifying the busy times. Let $(\bar{V}_k^{\prime(m)}, m = 1, ..., N_r)$ denote the busy time the facility spends producing product k in cycles 1 to N_r.

(1) For product k with $\delta_k < 0$, we take $\bar{V}_k^{\prime(m)} := \bar{V}_k^{(m)}$ for $m = 2, ..., N_r$, but $\bar{V}_k^{\prime(1)} := \bar{V}_k^{(1)} + |\delta_k|$. This means that the N_r-cycle PE-candidate x^ω has identical busy times for product k as those in $\bar{X}$, except that the first busy time is prolonged by $|\delta_k|$.

(2) For product k with $\delta_k > 0$, we take $\bar{V}_k^{\prime(\hat{m})} := \bar{V}_k^{(\hat{m})} - \delta_k$ for some $\hat{m} \in \{1, ..., N_r\}$ with $\bar{V}_k^{(\hat{m})} \geq \delta_k$. (We know such $\hat{m}$ exists for sufficiently small r due to (D.19) and (D.22).) In addition, we take $\bar{V}_k^{\prime(m)} := \bar{V}_k^{(m)}$ for all $m = 1, ..., N_r$, $m \neq \hat{m}$. Thus, the N_r-cycle PE-candidate x^ω has identical busy times for product k as those in $\bar{X}$, except that one busy time is shortened by δ_k.

(3) For product k with $\delta_k = 0$, we take $\bar{V}_k^{\prime(m)} := \bar{V}_k^{(m)}$ for all $m = 1, ..., N_r$. Namely, the N_r-cycle PE-candidate x^ω has identical busy times for product k as those in $\bar{X}$.

Observe that the constructed busy times satisfy flow balance for the inventory level of all products, i.e.,

$$\sum_{m=1}^{N_r} \bar{V}_k^{\prime(m)} = \rho_k \tau_{N_r} \quad k \in \mathcal{K},$$

so that $x^\omega(\tau_{N_r}) = x^\omega(0)$. Moreover, we will show that

$$||\bar{X} - x^\omega||_{\bar{R}_r \vee \tau_{N_r}} \leq \left(\frac{2rK}{1-\rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\}. \tag{D.23}$$

To see (D.23), note that with the same initialization, $\bar{X}$ and x^ω follow identical trajectory until $\bar{V}_k^{\prime(m)} \neq \bar{V}_k^{(m)}$, for some $k \in \mathcal{K}$, $m \in \{1, ..., N_r\}$. Namely, $\bar{X}(t) = x^\omega(t)$ for all $t \in [0, \bar{A}_k^{(m)}]$, where $\bar{A}_k^{(m)}$ is the polling epoch of product k in the mth cycle. By construction, $|\bar{V}_k^{\prime(m)} - \bar{V}_k^{(m)}| = |\delta_k|$. Since the inventory level of product ℓ always decreases at rate λ_ℓ and increases at rate $\mu_\ell - \lambda_\ell$, it holds for product ℓ that

$$||\bar{X}_\ell - x_\ell^\omega||_{\bar{D}_k^{(m)} \vee \bar{D}_k^{\prime(m)}} \leq |\delta_k| \left(\lambda_\ell + \mu_\ell - \lambda_\ell\right) = |\delta_k| \mu_\ell, \quad \ell \in \mathcal{K},$$

and that

$$||\bar{X} - x^\omega||_{\bar{D}_k^{(m)} \vee \bar{D}_k^{\prime(m)}} \leq |\delta_k| \max_{k \in \mathcal{K}} \{\mu_k\},$$

where $\bar{D}_k^{(m)}$ and $\bar{D}_k^{\prime(m)}$ are the departure epochs of product k in the mth cycle, respectively.

After time $\bar{D}_k^{(m)}$ (or $\bar{D}_k^{\prime(m)}$), $\bar{X}$ and x^ω increase and decrease at the same rate over the same time intervals, until another perturbed busy time $\bar{V}_{\hat{k}}^{\prime(\hat{m})} \neq \bar{V}_{\hat{k}}^{(\hat{m})}$, for some $\hat{k} \in \mathcal{K}$, $\hat{m} \in \{1, ..., N_r\}$. Following similar derivation as above, the second perturbation can *further* enlarge the difference between $\bar{X}$ and x^ω (from time zero to the departure epoch after the busy time in consideration) by a maximum of $(|\delta_k| + |\delta_{\hat{k}}|) \max_{k \in \mathcal{K}} \{\mu_k\}$. Formally, we have

$$||\bar{X} - x^\omega||_{\bar{D}_{\hat{k}}^{(\hat{m})} \vee \bar{D}_{\hat{k}}^{\prime(\hat{m})}} \leq (2|\delta_k| + |\delta_{\hat{k}}|) \max_{k \in \mathcal{K}} \{\mu_k\}.$$

The same arguments continue to the end of the N_rth cycle. In particular,

$$\begin{aligned}
||\bar{X} - x^\omega||_{\bar{R}_r \vee \tau_{N_r}} &\leq \sum_{k \in \mathcal{K}} K |\delta_k| \max_{k \in \mathcal{K}} \{\mu_k\} \\
&\leq K \sum_{k \in \mathcal{K}} 2r \left(\frac{1}{\mu_k} + \frac{\rho_k}{1-\rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\} \quad \text{by (D.22)} \\
&= \left(\frac{2rK}{1-\rho} \sum_{k \in \mathcal{K}} \frac{1}{\mu_k} \right) \max_{k \in \mathcal{K}} \{\mu_k\},
\end{aligned}$$

and (D.23) is established.

Setting $d_2 := \left(\frac{2K}{1-\rho}\sum_{k\in\mathcal{K}}\frac{1}{\mu_k}\right)\max_{k\in\mathcal{K}}\{\mu_k\}$ gives (D.9) in Lemma 12. Q.E.D.

D.2. Proof of Lemma 5

Proof: We proceed the proof with the following steps: In **Step 1**, we apply Lemmas 3 and 11 to show that under the sequence of BSI policies with parameters defined in (C.7), the inventory level process $\{\bar{X}^n : n \geq 1\}$ first converges to the fluid model as setup times grow without bound, and the fluid model converges to a global limit cycle as time goes to infinity. In **Step 2**, we show that the stationary sequence of inventory levels converges to the same global limit cycle under the large-setup-time scaling. This is equivalent to interchanging the limits as in the first step, namely, first considering the steady-state of the stochastic system as time goes to infinity and then letting the setup times diverge. In **Step 3**, we employ the interchange of limits result in the first two steps and show that the statement of Lemma 5 holds under an interchange of limit and expectation. In **Step 4**, we justify the interchange of limit and expectation in the third step by proving that the sequence of fluid-scaled cumulative cost over one stationary cycle is UI under Assumption 1.

Step 1. Global limit cycle of the fluid limit. Consider the sequence of inventory levels $\{\bar{X}^n : n \geq 1\}$ under the sequence of BSI policies with base-stock levels $\{\mathbf{B}^n : n \geq 1\}$ and idle times $\{\mathbf{W}^n : n \geq 1\}$ in (C.7). Let x be the fluid inventory level process under the BSI policy with base-stock levels $\mathbf{b}$ and idle times $\mathbf{w}$.

As a direct consequence of Lemma 11, it holds that if $\bar{X}^n(0) \Rightarrow x(0)$ in $\mathbb{R}^K$, then

$$\tilde{X}^n(m) \Rightarrow \tilde{x}(m) \quad \text{as } n \to \infty \quad \text{for all } m \geq 0. \tag{D.24}$$

Moreover, it follows from Lemma 3 that

$$\tilde{x}(m) \to x_e(a_1) \quad \text{as } m \to \infty, \tag{D.25}$$

where x_e is the global limit cycle for x, and a_1 denotes a generic polling epoch of product 1 in equilibrium. Combining (D.24) and (D.25), we get that for any real-valued, continuous and bounded function f on $\mathbb{R}^K$,

$$\lim_{m\to\infty}\lim_{n\to\infty}\mathbb{E}\left[f\left(\tilde{X}^n(m)\right)\right] = f\left(x_e(a_1)\right). \tag{D.26}$$

Step 2. Interchange of limits. We now interchange the order of taking limits in (D.26), namely, we first consider the steady state of the stochastic process ($m \to \infty$), and then apply the large-setup-time scaling ($n \to \infty$) on the steady state. Note that the steady state in the first limit is well defined due to the stability of the BSI policy, namely, for each $n \geq 1$,

$$\tilde{X}^n(m) \Rightarrow \tilde{X}^n(\infty) \quad \text{as } m \to \infty,$$

where the random variable $\tilde{X}^n(\infty)$ denotes the stationary distribution of the DTMC $\{\tilde{X}^n(m) : m \geq 0\}$. The next lemma further implies that

$$\tilde{X}^n(\infty) \Rightarrow x_e(a_1) \quad \text{as } n \to \infty.$$

Lemma 13 (interchange of limits) *Consider the sequence of stochastic systems under the sequence of BSI policies with base-stock levels* $\{\mathbf{B}^n : n \geq 1\}$ *and idle times* $\{\mathbf{W}^n : n \geq 1\}$ *defined in* (C.7)*. For any real-valued, continuous and bounded function* f *on* $\mathbb{R}^K$*,*

$$\lim_{m\to\infty}\lim_{n\to\infty} \mathbb{E}\left[f\left(\tilde{X}^n(m)\right)\right] = \lim_{n\to\infty}\lim_{m\to\infty} \mathbb{E}\left[f\left(\tilde{X}^n(m)\right)\right] = f\left(x_e(a_1)\right), \tag{D.27}$$

where x_e *is the global limit cycle under the BSI policy with parameters* $\mathbf{b}$ *and idle times* $\mathbf{w}$*, and* a_1 *denotes a generic polling epoch of product 1 in the limit cycle.*

The proof of Lemma 13 follows from Lemma 5.2 in Hu et al. (2020a) with the following modification. Recall from Section 4.2 that the shortfall level process is equivalent to the queue length process under the exhaustive policy for polling systems. In addition, the exhaustive policy for polling systems is equivalent to the binomial-exhaustive policy with parameters $(1, \mathbf{1})$ in Hu et al. (2020a). The idle times and the setup times can be considered together as the switchover times for polling systems. As a consequence, Lemma 5.2 in Hu et al. (2020a) implies that (D.27) holds for the shortfall levels embedded at the polling epochs, namely, for $\tilde{Y}^n(m) := \bar{Y}^n(\bar{U}^{(m),n})$ and $y_e := b_k - x_e$, it holds that

$$\lim_{m\to\infty}\lim_{n\to\infty} \mathbb{E}\left[f\left(\tilde{Y}^n(m)\right)\right] = \lim_{n\to\infty}\lim_{m\to\infty} \mathbb{E}\left[f\left(\tilde{Y}^n(m)\right)\right] = f\left(y_e(a_1)\right). \tag{D.28}$$

Then, with $\tilde{Y}^n(m) = \mathbf{B}^n - \tilde{X}^n(m)$ for all $m \geq 0$ and $n \geq 1$, we derive from (D.28) that

$$\lim_{m\to\infty}\lim_{n\to\infty} \mathbb{E}\left[f\left(\tilde{X}^n(m)\right)\right] = \lim_{n\to\infty}\lim_{m\to\infty} \mathbb{E}\left[f\left(\tilde{X}^n(m)\right)\right], \tag{D.29}$$

and (D.27) follows from (D.26) and (D.29).

Step 3. Long-run average cost. For the nth stochastic system under the base-stock policy with base-stock levels $\mathbf{B}^n$ and idle times $\mathbf{W}^n$, we follow similar convention as in Appendix D.2 and use α^n to denote the stationary distribution (alternatively, a stationary distribution if there exist more than one closed classes) of the embedded DTMC $\{\tilde{X}^n(m) : m \geq 0\}$. In addition, we write $\mathbb{E}_{\alpha^n}$ as the expectation operator under which $\mathbb{P}(\tilde{X}^n(0) \in \cdot) = \alpha^n(\cdot)$, namely, the DTMC is initiated in stationarity. Moreover, let $\bar{T}^n$ denote the fluid-scaled length of a stationary production cycle, which is finite w.p.1 under any stable control. Define the fluid-scaled cumulated cost over one stationary cycle

$$\bar{\Psi}^n := r + \int_0^{\bar{T}^n} \psi(\bar{X}^n(s))ds. \tag{D.30}$$

For each $n \geq 1$, let $\bar{M}^n(t) := \max\{m \geq 1 : \bar{U}^{(m),n} \leq t\}$. The fluid-scaled long-run average cost for the nth system can be written as

$$\begin{aligned}
\bar{C}^n &= \frac{1}{t}\left(r\bar{M}^n(t) + \sum_{k\in\mathcal{K}} r_k \mathbf{1}_{\{\bar{A}^{(\bar{M}^n(t)+1),n} \leq t\}} + \int_0^t \psi(\bar{X}^n(s))ds\right) \\
&= \frac{\sum_{m=1}^{\bar{M}^n(t)}\left(r + \int_{\bar{U}^{(m-1),n}}^{\bar{U}^{(m),n}} \psi(\bar{X}^n(s))ds\right) + \int_{\bar{U}^{(\bar{M}^n(t)),n}}^{t} \psi(\bar{X}^n(s))ds + \sum_{k\in\mathcal{K}} r_k \mathbf{1}_{\{\bar{A}^{(\bar{M}^n(t)+1),n} \leq t\}}}{\sum_{m=1}^{\bar{M}^n(t)} \bar{T}^{(m),n} + (t - \bar{U}^{(\bar{M}^n(t)),n})} \\
&= \frac{\frac{1}{\bar{M}^n(t)}\sum_{m=1}^{\bar{M}^n(t)}\left(r + \int_{\bar{U}^{(m-1),n}}^{\bar{U}^{(m),n}} \psi(\bar{X}^n(s))ds\right)}{\frac{1}{\bar{M}^n(t)}\sum_{m=1}^{\bar{M}^n(t)} \bar{T}^{(m),n} + (t - \bar{U}^{(\bar{M}^n(t)),n})} + o(1) \to \frac{\mathbb{E}_{\alpha^n}\left[\bar{\Psi}^n\right]}{\mathbb{E}_{\alpha^n}\left[\bar{T}^n\right]} \qquad \text{w.p.1,}
\end{aligned} \tag{D.31}$$

where the convergence follows from the convergence of the DTMC to the stationary distribution α^n.

The derivation above applies to any sequence of base-stock policies with base-stock levels $\{\mathbf{B}^n : n \geq 1\}$ and idle times $\{\mathbf{W}^n : n \geq 1\}$ in (C.7). Thus, for the sequence of fluid-translated BSI policies $\{\pi_*^n : n \geq 1\}$, (D.31) implies that

$$\lim_{n\to\infty} \bar{C}^n_{\pi_*^n} = \lim_{n\to\infty} \frac{\mathbb{E}_{\alpha^n}\left[\bar{\Psi}^n_{\pi_*^n}\right]}{\mathbb{E}_{\alpha^n}\left[\bar{T}^n_{\pi_*^n}\right]} = \lim_{n\to\infty} \frac{1}{\tau_*}\mathbb{E}_{\alpha^n}\left[\bar{\Psi}^n_{\pi_*^n}\right], \tag{D.32}$$

where the second equality follows from Lemma 8, i.e., $\mathbb{E}_{\alpha^n}\left[\bar{T}^n_{\pi_*^n}\right] = (sn + w_* n)/(n(1-\rho)) = \tau_*$.

Furthermore, if we can interchange the order of taking limit and expectation on the right-hand side of (D.32), then

$$\lim_{n\to\infty} \frac{1}{\tau_*}\mathbb{E}_{\alpha^n}\left[\bar{\Psi}^n_{\pi_*^n}\right] = \frac{1}{\tau_*}\mathbb{E}_{\alpha}\left[\lim_{n\to\infty} \bar{\Psi}^n_{\pi_*^n}\right] = \frac{1}{\tau_*}\left(r + \int_0^{\tau_*} \psi\left(x_*(s)\right) ds\right) = c_*, \tag{D.33}$$

where the second equality follows from the interchange of limits in Lemma 13.

The statement of Lemma 5 follows by combining (D.32) and (D.33). The work left is to justify (D.33).

Step 4. Uniform integrability. We now complete the proof by justifying the interchange of limit and expectation in the first equality in (D.33). Let Θ_k, $k \in \mathcal{K}$, denote a generic busy period "generated" by one unit of product k, namely, Θ_k is the amount of time it takes to reduce the shortfall level Y^n of product k by one. Due to the facts that (i) the shortfall level process is equivalent to the queue length process under the exhaustive policy for polling systems, (ii) the exhaustive policy is a special case of the binomial-exhaustive policy, and (iii) the sum of the idle and setup times can be considered as the switchover times for polling systems, the following lemma follows directly from Theorem 4 and Proposition 6.1 in Hu et al. (2020b).

Lemma 14 *Consider the sequence of stochastic systems under the sequence of BSI policies with base-stock levels* $\{\mathbf{B}^n : n \geq 1\}$ *and idle times* $\{\mathbf{W}^n : n \geq 1\}$ *defined in* (C.7). *Assume that* $\bar{X}^n(0) \stackrel{\mathrm{d}}{=} \tilde{X}^n(\infty)$ *for all* $n \geq 1$, *namely, the inventory level process is stationary. Let* x_e *denote the global limit cycle of the fluid model under the BSI policy with base-stock levels* $\mathbf{b}$ *and idle times* $\mathbf{w}$, *and* y_e *be the shortfall level process corresponding to* x_e. *Under Assumption 1, the steady-state shortfall level satisfies* $\mathbb{E}\left[Y_k^n(A_i^n)^\ell\right] < \infty$ *for all* $n \geq 1$, *and*

$$\lim_{n\to\infty} \mathbb{E}\left[\bar{Y}_k^n(\bar{A}_i^n)^\ell\right] = (y_{e,k}(a_i))^\ell \quad \textit{for all } \ell \in \mathbb{N}, \quad k, i \in \mathcal{K}.$$

Furthermore, the steady-state busy times satisfy $\mathbb{E}\left[\left(\bar{V}_k^n\right)^\ell\right] < \infty$ *for all* $n \geq 1$ *and*

$$\lim_{n\to\infty} \mathbb{E}\left[\left(\bar{V}_k^n\right)^\ell\right] = (y_{e,k}(a_k)\mathbb{E}\left[\Theta_k\right])^\ell \quad \textit{for all } \ell \in \mathbb{N}, \quad k \in \mathcal{K}.$$

By Assumption 2, the holding and backlog costs satisfy $h_k(x) = O(x^p)$ and $p_k(x) = O(x^p)$, for some $p > 1$. For the purpose of this proof, it is sufficient to restrict attention to cost functions that are polynomial of order p, namely, there exist constants c_{h_k} and c_{p_k} in $\mathbb{R}_+$ such that $h_k(x_k) = c_{h_k}(x_k^+)^p$ and $p_k(x_k) = c_{p_k}(x_k^-)^p$, $k \in \mathcal{K}$. The cumulated cost over one stationary cycle defined in (D.30) can be upper bounded by

$$\begin{aligned}
\bar{\Psi}^n &= r + \sum_{k\in\mathcal{K}} \int_0^{\bar{T}^n} \left[c_{h_k}\left(\bar{X}_k^n(s)^+\right)^p + c_{p_k}\left(\bar{X}_k^n(s)^-\right)^p\right] ds \\
&\le r + \sum_{k\in\mathcal{K}} \int_0^{\bar{T}^n} \left[c_{h_k}\left(\bar{B}_k^n\right)^p + c_{p_k}\left(\bar{Y}_k^n(s)\right)^p\right] ds \\
&\le r + \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{K}} \int_{\bar{A}_i^n}^{\bar{A}_{i+1}^n} \left[c_{h_k}\left(\bar{B}_k^n\right)^p + c_{p_k}\left(\bar{Y}_k^n(\bar{A}_i^n) + \bar{P}_k^n(\bar{V}_i^n + \bar{W}_i^n + \bar{S}_{i+1}^n)\right)^p\right] ds \\
&= r + \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{K}} \left(\bar{V}_i^n + \bar{W}_i^n + \bar{S}_{i+1}^n\right)\left[c_{h_k}\left(\bar{B}_k^n\right)^p + c_{p_k}\left(\bar{Y}_k^n(\bar{A}_i^n) + \bar{P}_k^n(\bar{V}_i^n + \bar{W}_i^n + \bar{S}_{i+1}^n)\right)^p\right] \\
&= r + \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{K}} \left(\bar{V}_i^n + w_i + \bar{S}_{i+1}^n\right)\left[c_{h_k}\left(\lceil b_k \rceil\right)^p + c_{p_k}\left(\bar{Y}_k^n(\bar{A}_i^n) + \bar{P}_k^n(\bar{V}_i^n + w_i + \bar{S}_{i+1}^n)\right)^p\right].
\end{aligned} \tag{D.34}$$

In (D.34) above, the first inequality is due to the relationship that $X_k^n(t) = B_k^n - Y_k^n(t)$ for all $k \in \mathcal{K}$ and $t \ge 0$. The second inequality is due to the omission of the production process, which decreases the value of $\bar{Y}^n$. The second-to-last equality follows from the fact that the time between two consecutive polling epochs, e.g., $\bar{A}_i^n$ and $\bar{A}_{i+1}^n$, consists of the corresponding production time, idle time, and the subsequent setup time. Finally, the last equality follows from the translation in (C.7) that $\mathbf{B}^n := \lceil n\mathbf{b} \rceil$ and $\mathbf{W}^n := n\mathbf{w}$.

Based on Assumption 1 and Lemma 14, it follows from standard arguments that each summand on the right-hand side of (D.34) is UI (see similar derivation in the proof of Theorem 2 in Hu et al. (2020a)). Since UI is preserved under finite summation, $\{\bar{\Psi}^n : n \ge 1\}$ is UI. Thus, the first equality in (D.33) follows by taking the sequence of BSI policies to be $\{\pi_{BSI,*}^n : n \ge 1\}$. Q.E.D.

D.3. Proof of Lemma 6

Proof: To establish the result, we first prove the following lemma, whose proof follows similar steps as in the proof of Lemma 5. In particular, **Step 1**, **Step 2**, and **Step 3** follow identical arguments as those in the proof of Lemma 5. We now explain how to modify **Step 4**. First, an analogue of Lemma 14 on the steady-state shortfall level and busy times holds under Assumption 3 follows again from Theorem 4 and Proposition 6.1 in Hu et al. (2020b).

Lemma 15 *Consider the sequence of stochastic systems under the sequence of BSI policies with base-stock levels $\{\mathbf{B}^n : n \ge 1\}$ and idle times $\{\mathbf{W}^n : n \ge 1\}$ defined in* (C.7)*. Assume that $\bar{X}^n(0) \stackrel{\mathrm{d}}{=} \tilde{X}^n(\infty)$ for all $n \ge 1$, namely, the inventory level process is stationary. Let x_e denote the global limit cycle of the fluid model under the BSI policy with parameters $\mathbf{b}$ and idle times $\mathbf{w}$, and y_e be the shortfall level corresponding to x_e. Under Assumption 3, the steady-state shortfall level satisfies $\mathbb{E}\left[Y_k^n(A_i^n)Y_j^n(A_i^n)\right] < \infty$ for all $n \ge 1$, and*

$$\lim_{n\to\infty} \mathbb{E}\left[\bar{Y}_k^n(\bar{A}_i^n)\bar{Y}_j^n(\bar{A}_i^n)\right] = y_{e,k}(a_i)y_{e,j}(a_i), \quad \textit{for all } k, j, i \in \mathcal{K}.$$

Furthermore, the steady-state busy times satisfy $\mathbb{E}\left[\left(\bar{V}_k^n\right)^2\right] < \infty$ for all $n \ge 1$ and

$$\lim_{n\to\infty} \mathbb{E}\left[\left(\bar{V}_k^n\right)^2\right] \to \left(y_{e,k}(a_k)\mathbb{E}\left[\Theta_k\right]\right)^2, \quad \textit{for all } k \in \mathcal{K}.$$

With linear holding and backlog costs (i.e., $p=1$), the upper bound in (D.34) reduces to

$$\bar{\Psi}^n \le r + \sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{K}} \left(\bar{V}_i^n + w_i + \bar{S}_{i+1}^n\right)\left[c_{h_k}\lceil b_k\rceil + c_{p_k}\left(\bar{Y}_k^n(\bar{A}_i^n) + \bar{P}_k^n(\bar{V}_i^n + w_i + \bar{S}_{i+1}^n)\right)\right]. \tag{D.35}$$

Using Assumption 3 and Lemma (15), we can apply similar procedures as in the proof of Lemma 5 to conclude that each summand on the right-hand side of (D.35) is UI, and the statement follows. Q.E.D.

Appendix E: More on the Numerical Experiments in Sections 5 and 6

E.1. Fluid Optimal Solutions in Table 2

In this section, we formalize the observation that several systems in Table 2 (e.g., systems 1–4 and systems 5–6) have the same fluid-optimal base-stock levels and costs.

Lemma 16 *Let P_1 be a fluid optimization problem (2.3) with setup times $(s_1,...,s_K)\in\mathbb{R}_+^K$. Let $(\mathbf{b}_*, w_*)$ denote an optimal solution to P_1. Consider another fluid optimization problem P_2 that is identical to P_1, except that P_2 has setup times $(\tilde{s}_1,...,\tilde{s}_K)$, where $(\tilde{s}_1,...,\tilde{s}_K)$ satisfies $0\le \sum_{k\in\mathcal{K}}\tilde{s}_k - \sum_{k\in\mathcal{K}} s_k \le w_*$. Then $(\mathbf{b}_*, w_*-\Delta)$ is optimal to P_2 for $\Delta := \sum_{k\in\mathcal{K}}\tilde{s}_k - \sum_{k\in\mathcal{K}} s_k$.*

Proof: Let $c^{P_1}(\mathbf{b},w)$ (alternatively, $c^{P_2}(\mathbf{b},w)$) denote the objective value of problem P_1 (alternatively, P_2) for a given solution $(\mathbf{b},w)$.

Since $0\le\Delta\le w_*$, it holds that $w_*-\Delta\ge 0$, which implies that $(\mathbf{b}_*, w_*-\Delta)$ is a feasible solution to P_2. Suppose for the sake of contradiction that $(\mathbf{b}_*, w_*-\Delta)$ is not optimal to P_2, i.e., there exists solution $(\mathbf{b}'_*, w'_*)$, $\mathbf{b}'_*\ne\mathbf{b}_*$ or $w'_*\ne w_*$, whose objective value satisfies

$$c^{P_2}(\mathbf{b}'_*, w'_*) < c^{P_2}(\mathbf{b}_*, w_*-\Delta). \tag{E.1}$$

Define $w''_* := w'_* + \Delta \ge 0$. Note that $(\mathbf{b}'_*, w''_*)$ is a feasible solution to P_1.

Inspection of the fluid optimization problem (2.3) implies that the setup times and the total idle time affect the solution only through their sum. Thus, for problem P_1 and P_2, we have $c^{P_1}(\mathbf{b},w_1) = c^{P_2}(\mathbf{b},w_2)$ if $w_1 + \sum_{k\in\mathcal{K}} s_k = w_2 + \sum_{k\in\mathcal{K}}\tilde{s}_k$, i.e., if $w_1 = w_2+\Delta$, for some $\mathbf{b}\in\mathbb{R}^K$, $w_1, w_2\in\mathbb{R}_+$. In particular, we have

$$c^{P_1}(\mathbf{b}'_*, w''_*) = c^{P_2}(\mathbf{b}'_*, w'_*) \quad\text{and}\quad c^{P_1}(\mathbf{b}_*, w_*) = c^{P_2}(\mathbf{b}_*, w_*-\Delta) \tag{E.2}$$

By (E.1) and (E.2), we get that

$$c^{P_1}(\mathbf{b}'_*, w''_*) = c^{P_2}(\mathbf{b}'_*, w'_*) < c^{P_2}(\mathbf{b}_*, w_*-\Delta) = c^{P_1}(\mathbf{b}_*, w_*),$$

which contradicts the optimality of $(\mathbf{b}_*, w_*)$ to problem P_1. Therefore, $(\mathbf{b}_*, w_*-\Delta)$ is optimal to P_2. Q.E.D.

To interpret the fluid optimal solutions for systems 1–4 in Table 2, we can consider the fluid optimization problem for system 1 as P_1 and for systems 2, 3, 4 as P_2 in the context of Lemma 16, respectively. Similarly, to interpret the fluid optimal solutions for systems 5 and 6 in Table 2, we can consider the fluid optimization problem for system 5 as P_1 and for systems 6 as P_2.

E.2. MDP Formulation

In this section, we provide detailed formulation for MDP Idle-Fix and MDP General in Section 6. We use the relative value iteration method (see, e.g., (Das et al. 1999, Section 2)) with state truncation to solve MDP Idle-Fix and MDP General. In particular, we consider a two-product system where the inventory level is kept between -40 and 40. The transition rates and available actions are modified such that the inventory level does not exceed the boundary. In our setting, if $X_1 = 40$, the action to produce product 1 is not available; if $X_1 = -40$, the demand rate for product 1 is set to $\lambda_1 = 0$. Similar treatment is applied to product 2.

E.2.1. MDP Idle-Fix Let G denote the state space, i.e., $G = \{(X_1, X_2, Z) \in \mathbb{Z} \times \mathbb{Z} \times \{1, 2\} : -40 \le X_i \le 40, i = 1, 2\}$. The action space is $\mathcal{A}(X_1, X_2, Z) = \{$"produce", "idle + switch"$\}$ if $Z \in \{1, 2\}$. For $i, j \in G$, $a \in \mathcal{A}(i)$, let $P_{ij}(a)$ denote the probability of entering state j if action a is taken at state i, and let $c(i, j, a)$ denote the expected accrual of cost if the state transitions from i to j under action a. The relative value iteration is summarized as follows:

1. Initiate $V^0 = \mathbf{0}$, choose some (arbitrary) state $x^* \in G$, specify $\epsilon > 0$, and set $m = 0$.
2. For each $i \in G$, compute $V^{m+1}(i)$ by
$$V^{m+1}(i) = \min_{a \in \mathcal{A}(i)} \left\{ \sum_{j \in G} P_{ij}(a) \left[c(i, j, a) + V^m(j) \right] - V^m(x^*) \right\}. \tag{E.3}$$
3. If $\mathrm{sp}(V^{m+1} - V^m) < \epsilon$, go to step 4. Otherwise increment m by 1 and return to step 2. Here, sp denotes "span," which is defined as $\mathrm{sp}(V) := \max_{i \in G} V(i) - \min_{i \in G} V(i)$.
4. For each $i \in G$, retrieve the ϵ-optimal action in steady state, denoted by $a_\epsilon(i)$, by selecting
$$a_\epsilon(i) \in \arg\min_{a \in \mathcal{A}_i} \left\{ \sum_{j \in G} P_{ij}(a) \left[c(i, j, a) + v^m(j) \right] - v^m(x^*) \right\}.$$

In (E.3), the term $\sum_{j \in G} P_{ij}(a) c(i, j, a)$ is the expected accrual of cost during the sojourn time at state i if action a is taken. For $i = (X_1, X_2, Z)$, it can be calculated as follows.

1. If $a =$ "produce," then
$$P_{ij}(a) = \begin{cases} \frac{\lambda_1}{\lambda_1 + \lambda_2 + \mu_1} & \text{if } j = (X_1 - 1, X_2, Z) \\ \frac{\lambda_2}{\lambda_1 + \lambda_2 + \mu_1} & \text{if } j = (X_1, X_2 - 1, Z) \\ \frac{\mu_1}{\lambda_1 + \lambda_2 + \mu_1} & \text{if } j = (X_1 + 1, X_2, Z) \\ 0 & \text{otherwise} \end{cases}$$
$$c(i, j, a) = \frac{1}{\lambda_1 + \lambda_2 + \mu_1} \sum_{k \in \mathcal{K}} \left[h_k (X_k)^+ + p_k (X_k)^- \right].$$
2. If $a =$ "idle + switch," then
$$P_{ij}(a) = \begin{cases} \mathrm{Pois}(n_1, \lambda_1(s_{Z^c} + w_*/2)) \mathrm{Pois}(n_2, \lambda_2(s_{Z^c} + w_*/2)) & \text{if } j = (X_1 - n_1, X_2 - n_2, Z^c) \\ 0 & \text{otherwise,} \end{cases}$$
where $\mathrm{Pois}(n, \lambda)$ denotes the probability mass function of a Poisson random variable with count n and mean λ. Using the fact that conditional on count $= n$, the n (unordered) arrival times are i.i.d. uniform random variables over the idle and setup time, the expected accrual of cost is given by, for $k = 1, 2$,
$$c_k(i, j, a) := \begin{cases} \frac{1}{2} h_k (X_k + X_k - n_k)(s_{Z^c} + w_*/2) & \text{if } X_k - n_k \ge 0 \\ -\frac{1}{2} p_k (X_k + X_k - n_k)(s_{Z^c} + w_*/2) & \text{if } X_k \le 0 \\ \frac{1}{2} h_k X_k \frac{X_k (s_{Z^c} + w_*/2)}{X_k - (X_k - n_k)} + \frac{1}{2} p_k (n_k - X_k) \frac{(n_k - X_k)(s_{Z^c} + w_*/2)}{X_k - (X_k - n_k)} & \text{if } X_k \ge 0, X_k - n_k \le 0 \\ 0 & \text{otherwise,} \end{cases}$$

and $c(i,j,a) = c_1(i,j,a) + c_2(i,j,a)$.

E.2.2. MDP General Let G denote the state space, i.e., $G = \{(X_1, X_2, Z) \in \mathbb{Z} \times \mathbb{Z} \times \{1, 2, 1^I, 2^I\} : -40 \le X_k \le 40, k = 1, 2\}$. The action space is $\mathcal{A}(X_1, X_2, Z) = \{$"produce", "switch", "idle"$\}$ if $Z \in \{1, 2\}$, and $\mathcal{A}(X_1, X_2, Z) = \{$"switch", "idle"$\}$ if $Z \in \{1^I, 2^I\}$. Similar to the case of MDP Idle-Fix, we can apply the relative value iteration algorithm, where the cost and transition probability are characterized as follows.

1. If $a =$ "produce," then

$$P_{ij}(a) = \begin{cases} \frac{\lambda_1}{\lambda_1+\lambda_2+\mu_1} & \text{if } j = (X_1 - 1, X_2, Z) \\ \frac{\lambda_2}{\lambda_1+\lambda_2+\mu_1} & \text{if } j = (X_1, X_2 - 1, Z) \\ \frac{\mu_1}{\lambda_1+\lambda_2+\mu_1} & \text{if } j = (X_1 + 1, X_2, Z) \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad c(i,j,a) = \frac{1}{\lambda_1 + \lambda_2 + \mu_1} \sum_{k=1,2} \left[h_k (X_k)^+ + p_k (X_k)^- \right].$$

2. If $a =$ "switch," define $Z^c := 1$ if $Z \in \{2, 2^I\}$, and $Z^c := 2$ if $Z \in \{1, 1^I\}$. Then

$$P_{ij}(a) = \begin{cases} \text{Pois}(n_1, \lambda_1 s_{Z^c}) \text{Pois}(n_2, \lambda_2 s_{Z^c}) & \text{if } j = (X_1 - n_1, X_2 - n_2, Z^c) \\ 0 & \text{otherwise.} \end{cases}$$

For $k = 1, 2$, define

$$c_k(i,j,a) := \begin{cases} \frac{1}{2} h_k (X_k + X_k - n_k) s_{Z^c} & \text{if } X_k - n_k \ge 0 \\ -\frac{1}{2} p_k (X_k + X_k - n_k) s_{Z^c} & \text{if } X_k \le 0 \\ \frac{1}{2} h_k X_k \frac{X_k s_{Z^c}}{X_k - (X_k - n_k)} s_{Z^c} + \frac{1}{2} p_k (n_k - X_k) \frac{(n_k - X_k) s_{Z^c}}{X_k - (X_k - n_k)} & \text{if } X_k \ge 0, X_k - n_k \le 0 \\ 0 & \text{otherwise.} \end{cases}$$

Then the expected cumulated cost is given by $c(i,j,a) = c_1(i,j,a) + c_2(i,j,a)$.

3. If $a =$ "idle," define $Z^c := 1^I$ if $Z \in \{1, 1^I\}$, and $Z^c := 2^I$ if $Z \in \{2, 2^I\}$. Then

$$P_{ij}(a) = \begin{cases} \frac{\lambda_1}{\lambda_1+\lambda_2} & \text{if } j = (X_1 - 1, X_2, Z^c) \\ \frac{\lambda_2}{\lambda_1+\lambda_2} & \text{if } j = (X_1, X_2 - 1, Z^c) \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad c(i,j,a) = \frac{1}{\lambda_1 + \lambda_2} \sum_{k=1,2} \left[h_k (X_k)^+ + p_k (X_k)^- \right].$$