Submitted to *Operations Research*
manuscript (Please, provide the manuscript number!)

Optimal Policies for a Dual-Sourcing Inventory Problem with Endogenous Stochastic Lead Times

Jing-Sheng Song
The Fuqua School of Business, Duke University, Durham, NC 27708, USA, jssong@duke.edu

Li Xiao
CUHK Business School, The Chinese University of Hong Kong, Shatin, Hong Kong, xiaoli@baf.cuhk.edu.hk

Hanqin Zhang
Business School, The National University of Singapore, 119245, Singapore, bizzhq@nus.edu.sg

Paul Zipkin
The Fuqua School of Business, Duke University, Durham, NC 27708, USA, paul.zipkin@duke.edu

We consider a single-product, two-source inventory system with Poisson demand and backlogging. Inventory can be replenished through a normal supply source, which consists of a two-stage tandem queue with exponential production time at each stage. We can also place an emergency order by skipping the first stage, for a fee. There is no fixed order cost. There are linear order, holding and backorder costs. Through a new approach, we obtain optimal ordering policies for the discounted or long-run average cost, and also characterize near-optimal heuristic policies. The approach consists of four steps. The first step is to establish an equivalent system, in the sense that it has the same optimal policy as the original system. The second step is to construct a tandem queueing system, where costs are charged in accord with the equivalent system's cost structure. The third step derives an optimal control of the service rate at each server so as to minimize the tandem queue's system-wide cost. The fourth and final step is to translate the queue's optimal policy to an optimal policy for the equivalent system and hence the original system.

Key words: dual-source, stochastic lead time, inventory policy, dynamic programming, tandem queue, optimal control

History: April 30, 2015; revised February 16, 2016; June 12, 2016; August 8, 2016

---

1. Introduction

This paper analyzes a single-product, two-source inventory system. The system is standard in several ways: Demands follow a Poisson process and all stockouts are backlogged. Inventory can be replenished from either a normal source or an emergency source. The normal source has a longer lead time than the emergency source but a lower unit purchase cost. There is no fixed order cost. There are a linear order cost for each source and linear holding and backorder costs at the demand point. We consider an infinite time-horizon with the discounted-cost or the average-cost objective. The novel feature in this system is the supply sources: The normal source consists of a two-stage tandem queue with exponential production time at each stage; the emergency source skips the first stage of the queue. See Figure 1. In addition to the usual holding cost at the inventory location, there is a holding cost at each stage. As will be detailed later, these new modeling features enable us to derive structural results and insights that other modeling choices cannot. Our approach thus opens the door for further development of this line of research.

Our goal is to find an optimal order policy that minimizes the discounted or long-run average system cost. As we shall see, a base-stock policy is not optimal here. Decisions, including when to order, which source to use, and how much to order, critically rely on the system status at both sources as well as the demand point.

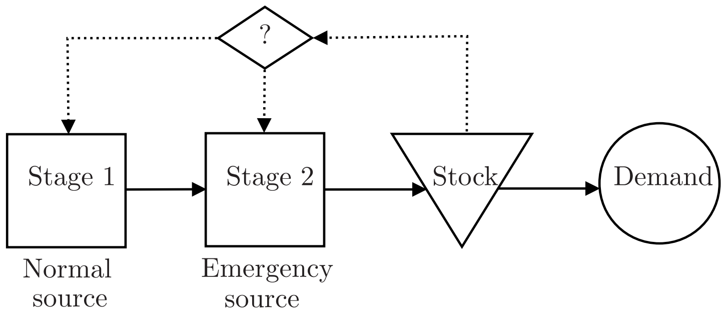

Figure 1 Two-Source Inventory System

In practice, multi-source systems provide many advantages and have been widely adopted. One prominent advantage is to assist companies to hedge against the risks of disruptions and long delays in their suppliers' production or transportation processes. Suppose we view the company's own production (e.g., final assembly) as the second stage and the supplier's production (e.g., components or subassemblies) and transportation as the first stage. In normal circumstances, it is sufficient to use the normal source to provide the inputs to produce and build up the inventory of finished goods. However, when the supplier suffers a disruption, he can no longer deliver the components that the company needs on time. The company itself has to buy the components or subassemblies, either from other suppliers or a spot market, paying a higher price.

For example, in 2011, when disastrous floods hit Thailand, many hard-disk-drive factories, such as Western Digital and Toshiba, were significantly impacted. Their operations were suspended and their abilities to meet customer demands were damaged. Because Thailand is the second-largest exporter of hard disk drives, most of which end up in personal computers, computer makers feared that the supply disruption would heavily affect their production of computers. They had to resort to other providers or the spot market. Drive prices nearly doubled all over the world. See Makan and Simon (2011) and Wikipedia (2011) for more details. Here, a U.S. assembly plant can be viewed as stage 2 in our model, while the Thailand factory, together with the transportation process from Thailand to the U.S. plant, is represented by the first stage. The transportation time from the emergency source (such as the spot market) is considered negligible, compared to that from the normal source.

Another valuable feature of dual sourcing is to help a company realize an agile, responsive manufacturing process, to meet surges in market demand. For example, in the cosmetics industry, it's hard to predict customer tastes, especially in emerging markets. Often, when one product becomes popular, its inventory runs out quickly. To capture such fast-moving demands, companies aim to adjust production rates as much as possible. Recently, L'Oréal, the world's largest cosmetics company, implemented an effective dual-sourcing strategy. Most products are produced at their own factories, but suppliers provide packaging and components. Their new platform collects data, including inventories and production capacities from both their suppliers and their own factories. Then they use these real-time data to determine when to ramp up production, and also to notify relevant suppliers, who may also need to increase production. When a supplier is unable to respond sufficiently, L'Oréal seeks alternative sources, see Schectman (2013). Similarly, many companies employ dual transportation modes for material supply to respond to demand surges; see Veeraraghavan and Scheller-Wolf (2008), who discuss several examples including Caterpillar and Amazon. In this context, we may view the cheaper, standard, land or sea transportation process as the first stage in our model, while normalizing the expedited air shipping time to zero, so that after paying the expediting fee, the material directly arrives at stage 2. The second stage may be a manufacturing facility or a port plus local transportation. If the material is bought from overseas, the processing time at the port may include custom clearance, among other things.

To reap the full benefit of the dual-sourcing strategy, managers need to decide how exactly to deploy the two sources when making replenishment orders. The problem turns out to be quite complex and has drawn attention from many researchers over several decades; see §2 for more detail. The majority of the literature assumes *constant* lead times at both sources. Even under such a simplification, the form of the optimal policy lacks a clear structure when the difference between the lead times of the two sources is more than one period. Thus, previous studies concentrate on simple heuristic policies, such as single-index and dual-index policies. Here, we develop and analyze a different, though still stylized, model of a dual-source system. In this model, the lead times at both sources are stochastic and endogenous, namely, sojourn times in a queueing system. The lead-time difference here is the time spent at the second stage, which is stochastic and can be of arbitrary length. The system state is a three-dimensional vector, the inventory level, and the numbers of orders waiting or being processed at each of the two stages. At any time, based on this information, we need to decide whether or not to order, which source to use, and how much to order. Although the system is fairly complex, we are able to identify an optimal policy, the first of its kind in the literature. Our analysis provide several insights for understanding dual-sourcing systems.

Our main results are as follows:

• We fully characterize the optimal order policy under a certain condition on the system parameters. The policy is specified by a number called the threshold and a switching curve given by a function of the number of orders at the second stage. There are several cases. When there are orders at the first stage and no order at the second stage, order directly from the second stage, if the inventory level drops below the threshold; when there is no order at the first stage but there are orders at the second stage, place an order to the first stage, if the inventory level is below the switching curve; when there is no order at either stage, order from both stages, if the inventory level drops below the threshold, and place an order to the first stage only, if the inventory level lies between the threshold and the switching curve; in all other cases, do not order.

• Some properties of the threshold and switching curve are established. They help us to better understand the impact of the system parameters on the optimal policy.

• For cases where the parameter condition does not hold, we propose a heuristic policy. Numerical experiments show that the heuristic policy performs quite well. We compare the policy to two other heuristics, a dual-index policy and the tailored base-surge policy of Allon and Van Mieghem (2010). The results show that our policy performs better than the others in nearly all cases.

To establish these results, we take a fairly novel approach: The basic idea is to reduce the original three-dimensional system to an equivalent and tractable two-dimensional system. To do this, we proceed in four steps. Let **O** denote the original system. First, we change the way to charge the order cost at the first stage to generate a new system, called System **M**. We show that its optimal policy is the same as that of System **O**, and the optimal costs of the two systems have a very simple relationship. Second, we reduce the three-dimensional System **M** to a two-dimensional system, called System **A**. To construct System **A**, we change the action of placing orders at the first stage of System **M** to controlling its service rate. Thus, we give more flexibility in using the normal source in System **A**. This change simplifies the system dynamics of System **A**. Third, we fully characterize the optimal policy of System **A**. Fourth and last, we show that the optimal policy of System **A** is also optimal for System **O**, under a certain condition on the parameters.

Another insight from our analysis is the role of pipeline inventory holding costs in the optimal policy. With a single source and backorders, the expected pipeline inventory cost is a constant, and thus does not affect the policy. This is no longer true in dual sourcing. It seems that this effect is not considered in the extant literature.

In our model, the single exponential server at each stage captures the randomness in supply lead times. It is not unusual that a shipment needs to go through a transportation network, possibly including multiple transportation modes, such as rail, ocean freight, and trucks, where congestion may occur. Indeed, according to a 2014 survey conducted by the International Air Transport Association (IATA) to major freight forwarders and their customers, low transport time reliability is perceived as the second most important factor (next to transportation cost). Using a half-year of air cargo data between 2012 and 2013 provided by a leading forwarder on 1336 routes served by 20 airlines, Shang et al. (2014) demonstrate that air cargo transport delay follows a multimodal distribution. Corman and Baertlein (2014) document data on long shipment delays at the U.S. West Coast ports due to congestion. Modeling the entire production and transportation process as an aggregate single-server node is a coarse approximation, but reasonable as a starting point for tractability and for gaining qualitative insights. (See Section 7 for a discussion of possible extensions to more general lead times.)

2. Related Literature

An overview of the literature on dual sourcing is in order. Studies of the form of optimal policy mostly focus on period-review systems with deterministic lead times. Barankin (1961) studies the single-period case. Fukuda (1964) extends it to the multi-period case, where the normal and emergency lead times differ by only one period. They prove that a single-index base-stock policy is optimal. This policy has two base-stock levels, one for each source, with the emergency base-stock level lower than the normal one. The single index is the inventory position of the system, which is the net inventory plus the total outstanding orders in both sources. The policy works mostly as a regular base-stock policy; in each period it orders enough to keep the inventory position at the normal base-stock level. It orders from the normal source unless the inventory position drops below the emergency base-stock level. In that case, it first orders up to the emergency base-stock from the emergency source and then orders the rest from the normal source. When demand information can be updated and fixed order costs are added, Sethi et al. (2003) prove that a single-index (s, S) policy is optimal. However, if the lead times differ more than one period, assuming no fixed order cost, Whittemore and Saunders (1977) show that the optimal policy is neither base-stock nor a simple function of backorders. By constructing an equivalent single source lost sales problem with a positive lead time, Sheopuri et al. (2010) show that the optimal policy does not have a simple structure; the optimal order quantities from the two sources are functions of the entire state vector. More recently, Hua et al. (2015) use $L^\natural$-convexity to derive some structural properties of the optimal policy. They show that the optimal normal order is more sensitive to the late-to-arrive outstanding orders, but the optimal emergency order is more sensitive to the soon-to-arrive outstanding orders. Li and Yu (2014) obtain related results using multimodularity. Lawson and Porteus (2000) study a related but different problem. They consider a periodic-review serial inventory system, where the regular lead time between adjacent stages is one period, but a shipment can be expedited instantaneously at a higher cost, so the lead-time difference between the two transportation modes is one period. They identify an optimal policy for each stage to decide how much to ship in by the regular lead time and how much to expedite in each period. In this setting, the amount to expedite from a stage depends on the inventory level at that stage, a feature absent in the dual-source problem here.

Since the optimal policy is complex for general lead times, several studies have developed heuristic policies. Scheller-Wolf et al. (2007) adopt the above single-index base-stock policy as a heuristic and develop an efficient computational method to find the best base-stock levels. Veeraraghavan and Scheller-Wolf (2008) propose a dual-index base-stock policy. This policy also has two base-stock levels as in the single-index policy. However, now it has another index – the inventory position including only the emergency source. It works mostly like the single-index policy, except that it orders from the emergency source whenever the emergency inventory position is below the lower base-stock level. In that case, it orders up to the lower base-stock level from the emergency source and the rest from the normal source. The use of the second index keeps the responsive stock under closer watch. Using numerical experiments, these authors demonstrate that the best dual-index base-stock policy performs very well compared to the true optimal policy (identified by complete search). Both Sheopuri et al. (2010) and Hua et al. (2015) develop heuristics that employ a base-stock policy for one source and some bounds to determine the order quantity for the other source.

In a continuous-review setting with deterministic lead times, Moinzadeh and Nahmias (1988) develop an extension of the standard (Q, R) policy to approximate problems with fixed costs. Assuming Poisson demand and no fixed order cost, Moinzadeh and Schmidt (1991) analyze system performance and optimization under the dual-index base-stock policy. Song and Zipkin (2009) consider systems with random lead times, including i.i.d. stochastic lead times, exogenous and sequential stochastic lead times, and endogenous stochastic supply networks. They also consider more than two supply sources. For stochastic lead times, Allon and Van Mieghem (2010) propose a tailored base-surge (TBS) policy with two parameters, which orders a constant quantity from the normal source and uses a base-stock policy to govern the orders from the emergency source.

Assuming renewal demand and supply processes, they develop closed form expressions for the asymptotically optimal policy parameters in heavy traffic. An earlier appearance of the TBS policy can be found in Rosenshine and Obee (1976) in a period-review system with a deterministic lead time at the normal source and zero lead time at the emergency source. Janakiraman et al. (2015) adopt this policy form in the periodic-review setting with general deterministic lead times and show that the best *TBS* policy is close to optimal if demand includes a random surge with a small probability. They also show that the performance of the best *TBS* policy improves as the lead time difference between the two sources grows. Furthermore, assuming deterministic lead times, Xin and Goldberg (2015) prove that a *TBS* policy is asymptotically optimal, when the lead time from the normal source grows large while the lead time from the emergency source is fixed. More recently, Arts et al. (2016) study a system in which the lead time from the emergency source is deterministic while the lead time from the normal source is the sum of a phase-type random variable and the lead time from the emergency source. The demand is modeled as a Markov modulated Poisson process. Assuming the normal ordering decision follows a fixed base-stock policy, they partially characterize the optimal order policy from the emergency source.

The system we consider here is a special case of the endogenous stochastic supply networks in Song and Zipkin (2009), i.e., the supply network at each source is a single exponential server queue. Unlike Song and Zipkin (2009) who assume a dual index base-stock policy, in the present paper we identify an optimal policy among all possible policies. Song and Zipkin (2009) conjecture that their policy is optimal for some systems including ours. The results here do *not* support this conjecture. To gain deeper insights, we also compare the true optimal policy with the best dual-index policy.

The methodology we use to identify the optimal control policy is related to the literature on optimal control of queueing systems and make-to-stock queues. For example, Gross and Harris (1971) consider a one-server queue in which the service rate depends on the number of orders in the queue. Li (1992) studies the choice between make-to-order and make-to-stock systems. Veatch and Wein (1994) consider a two-stage tandem inventory system and investigate conditions to support different policies. To the best of our knowledge, the present work is the first to characterize the optimal policies of a two-stage tandem queue with the option of emergency orders. The approach, which involves several system transformations, appears to be new.

The rest of the paper is organized as follows: Section 3 formulates the model (System **O**) and introduces the equivalent System **M**. In Section 4, we transform the problem into a tandem-queue control problem (System **A**) and establish its optimal policy. Based on System **A**'s optimal policy, in Section 5, we construct an optimal policy for System **M** (and consequently System **O**) under a mild assumption. When the assumption does not hold, we use the optimal policy of System **A** to develop a heuristic policy for System **O** in Section 6, and compare it with several existing heuristic policies such as the dual-index policy and the tailored base-surge policy. Finally, Section 7 summarizes our findings. All proofs are contained in the Electronic Companion online.

3. Problem Formulation

Consider a single-item inventory system where the item can be ordered from two different sources, a normal source and an emergency source. See Figure 2. The sources can be represented as two servers in tandem. When an order is placed at the normal source, it first goes to server 1 to be processed and then to server 2. An emergency order skips server 1 and goes directly to server 2. After an order is processed at server 2, it is delivered to the stock, available to satisfy customer demands. Each server adopts a work-conserving rule (it is never idle as long as there are orders waiting to be processed in its queue) and any service discipline. The processing times at server i $(i = 1, 2)$ are independent and exponentially distributed with rate μ_i. Let T_i denote a generic processing time at server i $(i = 1, 2)$. The customer demands follow a Poisson process with rate λ. The system costs include ordering costs, holding and backorders costs. Specifically,

- c_1: order cost per unit for the normal source;
- h_1: holding cost per unit per unit time for units waiting or being processed at server 1;
- c_2: order cost per unit for the emergency source, where $c_2 > c_1$;
- h_2: holding cost per unit per unit time for units waiting or being processed at server 2;
- h: finished goods holding cost per unit per unit time;
- b: finished goods backorder cost per unit per unit time.

At time t, let

- $N_1(t)$: the number of orders waiting or being processed at server 1;
- $N_2(t)$: the number of orders waiting or being processed at server 2;
- $IN(t)$: the net inventory (on-hand inventory minus backorders).

The system state at time t is $(IN(t), N_2(t), N_1(t))$. As usual, the order costs from both sources are charged at the moment the order is placed. Costs are discounted at continuous rate $\alpha > 0$. We call this the original system, or System **O** for short.

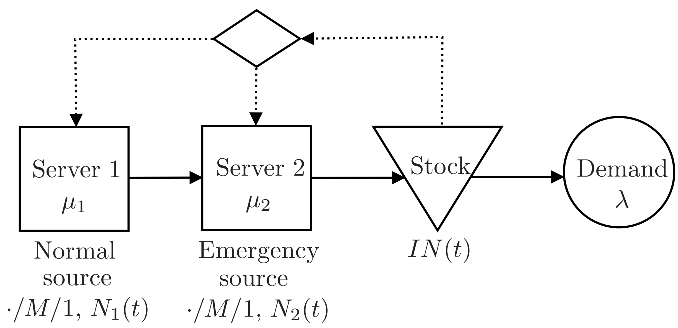

Figure 2 Open Network of Queues for Backorders

Let $\mathbb{Z}$ denote the set of integers, and $\mathbb{Z}_+$ the nonnegative integers. We uniformize the process as in Lippman (1975) by rescaling the parameters so that $\alpha+\lambda+\mu_1+\mu_2=1$. Let $f(x,y,z)$ be the optimal expected discounted cost over an infinite time horizon when the system's initial state is $(IN(0),N_2(0),N_1(0))=(x,y,z)\in\mathbb{Z}\times\mathbb{Z}_+^2$. This function satisfies the optimality equation

$$\begin{aligned} f(x,y,z) &= hx^+ + bx^- + h_2 y + h_1 z \\ &\quad + \min_{u\geq y, v\geq z}\Big\{(h_2+c_2)(u-y)+(h_1+c_1)(v-z)+\lambda f(x-1,u,v) \\ &\qquad + \mu_1[\mathbf{I}_{\{v>0\}} f(x,u+1,v-1)+\mathbf{I}_{\{v=0\}} f(x,u,v)] \\ &\qquad + \mu_2[\mathbf{I}_{\{u>0\}} f(x+1,u-1,v)+\mathbf{I}_{\{u=0\}} f(x,u,v)]\Big\} \\ &:= hx^+ + bx^- + h_2 y + h_1 z + \mathcal{T} f(x,y,z). \end{aligned} \tag{1}$$

For any function $g(x,y,z)$ defined on $\mathbb{Z}\times\mathbb{Z}_+^2$, define

$$\begin{aligned} \triangle_x g(x,y,z) &= g(x+1,y,z)-g(x,y,z), \\ \triangle_y g(x,y,z) &= g(x,y+1,z)-g(x,y,z), \\ \triangle_z g(x,y,z) &= g(x,y,z+1)-g(x,y,z). \end{aligned}$$

LEMMA 1. *The optimal expected discounted cost of System* **O**, $f(x,y,z)$, *satisfies, for* $y\geq 0$ *and* $z\geq 0$,

$$c_2+\triangle_y f(x,y,z)\geq 0, \quad c_1+\triangle_z f(x,y,z)\geq 0. \tag{2}$$

It is optimal not to place an emergency order when $y>0$ *and not to place a normal order when* $z>0$.

So, we should never place an order at a busy server. Thus, except perhaps for an initial transient period, there is at most one unit at server 1, and no order there ever waits to begin processing. Hence, the expected discounted holding cost incurred by an order at server 1 is $h_1/(\alpha+\mu_1)$. We can therefore add this constant to c_1 and set h_1 to 0 and obtain an equivalent formulation. From here on we presume this has been done.

Next we introduce another model (denoted by System **M**) to help us analyze System **O**. System **M** is identical to System **O**, except for the ordering cost at the normal source. For System **M**, that cost is $\tilde{c}_1$ (which will be specified shortly, see (5)), charged when the order completes service at server 1, not when the order is placed. For System **M**, let $\hat{f}(x,y,z)$ be the optimal expected discounted cost over an infinite horizon, when the initial state is $(x,y,z)\in\mathbb{Z}\times\mathbb{Z}_+^2$. Similar to (1), the optimality equation for System **M** is

$$\hat{f}(x,y,z) = hx^+ + bx^- + h_2 y$$

$$
\begin{aligned}
&+ \min_{u\geq y, v\geq z} \Big\{ (h_2+c_2)(u-y) + \lambda \hat{f}(x-1,u,v) \\
&\qquad + \mu_1[\mathbf{I}_{\{v>0\}}(\hat{f}(x,u+1,v-1)+\tilde{c}_1) + \mathbf{I}_{\{v=0\}}\hat{f}(x,u,v)] \\
&\qquad + \mu_2[\mathbf{I}_{\{u>0\}}\hat{f}(x+1,u-1,v) + \mathbf{I}_{\{u=0\}}\hat{f}(x,u,v)] \Big\} \\
:=\,& hx^+ + bx^- + h_2 y + \hat{\mathcal{T}}\hat{f}(x,y,z). \qquad (3)
\end{aligned}
$$

Lemma 2. *The optimal expected discounted cost of System* **M**, $\hat{f}(x,y,z)$, *satisfies, for* $y \geq 0$ *and* $z \geq 0$,

$$
c_2 + \triangle_y \hat{f}(x,y,z) \geq 0, \quad \triangle_z \hat{f}(x,y,z) \geq 0. \qquad (4)
$$

It is optimal not to place an emergency order when $y > 0$ *and not to place a normal order when* $z > 0$.

By Lemmas 1 and 2, for both Systems **O** and **M**, we need only consider $z=0$ and 1 in the optimality equations. Recalling that the processing time of server 1 is T_1, in order that these two systems have the same expected discounted ordering cost for the normal source, we set $c_1 = \mathsf{E}e^{-\alpha T_1}\tilde{c}_1$. Noting that T_1 is an exponential random variable with parameter μ_1, we have

$$
\tilde{c}_1 = \frac{\mu_1+\alpha}{\mu_1} c_1. \qquad (5)
$$

Hence, we have

$$
f(x,y,z) = \begin{cases} \hat{f}(x,y,z), & \text{if } z=0, \\ \hat{f}(x,y,z) - c_1, & \text{if } z=1. \end{cases} \qquad (6)
$$

Next we establish the equivalence of Systems **O** and **M** in terms of their optimal policies.

Lemma 3. *Systems* **O** *and* **M** *have the same optimal policy.*

Next, we introduce a simpler system that will help us find an optimal policy.

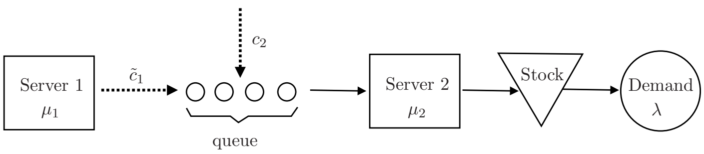

Figure 3 Auxiliary System

4. The Auxiliary System

4.1. Definition

The auxiliary system, called System **A** (see Figure 3), is similar to System **M**. However, we do not place normal orders, and there is no queue at server 1. Instead, we control the service rate at server 1. There is, in effect, an ample supply of costless units for the server to work on. Now, μ_1 means the maximal service rate. We still place emergency orders directly to server 2, which works as before.

As there is no service-rate charge at server 1, we can omit $N_1(t)$ from the state, leaving only $IN(t)$ and $N_2(t)$. Let $\tilde{f}(x,y)$ represent the optimal expected discounted cost over an infinite horizon with the initial condition $(IN(0),N_2(0))=(x,y)\in\mathbb{Z}\times\mathbb{Z}_+$. Let us apply the uniformization technique above. Note that the optimality equation for System **A** is linear in the processing rate at server 1. Therefore, the optimal service rate at server 1 is a bang-bang rule, either 0 or μ_1. The optimality equation thus can be written as

$$\begin{aligned}\tilde{f}(x,y) &= hx^+ + bx^- + h_2y + \min_{u\geq y}\Big\{(h_2+c_2)(u-y) + \lambda\tilde{f}(x-1,u)\\ &\qquad + \mu_1\min\{\tilde{f}(x,u+1)+\tilde{c}_1, \tilde{f}(x,u)\}\\ &\qquad + \mu_2[\mathbf{I}_{\{u>0\}}\tilde{f}(x+1,u-1)+\mathbf{I}_{\{u=0\}}\tilde{f}(x,u)]\Big\},\\ &:= c(x,y) + \tilde{\mathcal{T}}\tilde{f}(x,y). \end{aligned} \tag{7}$$

We remark that, compared to the system considered by Veatch and Wein (1994), System **A** is more complicated. There is one more decision, namely, when to place an emergency order which goes directly to server 2.

In the following we aim to find an optimal policy for System **A**. Specifically, in Subsection 4.2, we first identify some properties of $\tilde{f}(x,y)$ and then use these properties to characterize the structure of the optimal policy. In Subsection 4.3, we prove that the operator $\tilde{\mathcal{T}}$ defined in (7) preserves these properties. Then we conclude that $\tilde{f}(x,y)$ does have the properties used in Subsection 4.2.

4.2. Structure of Optimal Policy

For any function $g(x,y)$ defined on $\mathbb{Z}\times\mathbb{Z}_+$, let

$$\begin{aligned}\triangle_{xy}g(x,y) &= \triangle_x g(x,y+1) - \triangle_x g(x,y),\\ \triangle_{xx}g(x,y) &= \triangle_x g(x+1,y) - \triangle_x g(x,y),\\ \triangle_{yy}g(x,y) &= \triangle_y g(x,y+1) - \triangle_y g(x,y).\end{aligned}$$

It is easy to verify that $\triangle_{xy}g(x,y) = \triangle_{yx}g(x,y)$.

Suppose a function $g(x,y)$ has the following properties:

- Property (P1): Supermodularity: $\triangle_{xy}g(x,y)\geq 0$;
- Property (P2): Diagonal dominance in x: $\triangle_{xx}g(x,y)\geq\triangle_{xy}g(x,y)$;
- Property (P3): Diagonal dominance in y: $\triangle_{yy}g(x,y)\geq\triangle_{xy}g(x,y)$.

Remark 1. A two-variable function satisfying Properties (P2)-(P3) is called super-convex by Koole (2007). Moreover, according to Altman et al. (2000), such a function is multimodular. Also, let $\bar{g}(x,y)=g(x,-y)$. Then $\bar{g}(x,y)$ is $L^\natural$-convex. See Murota (2003).

Intuitively, as in System **O**, at state (x,y) with $y>0$, we will not place an emergency order at server 2, because server 2 is already busy, and deferring the decision reduces the discounted order cost without affecting other costs. We verify this intuitive observation in the following lemma.

Lemma 4. *The expected discounted optimal cost $\tilde{f}(\cdot,\cdot)$ satisfies*

$$c_2+\triangle_y\tilde{f}(x,y)\geq 0 \quad \textit{for any} \quad y\geq 0. \tag{8}$$

Thus, it is optimal not to order more at server 2 when it is busy.

Using the optimality equation (7),

$$\begin{aligned}\tilde{\mathcal{T}}\tilde{f}(x,y) &= \lambda\tilde{f}(x-1,y)+\mu_1\min\left\{\tilde{f}(x,y),\tilde{f}(x,y+1)+\tilde{c}_1\right\}\\ &\quad +\mu_2\Big(\mathbf{I}_{\{y>0\}}\tilde{f}(x+1,y-1)+\mathbf{I}_{\{y=0\}}\tilde{f}(x+1,y)\Big)\\ &= \lambda\tilde{f}(x-1,y)+\mu_1\tilde{f}(x,y)+\mu_1\min\left\{\triangle_y\tilde{f}(x,y)+\tilde{c}_1,0\right\}\\ &\quad +\mu_2\Big(\mathbf{I}_{\{y>0\}}\tilde{f}(x+1,y-1)+\mathbf{I}_{\{y=0\}}\tilde{f}(x+1,y)\Big).\end{aligned}$$

Thus to determine the optimal policy, it suffices to compare $\triangle_y\tilde{f}(x,y)+\tilde{c}_1$ to 0. Accordingly, for $y\geq 0$, define

$$S(y)=\min\{x\ :\ \triangle_y\tilde{f}(x,y)+\tilde{c}_1\geq 0\} \tag{9}$$

with the convention that $S(y)=\infty$ if the set given by the right-hand side is empty. The following result describes the optimal action for server 1.

Lemma 5. *Assume that Properties* (P1) *and* (P3) *hold for $\tilde{f}(\cdot,\cdot)$. $S(y)$ is uniquely determined. Moreover it is decreasing in y, that is, $S(y)\leq S(y-1)$. In particular, if both $S(y)$ and $S(y-1)$ are finite, then $S(y)<S(y-1)$. Thus for state $(x,y)\in\mathbb{Z}\times\mathbb{Z}_+$, when $x<S(y)$, it is optimal to keep server* 1 *on; when $x\geq S(y)$, it is optimal to keep server* 1 *off.*

Since $S(y)$ is decreasing in y, a larger x or a larger y means it less likely we want to operate server 1. This makes intuitive sense. Both x and y represent stock in the system, now or later. So, the larger these are, the less need there is to add more stock, which is the effect of operating server 1.

Now, we consider states $(x,0)$ where server 2 is idle. We need to decide whether to order directly at server 2. By the optimality equation (7),

$$\begin{aligned}
\tilde{T}\tilde{f}(x,0) = &\min_{u\in\{0,1\}} \Big\{(h_2+c_2)u + \lambda\tilde{f}(x-1,u) + \mu_1 \min\{\tilde{f}(x,u+1)+\tilde{c}_1, \tilde{f}(x,u)\} \\
&\qquad + \mu_2[\mathbf{I}_{\{u>0\}}\tilde{f}(x+1,u-1) + \mathbf{I}_{\{u=0\}}\tilde{f}(x,u)]\Big\} \\
= &\min\Big\{\lambda\tilde{f}(x-1,0)+\mu_1\tilde{f}(x,0)+\mu_2\tilde{f}(x,0), \\
&\quad \mu_1\tilde{c}_1 + \lambda\tilde{f}(x-1,0)+\mu_1\tilde{f}(x,1)+\mu_2\tilde{f}(x,0), \\
&\quad h_2+c_2+\mu_1\tilde{c}_1+\lambda\tilde{f}(x-1,1)+\mu_1\tilde{f}(x,2)+\mu_2\tilde{f}(x+1,0), \\
&\quad h_2+c_2+\lambda\tilde{f}(x-1,1)+\mu_1\tilde{f}(x,1)+\mu_2\tilde{f}(x+1,0)\Big\} \\
= &\lambda\tilde{f}(x-1,0)+\mu_1\tilde{f}(x,0)+\mu_2\tilde{f}(x,0) \\
&+\min\Big\{0,\ \mu_1\tilde{c}_1+\mu_1\triangle_y\tilde{f}(x,0),\ h_2+c_2+\mu_1\tilde{c}_1+\lambda\triangle_y\tilde{f}(x-1,0) \\
&\qquad +\mu_1(\triangle_y\tilde{f}(x,1)+\triangle_y\tilde{f}(x,0))+\mu_2\triangle_x\tilde{f}(x,0), \\
&\qquad h_2+c_2+\lambda\triangle_y\tilde{f}(x-1,0)+\mu_1\triangle_y\tilde{f}(x,0)+\mu_2\triangle_x\tilde{f}(x,0)\Big\}. \qquad (10)
\end{aligned}$$

In (10), the minimum of the four terms is just a comparison between four possible actions, namely, server 1 off and no order at server 2, server 1 on and no order at server 2, server 1 on and ordering at server 2, and server 1 off and ordering at server 2. Because $S(1)<S(0)$ by Lemma 5, we partition the stock level x into three intervals, $(-\infty,S(1))$, $[S(1),S(0))$, and $[S(0),\infty)$. Again by Lemma 5, instead of comparing four actions in (10), we need only compare two actions in each interval. Table 1 summarizes the possibilities. Each possible action is represented by a pair of phrases in parentheses; the first phrase gives the action at server 1, and the second the action at server 2.

Table 1 At state $(x,0)$, two actions comparison in each interval

$x<S(1)$	$S(1)\le x<S(0)$	$x\ge S(0)$
(on, feed) or (on, do nothing)	(off, feed) or (on, do nothing)	(off, feed) or (off, do nothing)

For example, in the first interval ($x<S(1)$), action (on, feed) is better than (off, feed), by the definition of $S(1)$; since $S(1)<S(0)$, we also have $x<S(0)$, which implies that (on, do nothing) is better than (off, do nothing), in view of the definition of $S(0)$. Therefore, on the first interval, we do not need to consider actions (off, feed) and (off, do nothing); we need only compare the two actions shown in the table, (on, feed) and (on, do nothing). Using the optimality equation (10), (on, feed) is better than (on, do nothing), if and only if

$$\begin{aligned}
\mu_1\tilde{c}_1+\mu_1\triangle_y\tilde{f}(x,0) > &\ h_2+c_2+\mu_1\tilde{c}_1+\lambda\triangle_y\tilde{f}(x-1,0)+\mu_2\triangle_x\tilde{f}(x,0) \\
&+\mu_1\Big(\triangle_y\tilde{f}(x,1)+\triangle_y\tilde{f}(x,0)\Big). \qquad (11)
\end{aligned}$$

In the second interval ($[S(1), S(0))$), similarly, we need only compare (off, feed) and (on, do nothing). Here, (off, feed) is better, if and only if

$$h_2 + c_2 + \lambda \triangle_y \tilde{f}(x-1,0) + \mu_1 \triangle_y \tilde{f}(x,0) + \mu_2 \triangle_x \tilde{f}(x,0) < \mu_1 \tilde{c}_1 + \mu_1 \triangle_y \tilde{f}(x,0). \tag{12}$$

Finally, in the third interval ($x > S(0)$), it is sufficient to compare (off, feed) and (off, do nothing), and (off, feed) is better, if and only if

$$h_2 + c_2 + \lambda \triangle_y \tilde{f}(x-1,0) + \mu_1 \triangle_y \tilde{f}(x,0) + \mu_2 \triangle_x \tilde{f}(x,0) < 0. \tag{13}$$

Inequalities (11)-(13) can be given by a uniform inequality

$$\begin{aligned} &h_2 + c_2 - \mu_1 \tilde{c}_1 + \lambda \triangle_y \tilde{f}(x-1,0) + \mu_2 \triangle_x \tilde{f}(x,0) \\ &\quad < -\mu_1 \Big(\triangle_y \tilde{f}(x,1) + \tilde{c}_1 \Big) \mathbf{I}_{\{x < S(1)\}} - \mu_1 \Big(\triangle_y \tilde{f}(x,0) + \tilde{c}_1 \Big) \mathbf{I}_{\{x \geq S(0)\}}. \end{aligned}$$

We summarize this discussion in the following lemma.

LEMMA 6. *Assume that Properties* (P1)-(P3) *hold for* $\tilde{f}(\cdot,\cdot)$. *The threshold* R *defined by*

$$\begin{aligned} R = \min \Big\{ x : &\, h_2 + c_2 - \mu_1 \tilde{c}_1 + \lambda \triangle_y \tilde{f}(x-1,0) + \mu_2 \triangle_x \tilde{f}(x,0) \\ &\geq -\mu_1 (\triangle_y \tilde{f}(x,1) + \tilde{c}_1) \mathbf{I}_{\{x < S(1)\}} - \mu_1 (\triangle_y \tilde{f}(x,0) + \tilde{c}_1) \mathbf{I}_{\{x \geq S(0)\}} \Big\} \end{aligned}$$

is uniquely determined with the convention that $R = \infty$ *if the set given by the right-hand side is empty. Moreover, for state* $(x,0)$, *if* $x < R$, *it is optimal to feed server* 2; *if* $x \geq R$, *it is optimal to do nothing at server* 2.

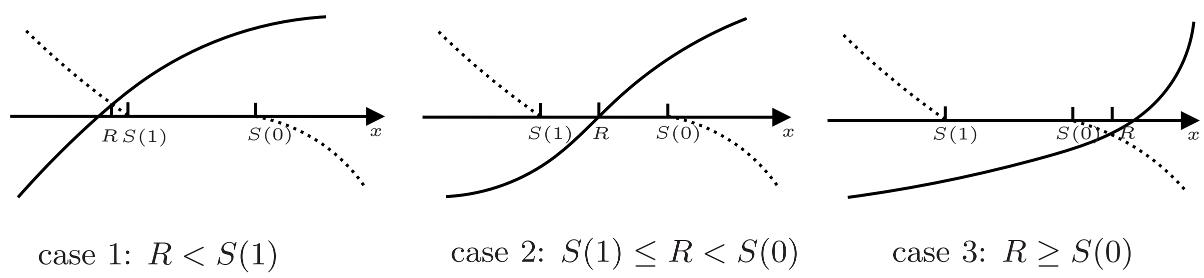

Figure 4 Different Locations of R

Now, R can lie in any of the three intervals. See Figure 4, where the dotted curve is $-\mu_1(\triangle_y \tilde{f}(x,1) + \tilde{c}_1)\mathbf{I}_{\{x<S(1)\}} - \mu_1(\triangle_y \tilde{f}(x,0) + \tilde{c}_1)\mathbf{I}_{\{x \geq S(0)\}}$, and the solid curve is $h_2 + c_2 - \mu_1 \tilde{c}_1 + \lambda \triangle_y \tilde{f}(x-1,0) + \mu_2 \triangle_x \tilde{f}(x,0))$. With the help of Lemmas 5-6, we can immediately obtain the optimal policy.

THEOREM 1. *Assume that Properties* (P1)-(P3) *hold for* $\tilde{f}(\cdot,\cdot)$. *The following actions are optimal in each of the cases identified*:

For $N_2 = 0$:

(i) $R < S(1)$ (*Figure* 5) : (a) $IN < R$: *Keep server* 1 *on, and feed a unit to server* 2; (b) $R \le IN < S(0)$: *Keep server* 1 *on, and do nothing at server* 2; (c) $IN \ge S(0)$: *Keep server* 1 *off and do nothing at server* 2.

(ii) $S(1) \le R < S(0)$ (*Figure* 6) : (a) $IN < S(1)$: *Keep server* 1 *on and feed a unit to server* 2; (b) $S(1) \le IN < R$: *Keep server* 1 *off and feed a unit to server* 2; (c) $R \le IN < S(0)$: *Keep server* 1 *on and do nothing at server* 2; (d) $IN \ge S(0)$: *Keep server* 1 *off and do nothing at server* 2.

(iii) $R \ge S(0)$ (*Figure* 7) : (a) $IN < S(1)$: *Keep server* 1 *on and feed a unit to server* 2; (b) $S(1) \le IN < R$: *Keep server* 1 *off and feed a unit to server* 2; (c) $IN \ge R$: *Keep server* 1 *off and do nothing at server* 2.

For $N_2 > 0$:

(iv) $IN < S(N_2)$: *Keep server* 1 *on and do nothing at server* 2.

(v) $IN \ge S(N_2)$: *Keep server* 1 *off and do nothing at server* 2.

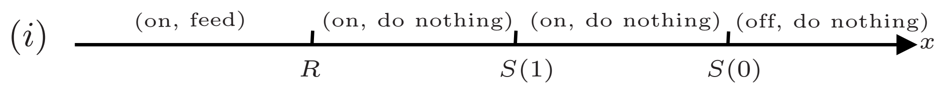

Figure 5 Optimal Policy When $R < S(1)$ and $N_2 = 0$

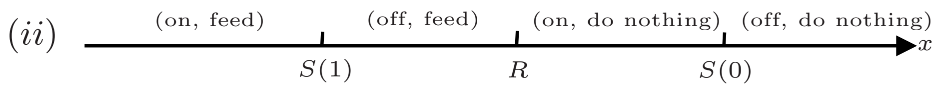

Figure 6 Optimal Policy When $S(1) \le R < S(0)$ and $N_2 = 0$

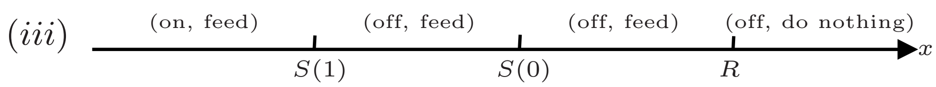

Figure 7 Optimal Policy When $R \ge S(0)$ and $N_2 = 0$

4.3. Preservation

In this subsection, we show that the optimal cost function $\tilde{f}(x,y)$ has Properties (P1)-(P3). As usual (see Weber and Stidham (1987)), we prove that the operator $\tilde{\mathcal{T}}$ defined in (7) preserves these properties. The argument is divided into several parts. Lemma 8 considers the case $y>0$. Lemma 9 is devoted to the case $y=0$. This case needs an additional property (see (16)) which is preserved under $\tilde{\mathcal{T}}$, as shown by Lemma 10. Lemma 11 shows that $bx^- + hx^+ + h_2 y$ also satisfies Properties (P1)-(P3). The argument concludes with Theorem 2.

Lemma 7. *If $\psi(x,y)$ satisfies*

$$c_2 + \triangle_y \psi(x,y) \geq 0 \tag{14}$$

for $(x,y) \in \mathbb{Z} \times \mathbb{Z}_+$, then for any positive integer y,

$$\begin{aligned}\min_{u \geq y} \Big\{ &(h_2 + c_2)(u-y) + \lambda\psi(x-1,u) + \mu_1 \min\{\psi(x,u+1) + \tilde{c}_1, \psi(x,u)\} \\ &+ \mu_2 [\mathbf{I}_{\{u>0\}} \psi(x+1,u-1) + \mathbf{I}_{\{u=0\}} \psi(x,u)] \Big\} \\ = {}& \lambda\psi(x-1,y) + \mu_1 \min\Big\{\psi(x,y+1) + \tilde{c}_1, \psi(x,y)\Big\} + \mu_2 \psi(x+1,y-1). \end{aligned} \tag{15}$$

Lemma 8. *If $\psi(x,y)$ satisfies Properties* (P1)-(P3)*, and condition* (14) *in Lemma* 7*, then $\tilde{\mathcal{T}}\psi(x,y)$ satisfies Properties* (P1)-(P3) *for $y>0$.*

Lemma 9. *If $\psi(x,y)$ satisfies Properties* (P1)-(P3)*, condition* (14) *in Lemma* 7*, and*

$$\triangle_x \psi(x,0) \leq \triangle_y \psi(x,0), \tag{16}$$

then $\tilde{\mathcal{T}}\psi(x,y)$ also satisfies Properties (P1)-(P3) *for $y=0$.*

Lemma 10. *If $\psi(x,y)$ satisfies condition* (14) *in Lemma* 7*, and condition* (16) *in Lemma* 9*, then*

$$\triangle_x \tilde{\mathcal{T}}\psi(x,y) \leq \triangle_y \tilde{\mathcal{T}}\psi(x,y).$$

For the remaining results we need to impose an assumption on the holding costs, namely, $h_2 \geq h$. We shall discuss it later on (see Subsection 5.2.1).

Lemma 11. *The function $c(x,y) = hx^+ + bx^- + h_2 y$ with $h_2 \geq h$ satisfies Properties* (P1)-(P3)*, condition* (14) *in Lemma* 7*, and condition* (16) *in Lemma* 9*.*

Now, suppose we start with $c(x,y) = hx^+ + bx^- + h_2 y$ as an estimate of $\tilde{f}(x,y)$ in the right-hand side of (7) and iterate as usual. Standard dynamic-programming arguments verify that, in the limit, we obtain $\tilde{f}(x,y)$. By Lemmas 8-11, we have the following theorem.

Theorem 2. *If $h_2 \geq h$, then $\tilde{f}(x,y)$ satisfies Properties* (P1)-(P3)*.*

4.4. Long-run Average Cost Criterion

This subsection shows that the previous results for discounted costs extend to the long-run average-cost problem. For average costs, the optimality equation becomes

$$\begin{aligned}&\bar{f}(x,y)+\gamma\\&=hx^{+}+bx^{-}+h_2y+\min_{u\ge y}\Big\{(c_2+h_2)(u-y)+\lambda\bar{f}(x-1,u)\\&+\mu_1\min\{\bar{f}(x,u+1),\bar{f}(x,u)\}+\mu_2[\mathbf{I}_{\{u>0\}}\bar{f}(x+1,u-1)+\mathbf{I}_{\{u=0\}}\bar{f}(x,u)]\Big\}.\end{aligned}$$

Here, γ is the optimal average cost rate. To characterize the optimal policy in this case, we verify certain conditions in Weber and Stidham (1987) (page 212, (a)-(g)). It is easy to check that our problem satisfies all their conditions except (c) and (d). But actually we do not need these two conditions here, due to the fact that the set of feasible actions is finite. Their results then imply that the long-run average-cost optimal policy can be obtained as the limit of discounted-cost optimal policies, as the discount rate becomes small. Hence, there exist $\bar{R}$, and a function $\bar{S}(\cdot)$ defined on $\mathbb{Z}_+$, such that the optimal policy can be described just as in Theorem 1, with R replaced by $\bar{R}$ and $S(\cdot)$ replaced by $\bar{S}(\cdot)$.

5. The Optimal Policy for the Original System

In this section, using the optimal policy of System **A** given by Theorem 1, we describe the optimal policy for the original system.

5.1. Optimal Policy for System M

To get the optimal policy for System **M**, we first establish an equivalence between System **M** and System **A** under certain conditions.

PROPOSITION 1. *If $R<S(1)$ in System* **A***, then the optimal cost functions $\tilde{f}(x,y)$ and $\hat{f}(x,y,z)$ satisfy*

$$\tilde{f}(x,y)=\hat{f}(x,y,z).$$

Moreover, the optimal policy of System **A** *is also optimal for System* **M**.

Next, we identify when the condition $R<S(1)$ in Proposition 1 holds. By the definition of $S(y)$ in (9), for $x<S(0)$,

$$-\mu_1(\triangle_y\tilde{f}(x,1)+\tilde{c}_1)\mathbf{I}_{\{x<S(1)\}}-\mu_1(\triangle_y\tilde{f}(x,0)+\tilde{c}_1)\mathbf{I}_{\{x\ge S(0)\}}\ge 0,$$

and for $x\ge S(0)$,

$$-\mu_1(\triangle_y\tilde{f}(x,1)+\tilde{c}_1)\mathbf{I}_{\{x<S(1)\}}-\mu_1(\triangle_y\tilde{f}(x,0)+\tilde{c}_1)\mathbf{I}_{\{x\ge S(0)\}}\le 0.$$

Hence, from the definition of R in Lemma 6, a sufficient condition for $R < S(1)$ is, for $x < S(1)$,

$$h_2 + c_2 - \mu_1 \tilde{c}_1 + \lambda \triangle_y \tilde{f}(x-1,0) + \mu_2 \triangle_x \tilde{f}(x,0) \geq 0.$$

Note, again by Lemma 4,

$$\begin{aligned} & h_2 + c_2 - \mu_1 \tilde{c}_1 + \lambda \triangle_y \tilde{f}(x-1,0) + \mu_2 \triangle_x \tilde{f}(x,0) \\ & = h_2 + c_2(1-\lambda) - \mu_1 \tilde{c}_1 + \lambda \Big[\triangle_y \tilde{f}(x-1,0) + c_2 \Big] + \mu_2 \triangle_x \tilde{f}(x,0) \\ & \geq h_2 + c_2(1-\lambda) - \mu_1 \tilde{c}_1 + \mu_2 \triangle_x \tilde{f}(x,0). \end{aligned}$$

Hence, if the optimal cost $\tilde{f}(x,y)$ satisfies

$$h_2 + c_2(1-\lambda) - \mu_1 \tilde{c}_1 + \mu_2 \triangle_x \tilde{f}(x,0) \geq 0, \tag{17}$$

then $R < S(1)$. To establish (17), we need the following assumption on the system primitives.

- **Assumption A:** $h_2 + c_2(1-\lambda) - \mu_1 \tilde{c}_1 \geq \mu_2 b/\alpha$.

The next proposition states that Assumption A is a sufficient condition for (17).

PROPOSITION 2. *If Assumption A holds, then, for any* $(x,y) \in \mathbb{Z} \times \mathbb{Z}_+$,

$$h_2 + (1-\lambda) c_2 - \mu_1 \tilde{c}_1 + \mu_2 \triangle_x \tilde{f}(x,y) \geq 0. \tag{18}$$

Now we discuss Assumption A. This is essentially a sufficient condition for ensuring that, for system **A**, if the entire system is idle, then it is better to turn server 1 on than to turn it off, and it is better not to feed a unit to server 2 than to feed a unit to it. Simply put, (on, do nothing) is better than (off, feed). To see this, at state (x,y), if (off, feed) is used at the beginning, the cost is

$$hx^+ + bx^- + c_2 + h_2(y+1) + \lambda \tilde{f}(x-1,y+1) + \mu_1 \tilde{f}(x,y+1) + \mu_2 \tilde{f}(x+1,y);$$

while if decision (on, do nothing) is used, the cost is

$$hx^+ + bx^- + h_2 y + \lambda \tilde{f}(x-1,y) + \mu_1(\tilde{c}_1 + \tilde{f}(x,y+1)) + \mu_2 \Big(\mathbf{I}_{\{y=0\}} \tilde{f}(x,y) + \mathbf{I}_{\{y>0\}} \tilde{f}(x+1,y-1) \Big).$$

The difference between these two costs is

$$c_2 + h_2 - \mu_1 \tilde{c}_1 + \lambda \Delta_y \tilde{f}(x-1,y) + \mu_2 \Big(\tilde{f}(x+1,y) - \mathbf{I}_{\{y=0\}} \tilde{f}(x,y) - \mathbf{I}_{\{y>0\}} \tilde{f}(x+1,y-1) \Big). \tag{19}$$

When $y = 0$, the above difference can be written as

$$\lambda(c_2 + \Delta_y \tilde{f}(x-1,0)) + (1-\lambda) c_2 + h_2 - \mu_1 \tilde{c}_1 + \mu_2 \Delta_x \tilde{f}(x,0).$$

Using Lemma 4, the difference is larger than

$$h_2 + (1-\lambda) c_2 - \mu_1 \tilde{c}_1 + \mu_2 \triangle_x \tilde{f}(x,y).$$

By Proposition 2, this quantity is nonnegative. Thus, decision (off, feed) is worse than (on, do nothing) at state $(x,0)$.

When $y>0$, the difference (19) becomes

$$c_2(\alpha+\mu_1)-\tilde{c}_1\mu_1+h_2+\lambda(c_2+\Delta_y\tilde{f}(x-1,y))+\mu_2(c_2+\Delta_y\tilde{f}(x+1,y-1)).$$

Again by Lemma 4, the difference is larger than

$$c_2(\alpha+\mu_1)-\tilde{c}_1\mu_1+h_2,$$

which, by equation (5), is nonnegative. Therefore, again, decision (on, do nothing) is better than decision (off, feed) at state (x,y) with $y>0$.

5.2. Optimal Policy for System O

In this subsection, we establish the optimal policy for the original system (System **O**). At the end, we consider the special case $c_1=c_2$.

We summarize Lemma 3, Theorem 1, and Propositions 1-2 to obtain the optimal policy for System **O**.

Theorem 3. *Assume that $h_2\geq h$ and Assumption A hold. For the original problem (System* **O***), the following actions are optimal in each of the cases identified.*

For $N_2=0, N_1=0$:

(i) $IN<R$: *place one order at each source*;

(ii) $R\leq IN<S(0)$: *place one order at the normal source, and do nothing at the emergency source*;

(iii) $IN\geq S(0)$: *do nothing at both sources.*

For $N_2>0, N_1=0$:

(iv) $IN<S(N_2)$: *place one order at the normal source, and do nothing at the emergency source*;

(v) $IN\geq S(N_2)$: *do nothing at both sources.*

For $N_2=0, N_1>0$:

(vi) $IN<R$: *do nothing at the normal source, and place one order at the emergency source*;

(vii) $IN\geq R$: *do nothing at both sources.*

For $N_2>0, N_1>0$:

(viii) *do nothing at both sources.*

Remark 2. Similar to Subsection 4.4, by Lemma 3 and Proposition 1, if $\bar{R}<\bar{S}(1)$, then the long-run average-cost optimal policy for System **O** can be characterized by Theorem 3 with R replaced by $\bar{R}$, and $S(\cdot)$ replaced by $\bar{S}(\cdot)$.

Now we discuss the assumption $h_2 \geq h$. Recall that, in a discounted-cost model, the holding costs include direct physical handling costs only, not financing costs. The discounting process itself captures the time-value of money. So, in our model, there is no reason to expect universal relations between h_1, h_2, and h. The relations depend on the particular physical characteristics of the items in the system. For example, the finished good could be more delicate than the intermediate products and so require more expensive handling; or, just as likely, the intermediate products could be more delicate. Therefore, the assumption $h_2 \geq h$ does rule out some practical cases, but not all. In an average-cost model, the holding costs do include financing costs, so $h_2 \geq h$ is less likely to hold. Veatch and Wein (1992) and Weber and Stidham (1987) discuss similar assumptions in different contexts. Section 6 presents a heuristic for the case with $h_2 < h$.

In our analysis, the assumption $h_2 \geq h$ is used to guarantee that the optimal-cost function $\tilde{f}(x,y)$ for System **A** satisfies Property (P3); see Lemma 11. Table 2 presents some counterexamples (with $h_2 = 0$ and $S(y)$ given by (9)) showing that, if $h_2 < h$, then indeed Property (P3) may fail.

Table 2 Examples: $\triangle_{xy}\tilde{f}(S(y_0)-1,y_0) > \triangle_{yy}\tilde{f}(S(y_0)-1,y_0)$

λ	μ_1	μ_2	α	h	b	c_1	c_2	y_0	$S(y_0)$	$\triangle_{xy}\tilde{f}(S(y_0)-1,y_0)$	$\triangle_{yy}\tilde{f}(S(y_0)-1,y_0)$
0.3	0.3	0.39	0.01	2	90	10	30	2	10	2.937	2.008
0.3	0.3	0.39	0.01	2	120	10	30	3	10	2.086	1.811
0.3	0.34	0.35	0.01	2	90	10	30	2	13	2.316	1.629
0.3	0.34	0.35	0.01	2	120	10	30	2	14	2.287	1.725

5.2.1. Special Case $c_1 = c_2$ When $c_1 = c_2$, the normal source loses its cost advantage, so we never use it. The problem degenerates into a single-source problem involving the emergency -source only. It is easy to formulate this model and to show that the optimal policy is a simple threshold policy. It is instructive to compare this system's threshold with that of the system above. (This sort of dynamic-programming methodology appears in Heyman (1968), although his model differs from this one in some detail. By the time of Li (1992), the approach is treated as standard. For completeness, we include a brief analysis.)

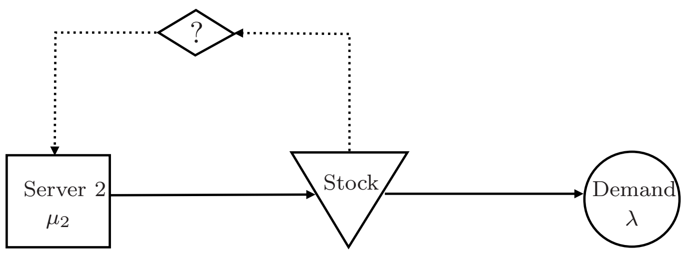

Figure 8 One Source with Backorders

We uniformize the process by rescaling the parameters, so that $\alpha+\lambda+\mu_2=1$. Given initial state $(IN(0),N_2(0))=(x,y)\in\mathbb{Z}\times\mathbb{Z}_+$, let $f^{(1)}(x,y)$ be the optimal expected discounted cost. This function satisfies the optimality equation

$$\begin{aligned} f^{(1)}(x,y) &= hx^+ + bx^- + h_2 y + \min_{u\geq y}\Big\{(c_2+h_2)(u-y)+\lambda f^{(1)}(x-1,u) \\ &\qquad + \mu_2[\mathbf{I}_{\{u>0\}}f^{(1)}(x+1,u-1)+\mathbf{I}_{\{u=0\}}f^{(1)}(x,0)]\Big\} \\ &:= hx^+ + bx^- + \mathcal{T}^{(1)}f^{(1)}(x,y). \end{aligned} \tag{20}$$

LEMMA 12. *If $\psi(x,y)$ satisfies*

$$c_2+\triangle_y\psi(x,y)\geq 0 \quad \textit{for any} \quad x\geq 0 \textit{ and } y\geq 0. \tag{21}$$

and Properties (P1)-(P2)*, then $\mathcal{T}^{(1)}\psi(x,y)$ also satisfies these conditions.*

THEOREM 4. *The expected discounted optimal cost function $f^{(1)}(\cdot,\cdot)$ satisfies* (21) *and Properties* (P1)-(P2)*. Let*

$$S=\min\{x\in\mathbb{Z}\ :\ c_2+h_2+\lambda\triangle_y f^{(1)}(x-1,0)+\mu_2\triangle_x f^{(1)}(x,0)\geq 0\}. \tag{22}$$

For $y>0$, it is optimal not to order from server 2; *For $y=0$, when $x<S$, it is optimal to place an order to server* 2; *when $x\geq S$, it is optimal not to order.*

REMARK 3. For this system, the optimal policy does not quite have a base-stock form, if we understand that to mean $IN(t)+N_2(t)=S$. The reason, again, is that in the discounted-cost setting, we can delay orders until the server becomes available. On the other hand, it is true that the optimal policy is described by one critical number, S. When the source is idle ($N_2(t)=0$) and the net inventory drops below S ($IN(t)<S$), order one unit; otherwise, do not order. By the way, Theorem 4 can be extended to the single-source lost-sales model in the same way. The optimal policy for that system is similar to the one given here.

REMARK 4. To compare this result with that of the two-source system, note that, for the original system with $c_1=c_2$, the test quantity in Lemma 6 becomes

$$\begin{aligned} & h_2+c_2-\mu_1\tilde{c}_1+\lambda\triangle_y\tilde{f}(x-1,0)+\mu_2\triangle_x\tilde{f}(x,0) \\ & +\mu_1(\triangle_y\tilde{f}(x,1)+\tilde{c}_1)\mathbf{I}_{\{x<S(1)\}}+\mu_1(\triangle_y\tilde{f}(x,0)+\tilde{c}_1)\mathbf{I}_{\{x\geq S(0)\}} \\ & \quad = (h_2+c_2)(\alpha+\lambda+\mu_1+\mu_2)-\mu_1(1+\frac{\alpha}{\mu_1})c_1+\lambda\triangle_y\tilde{f}(x-1,0)+\mu_2\triangle_x\tilde{f}(x,0) \\ & \quad\quad +\mu_1(\triangle_y\tilde{f}(x,1)+(1+\frac{\alpha}{\mu_1})c_1)\mathbf{I}_{\{x<S(1)\}}+\mu_1(\triangle_y\tilde{f}(x,0)+(1+\frac{\alpha}{\mu_1})c_1)\mathbf{I}_{\{x\geq S(0)\}}. \end{aligned}$$

Let $Q(\mu_1, x)$ denote the above expression. Then

$$\frac{\mathsf{d}Q(\mu_1,x)}{\mathsf{d}\mu_1} = \begin{cases} c_2 + h_2 + \triangle_y \tilde{f}(x,1), & \text{if } x < S(1), \\ c_2 + h_2 - c_1, & \text{if } S(1) \le x < S(0), \\ c_2 + h_2 + \triangle_y \tilde{f}(x,0), & \text{if } x \ge S(0). \end{cases}$$

In view of Lemma 4, $Q(\mu_1, x)$ is increasing in μ_1. Therefore, by Lemma 6, R is decreasing in μ_1. On the other hand, by Lemma 6 and Theorem 4, when $\mu_1 = 0$, we have $R = S$. Hence, for general $\mu_1 \ge 0$, we have $R \le S$. We summarize this discussion as follows.

COROLLARY 1. *The optimal threshold R for the dual-source problem is no more than that of the single-source problem, S.*

This result makes intuitive sense. Imagine that we start with the single-source system and then add another source, server 1. The condition to trigger orders at server 2 then becomes more stringent. For given values of IN, we use server 2 less than before.

6. Heuristic Policies and Comparisons

In this section, inspired by the analysis of Section 5, we first develop a heuristic policy for the case where Assumption A does not hold, but still $h_2 \ge h$. We compare this policy with two existing policies. Then, we construct a heuristic policy for the case where $h_2 < h$, and numerically test it.

6.1. Heuristic Policies and Numerical Experiments

Theorem 3 completely characterizes the optimal policy for System **O** when $h_2 \ge h$ and Assumption A holds. Also, for all cases, Theorem 1 describes the optimal policy for System **A**. By Proposition 1, it is only when $R \ge S(1)$ that the optimal policy of System **A** cannot be implemented in System **O**. (See Figures 6-7.) But there is only one scenario that prevents this translation. This happens when both servers are working and the net stock level IN happens to be $(S(1) - 1)$. If server 2 finishes its job first, by Theorem 1, System **A** turns server 1 off and directly places an order at server 2. But System **O** cannot turn server 1 off; once an order is placed there, it cannot be called back.

In view of this observation, suppose we modify this part of System **A**'s policy to obtain a feasible policy for System **O**. When $IN < S(1)$ and $N_2 = 0$, order from server 2, and place an order at server 1 also, if it becomes idle; when IN just reaches $S(1)$ from below and $N_2 = 0$, place an order at server 2, and do nothing at server 1. (Actually, in this case, server 1 is busy processing an order already placed before.) When $S(1) < IN < R$ and $N_2 = 0$, place an order at server 2 and do nothing at server 1. We call this the *Threshold-Curve* (or *TC*) Policy.

We now report a numerical test of this heuristic policy. Define the optimality gap as

$$\mathsf{gap}(IN, N_2, N_1) = \frac{\breve{f}(IN, N_2, N_1) - f(IN, N_2, N_1)}{f(IN, N_2, N_1)}, \tag{23}$$

where $f(IN, N_2, N_1)$ is the optimal cost of System **O**, and $\breve{f}(IN, N_2, N_1)$ is the expected discounted cost of System **O** using policy TC.

Table 3 shows the maximal gap over all states (IN, N_2, N_1) for several cases. In all these cases, $\alpha = 0.01, h = 2, h_2 = 3, c_1 = 10$ and $c_2 = 30$. Evidently, the gap is very small in all instances, so the policy performs well.

Table 3 Maximal Gaps of Heuristic Policy TC

λ	μ_1	μ_2	b	$\max \mathsf{gap}$
0.3	0.3	0.39	90	0.05 %
0.3	0.3	0.39	120	0.05 %
0.3	0.34	0.35	90	0.05 %
0.3	0.34	0.35	120	0.04 %
0.2	0.4	0.39	90	0.07 %
0.2	0.4	0.39	120	0.05 %

6.2. Dual-Index Policy

Song and Zipkin (2009) study a heuristic we call the *Dual-Index* (or *DI*) Policy. The name indicates two inventory positions,

$$IP_1(t) = IN(t) + N_2(t) + N_1(t), \tag{24}$$

$$IP_2(t) = IN(t) + N_2(t). \tag{25}$$

The policy is specified by two parameters, s_1 and s_2, with $s_1 \geq s_2$. The policy triggers orders so as to maintain $IP_1(t) = s_1$ and $IP_2(t) \geq s_2$.

We can immediately see that, under the optimal policy, $IP_1(t)$ is not always constant. No order is placed as long as both nodes are busy, even following several demands. Similarly, $IP_2(t)$ is not kept above a constant. So, the *DI* policy leads to quite different system dynamics.

An attractive feature of the *DI* policy is that it is easy to evaluate. Under such a policy, the system is equivalent to a tandem queue, where the first node has a capacity limit $u_1 = s_1 - s_2$. This system, moreover, has a product-form stationary distribution. Specifically,

$$\mathsf{Pr}(N_1 = n_1, N_2 = n_2) = \frac{(1-\rho_1)(1-\rho_2)}{1-\rho_1^{u_1+1}} \rho_1^{n_1} \rho_2^{n_2},$$

where $\rho_1 = \lambda/\mu_1$, $\rho_2 = \lambda/\mu_2$, $n_1 \in \{0, 1, \ldots, u_1\}$, and $n_2 \in \{0, 1, \ldots\}$. Define $N = N_1 + N_2$. N has stationary distribution

$$\mathsf{Pr}(N = n) = \begin{cases} \frac{(1-\rho_1)(1-\rho_2)}{(\rho_1-\rho_2)(1-\rho_1^{u_1+1})}(\rho_1^{n+1} - \rho_2^{n+1}), & \text{for } n \leq u_1, \\ \frac{(1-\rho_1)(1-\rho_2)}{(\rho_1-\rho_2)(1-\rho_1^{u_1+1})}(\rho_1^{u_1+1} - \rho_2^{u_1+1})\rho_2^{n-u_1}, & \text{for } n > u_1. \end{cases}$$

The utilization of the emergency source $\eta(u_1)$ is

$$\eta(u_1) = \mathsf{Pr}(N_1 = u_1) = \frac{1-\rho_1}{1-\rho_1^{u_1+1}}\rho_1^{u_1}.$$

With these formulas it is not hard to compute the average cost:

$$f_{DI}(s_1, u_1) = c_1\lambda + (c_2 - c_1)\lambda\eta(u_1) + h\mathsf{E}(s_1 - N)^+ + b\mathsf{E}(N - s_1)^+ + h_2\mathsf{E}N_2. \tag{26}$$

We can find the best values of (s_1, u_1) by complete search or some other means.

Inspired by the analysis in Sections 3 and 5, we propose the following modified dual-index policy, called the *Smart-Dual-Index* (or *SDI*) policy:

- when $IP_2(t) < s_2$ and server 2 is idle, order one unit at server 2,
- when $IP_1(t) < s_1$ and server 1 is idle, order one unit at server 1,

where $IP_1(t)$ and $IP_2(t)$ are given by (24)-(25). The idea here is to follow the (optimal) rule never to order from a busy server, but otherwise to follow the *DI* policy.

Table 4 Optimal Policy vs. Heuristics

$c_1 = 10$, $c_2 = 30$, $\lambda = 1$, $\mu_1 = 1.3$, and $\mu_2 = 1.5$

h	b	h_2	Optimal Cost	*DI* Policy				*SDI* Policy				*TC* Policy	
				Cost	gap%	s_1	s_2	Cost	gap %	s_1	s_2	Cost	gap%
3	60	4	41.20	43.70	6.07	9	4	42.04	2.04	9	5	41.20	0
2	60	3	33.91	35.53	4.78	11	5	34.40	1.45	10	6	33.91	0
1	60	2	25.00	25.92	3.68	13	6	25.38	1.52	12	7	25.00	0
3	90	4	44.43	46.57	4.82	10	5	44.93	1.13	10	6	44.43	0
2	90	3	35.85	37.49	4.57	12	6	36.35	1.39	11	7	35.85	0
1	90	2	25.99	26.91	3.54	14	7	26.36	1.42	13	8	25.99	0
3	120	4	46.44	48.59	4.63	11	6	46.95	1.10	10	7	46.44	0
2	120	3	37.20	38.88	4.52	12	7	37.81	1.64	12	8	37.20	0
1	120	2	26.67	27.63	3.60	15	7	27.06	1.46	14	8	26.67	0

Table 5 Optimal Policy vs. Heuristics

$c_1 = 10$, $c_2 = 30$, $\lambda = 6$, $\mu_1 = 8$, and $\mu_2 = 7$

h	b	h_2	Optimal Cost	*DI* Policy				*SDI* Policy				*TC* Policy	
				Cost	gap %	s_1	s_2	Cost	gap %	s_1	s_2	Cost	gap %
3	60	4	137.56	146.27	6.33	23	10	143.70	4.46	22	15	137.60	0.03
2	60	3	119.20	124.77	4.67	26	11	123.24	3.39	24	16	119.24	0.03
1	60	2	96.51	99.89	3.50	30	12	99.34	2.93	29	18	96.53	0.02
3	90	4	145.28	153.86	5.91	25	12	151.30	4.14	24	17	145.33	0.03
2	90	3	124.34	129.86	4.44	28	13	128.43	3.29	27	19	124.37	0.02
1	90	2	99.71	102.49	2.79	33	15	102.03	2.33	32	21	99.72	0.01
3	120	4	150.77	159.28	5.64	27	14	156.81	4.01	26	19	150.79	0.01
2	120	3	127.82	133.52	4.46	30	15	132.19	3.42	29	20	127.85	0.02
1	120	2	101.76	104.34	2.54	35	16	103.95	2.15	33	23	101.79	0.03

We now compare *DI*, *SDI* and *TC* heuristics to the optimal policy under average cost. To compute the optimal policy, we use a direct adaption of the policy iteration algorithm described in Puterman (1994) (Chapter 8 Section 8.7). The state space is truncated to $(IN, N_2, N_1) \in [-L_1, L_1] \times [0, L_2] \times \{0, 1\}$, where L_1, L_2 are positive integers. We increase L_1, L_2 until the policy doesn't change. Note that, by the structure of the optimal policy, we can restrict the feasible values of N_1 to 0 and 1. For the heuristics (or any other fixed policy), the system becomes a Markov chain. Thus, its long-run average cost can be easily obtained by solving for the stationary distribution.

Tables 4-5 show the results for several parameter settings. Evidently, policy *TC* is nearly optimal in every case. Policy *DI* performs well, and policy *SDI* is better than *DI*, but *TC* is much better than *SDI*. The intuition here is, policy *TC* dynamically adjusts the threshold for triggering orders to the normal source by $S(N_2)$ based on the congestion (N_2) at the emergency source (see Theorem 3), while *DI* and *SDI* basically fix this threshold to s_1. This dynamic adjustment significantly improves the system performance. To see this dynamic adjustment, in Tables 6-7 we provide the optimal and *TC* policy parameters for the systems described in Tables 4-5. To save space, we report the switching curves $S(N_2)$ from $N_2 = 0$ to $N_2 = 7$. In Table 6, the optimal policy and *TC* are identical; Table 7 indicates a few cases where these two policies are not exactly the same.

Table 6 Optimal/*TC* Policy Parameters in Table 4

$c_1 = 10$, $c_2 = 30$, $\lambda = 1$, $\mu_1 = 1.3$, and $\mu_2 = 1.5$

h	b	h_2	R	$S(0)$	$S(1)$	$S(2)$	$S(3)$	$S(4)$	$S(5)$	$S(6)$	$S(7)$
3	60	4	5	10	8	7	5	3	-2	-18	-18
2	60	3	5	11	10	8	6	4	2	-15	-17
1	60	2	6	14	12	10	9	7	5	2	-13
3	90	4	5	11	9	8	6	4	0	-17	-18
2	90	3	6	12	11	9	7	5	3	-15	-17
1	90	2	7	15	13	11	10	8	6	3	-11
3	120	4	6	12	10	9	7	5	1	-17	-17
2	120	3	7	13	11	10	8	6	4	-14	-17
1	120	2	8	16	14	12	10	8	6	4	-5

Table 7 Optimal/*TC* Policy Parameters in Table 5

$c_1 = 10$, $c_2 = 30$, $\lambda = 6$, $\mu_1 = 8$, and $\mu_2 = 7$

h	b	h_2	R	$S(0)$	$S(1)$	$S(2)$	$S(3)$	$S(4)$	$S(5)$	$S(6)$	$S(7)$
3	60	4	13	26/25	24/23	22	20	18	16/15	14/13	9
2	60	3	15	29/30	27/28	25/26	23/24	21	19	17	14
1	60	2	16	35	32	30	27/28	25	23	20/21	18
3	90	4	15	28/27	26	24	22	20	18	15	12
2	90	3	17	31	29	27	25	23	21	18	15/16
1	90	2	18/19	37/38	35	32/33	30	28	25/26	23	20/21
3	120	4	18/17	29	27/28	26	24	22	20/19	18/17	15/14
2	120	3	19	34	32	30	28	25/26	23	21	18
1	120	2	20	39	36/37	34	32	29/30	27	25	22

6.3. Tailored Base-Surge Policy

As reviewed in the introduction, Allon and Van Mieghem (2010), and Janakiraman et al. (2015) develop the *TBS* policy, in which a constant order quantity is placed at the normal source and a base-stock policy governs the orders from the emergency source. We adapt this notion to our setting as follows:

- Before the system starts, we select a service rate for Server 1, constrained to be no more than μ_1. Once the system starts, Server 1 always works at the chosen rate;
- If $N_2(t) > 0$, server 2 is working, and we don't order from server 2; if $N_2(t) = 0$ and $IN(t) < s$, we order one unit at server 2, and server 2 starts working; if $N_2(t) = 0$ and $IN(t) \geq s$, we don't order at server 2, and server 2 stays idle; (It's possible that $IN(t) \geq s$, because server 1 is always working.)

The policy variable s is called the base-stock level. Since server 1 is always working, the system state at any time t is $(IN(t), N_2(t))$. Similar to (1), we can write the equation for the expected discounted cost of the *TBS* policy. Here, we focus on average cost. For any fixed s, the system becomes a Markov chain. We compute its stationary distribution and from that the average cost. We numerically search for the best s given the service rate at Server 1. We then select the best such service rate, denoted μ_1^*.

In Tables 8 and 9, we set $h = 1$, $b = 50$, $h_2 = 1.5$, $c_1 = 10$ and $c_2 = 30$. In Table 8, we fix $\mu_1 = 1$ for the optimal policy and the *TC* policy. For the *TBS* policy we optimally select μ_1^* over the interval $[0, \mu_1]$. Similarly, in Table 9, we fix the arrival rate $\lambda = 1$ and the service rate at server 2 $\mu_2 = 1.5$, but consider a range of service rates μ_1 at server 1 from 0.5 to 0.9. Again, for the *TBS* policy, we optimize over the interval $[0, \mu_1]$. Evidently, the best *TBS* policy works well sometimes but not always. In Table 8, its performance improves as the arrival rate increases. This is consistent with the observations of Allon and Van Mieghem (2010). There are two reasons: (i) when the arrival rate is low, server 1 is not fully utilized, even with its optimal service rate; (ii) as discussed above, the TC policy's dynamic adjustment of the threshold is more effective than other means of controlling server 1. In Table 9, the performance of the *TBS* policy improves as the lead time difference between normal source and emergency source becomes larger. This is consistent with the findings of Janakiraman et al. (2015) and Xin and Goldberg (2015) for deterministic lead times.

Table 8 Optimal Policy vs. Heuristics

λ	μ_1	μ_2	Optimal Cost	*TC* Policy		*TBS* Policy			
				Cost	gap %	Cost	gap %	μ_1^*	s
1.1	1	7	19.19	19.19	0	20.81	8.44	0.87	1
1.2	1	7	21.34	21.34	0	22.30	4.50	0.96	1
1.3	1	7	23.76	23.76	0	24.00	1.01	1	1
1.4	1	7	26.40	26.40	0	26.47	0.27	1	1
1.5	1	7	29.25	29.25	0	29.26	0.03	1	1

Table 9 Optimal Policy vs. Heuristics

λ	μ_1	μ_2	Optimal Cost	*TC* Policy		*TBS* Policy			
				Cost	gap %	Cost	gap %	μ_1^*	s
1	0.9	1.5	26.09	26.09	0	27.56	5.63	0.81	7
1	0.8	1.5	27.09	27.09	0	27.57	1.77	0.8	7
1	0.7	1.5	28.29	28.29	0	28.33	0.14	0.7	8
1	0.6	1.5	29.77	29.77	0	29.77	0	0.6	8
1	0.5	1.5	31.43	31.43	0	31.43	0	0.5	9

6.4. Robustness of the TC Policy

In this section we demonstrate by numerical experiments that the *TC* policy still works well for deterministic processing times, compared to the *DI*, *SDI* and *TBS* policies. Now, the sources are described by two $\cdot/D/1$ queues in series (see Figure 2). Server 2 has constant service time L_2, and Server 1 has service time $L_1 = nL_2$, where n is a nonnegative integer.

In Table 10, we consider a small value of n ($n = 2$), and $L_1 = 2$, $L_2 = 1$, $h = 1$, $b = 30$, $h_2 = 1.5$, $c_1 = 1$, and $c_2 = 30$. The *TBS* policy is allowed to choose any greater value of n; the best value is denoted n^*, and the resulting L_1 is $L_1^* = n^*L_2$. The table shows that the *TC* policy performs better than the *TBS* policy; however, when λ is large, the two policies perform almost the same. This is consistent with the observations above for the case of exponential lead times.

Table 10 *TC* Policy vs. *TBS* Policy

λ	*TC* Policy				*TBS* Policy		
	Cost	$S(0)$	$S(1)$	R	Cost	L_1^*	s
0.9	30.35	25	20	17	30.58	2	17
0.8	18.92	20	15	8	18.98	2	7
0.7	13.19	15	10	4	13.25	2	4
0.6	9.39	10	6	3	9.75	2	2

Table 11 *TC* Policy vs. *TBS* Policy

λ	*TC* Policy				*TBS* Policy		
	Cost	$S(0)$	$S(1)$	R	Cost	L_1^*	s
0.8	30.22	39	16	8	30.31	10	8
0.6	19.83	20	9	4	19.85	10	4
0.4	12.09	16	7	2	12.10	10	2
0.2	5.19	15	7	1	5.20	10	1

In Table 11, we consider a large value of n, ($n = 10$), with $L_1 = 10$, $L_2 = 1$, $h = 1$, $b = 30$, $h_2 = 1.5$, $c_1 = 1$, and $c_2 = 30$. Here, the *TC* policy performs almost the same as the *TBS* policy. This is mainly due to the long lead time $L_1 = 10$ at server 1. Again, this is consistent with the observations for exponential service times.

Next, we compare the *TC* policy to the *DI* and *SDI* policies. To make those policies stable, we require $\lambda < 1/L_1$. Using the same data as in Table 10, Table 12 demonstrates that the *TC* policy performs better than the *DI* policy, but the difference is small for small λ. The *TC* policy and the *SDI* policy perform almost the same. The reason is, when λ is small, the two servers are idle most of the time. Thus, the probability that either $IP_2(t) < s_2$ and Server 2 is idle or $IP_1(t) < s_1$ and Server 1 is idle becomes large. The three policies thus choose nearly the same actions most of the time.

Table 12 *TC* Policy vs. *DI* Policy and *SDI* Policy

λ	*TC* Policy				*DI* Policy			*SDI* Policy		
	Cost	$S(0)$	$S(1)$	R	Cost	s_1	s_2	Cost	s_1	s_2
0.4	5.01	4	3	1	6.00	5	1	5.02	4	1
0.3	3.88	3	2	0	4.40	4	1	3.88	3	0
0.2	2.86	2	1	0	3.18	2	0	2.86	2	0
0.1	2.04	1	0	0	2.06	1	0	2.04	1	0

Table 13 *TC* Policy vs. *DI* Policy and *SDI* Policy

λ	*TC* Policy				*DI* Policy			*SDI* Policy		
	Cost	$S(0)$	$S(1)$	R	Cost	s_1	s_2	Cost	s_1	s_2
0.09	2.58	3	1	0	2.97	3	1	2.56	3	0
0.08	2.30	2	1	0	2.64	2	0	2.30	2	0
0.07	2.25	1	1	0	2.33	2	0	2.48	1	0
0.06	2.01	1	1	0	2.10	1	0	2.11	1	0

Finally, using the same data settings as in Table 11, Table 13 shows that the *TC* policy still performs better than the *DI* policy, although the difference is smaller than in Table 12. The *TC* policy and the *SDI* policy still perform almost the same.

6.5. Heuristic Policy for $h_2 < h$

We now propose a heuristic for systems with $h_2 < h$. The idea is simple. Construct a new system, identical to the original, but with h_2 raised to h. Find an optimal policy for the new system as above, and use it as a heuristic for the original system. We now test this approach.

We compare this heuristic policy with the *DI* and *SDI* policies in Table 14. To find the optimal policy, we use a direct adaption of the policy-iteration algorithm used in Subsection 6.2. We fix $\lambda = 0.2, h_2 = 0, c_1 = 0, c_2 = 30$, and choose h and b reflecting various service levels. In the first part of the table (Columns 1-5), we set $\mu_1 = 0.3, \mu_2 = 0.4$. In the second part (Columns 6-10), we set $\mu_1 = 0.35, \mu_2 = 0.35$.

Evidently, the heuristic policy performs well. The maximal gap is below 3.2%. Notice that the heuristic performs better when the service level is high. Here is a possible explanation: A system with a high service level holds substantial finished-goods inventory. It is likely that such a system rarely uses the emergency source. Thus, the impact of the holding cost at the emergency source is small. *SDI* performs very well, which implies that the property of not ordering from a busy server is valuable.

Table 14 Gaps of Heuristic, *DI*, and *SDI* Policies

b	h	gap %			b	h	gap %		
		Heuristic	*DI*	*SDI*			Heuristic	*DI*	*SDI*
6.40	1.60	1.27	2.27	0.09	7.39	1.85	3.16	1.80	0.19
6.18	0.69	0.80	2.30	0.06	7.08	0.79	1.72	2.17	0.22
6.08	0.32	0.68	2.46	0.48	6.96	0.37	1.04	1.38	0.48

We compare this heuristic policy with the *TBS* policy in Table 15, where we fix $\lambda = 0.4, h_2 = 0, c_1 = 0, c_2 = 30$, and choose h and b to reflect various service levels. In the first part of the table (Columns 1-4), we set $\mu_1 = 0.3, \mu_2 = 0.6$. In the second part (Columns 5-8), we set $\mu_1 = 0.3, \mu_2 = 1.2$. The heuristic policy performs well, better than the *TBS* policy. The maximal gap is below 0.3%. Both heuristics perform better when the service level is high. The explanation above also applies here.

Table 15 Gaps of Heuristic and *TBS* Policies

b	h	gap %		b	h	gap %	
		Heuristic	*TBS*			Heuristic	*TBS*
6.40	1.60	0.22	6.22	7.39	1.85	0.30	13.33
6.18	0.69	0.05	1.80	7.08	0.79	0	5.37
6.08	0.32	0	0.20	6.96	0.37	0	0.68

7. Concluding Remarks

In this paper, we study a dual-source inventory system with endogenous stochastic lead times. We develop a new approach to find an optimal policy. This approach transforms the original problem to a new problem with service-rate controls for a tandem queue with two servers, Server 1 and Server 2. The optimal policy for the new problem can be characterized by a fixed threshold for sending orders to Server 2 (denoted R), and a threshold for sending orders to Server 1, which is given by a function of the state at Server 2 (denoted $S(\cdot)$). When the parameters are such that $R < S(1)$, the optimal policy for the tandem queue completely determines that of the original system. For the case $R \geq S(1)$, we use the tandem queue's policy to construct a good heuristic policy. We also study the special case with identical order costs, which reduces to a single-source system. We compare the optimal policies in the single- and dual-source systems.

Our single-server-queue model of each stage allows us to identify an optimal policy and therefore to sharply quantify the performance of various heuristics. The analysis here, based on exponential processing times, can be extended to Erlang processing times or other Markovian supply networks by expanding the state space of the dynamic programs. These more general supply systems can be constructed to approximate any stochastic lead times. We shall leave this extension to future research. It would also be interesting to extend this work to more than two sources and fixed as well as variable order costs.

Finally, the optimal policy and several of the heuristics use real-time information about the supply system, including which servers are busy. With the adoption of technologies such as RFID and mobile devices, such information is likely increasing. As mentioned earlier, most existing dual-source models assume the buyer has *no* information about conditions at the supplier. Most real systems are somewhere in between; the buyer has *some* information but not all. It would be an interesting future research topic to formalize and analyze notions of partial information in this context, as in, for example, Song and Zipkin (1996), and Guo and Zipkin (2007).

Acknowledgement: The authors would like to thank the Area Editor, the anonymous Associate Editor and referees for their valuable suggestions to improve the paper. This research was partially supported by the Chinese Natural Science Foundation through Grant No. 71390331.

References

Allon G, Van Mieghem J A (2010) Global dual sourcing: Tailored basesurge allocation to near- and offshore production. *Management Sci.* 56(1): 110-124.

Altman E, Gaujal B, Hordijk A (2000) Multimodularity, convexity, and optimization properties. *Math. Oper. Res.* 25(2): 324-347.

Arts J, Basten R, Van Houtum G (2016) Repairable stocking and expediting in a fluctuating demand environment: Optimal policy and heuristics. *Oper. Res.* to appear.

Barankin E W (1961) A delivery-lag inventory model with an emergency provision (the single-period case). *Naval Res. Logist. Quart.* 8(3): 285-311.

Fukuda Y (1964) Optimal policies for the inventory problem with negotiable leadtime. *Management Sci.* 10(4): 690-708.

Gorman S, Baertlein L (2014) The US is facing a shipping port congestion crisis. *Reuters* December 11, 2014, 11:35 pm.

Gross D, Harris C (1971) On one-for-one-ordering inventory policies with state-dependent leadtimes. *Oper. Res.* 19(3): 735-760.

Guo P, Zipkin P (2007) Analysis and comparison of queues with different levels of delay information. *Management Sci.* 53(6): 962-970.

Heyman D (1968) Optimal operating policies for M/G/1 queuing systems. *Oper. Res.* 16(2): 362-382.

Hua Z, Yu Y, Zhang W, Xu X (2015) Structure properties of the optimal policy for dual-sourcing systems with general lead times. *IIE Trans.* 47(8): 841-850.

Janakiraman G, Seshadri S, Sheopuri A (2015) Analysis of tailored base-surge policies in dual sourcing inventory systems. *Management Sci.* 61(7): 1547-1561.

Koole G (2007) *Monotonicity in Markov Reward and Decision Chains*: *Theory and Applications.* Volume **1**, Now Publishers Inc.

Lawson D G, Porteus E L (2000) Multistage inventory management with expediting, *Oper. Res.* 48(6): 878-893.

Li L (1992) The role of inventory in delivery-time competition. *Management Sci.* 38(2): 182-197.

Li Q, Yu P (2014) Multimodularity and its applications in three stochastic dynamic inventory problems. *Manufacturing Service Oper. Management.* 16(3): 455-463

Lippman S (1975) Applying a new device in the optimization of exponential queueing systems. *Oper. Res.* 23(3): 687-710.

Makan A, Simon B (2011) Thai floods hit global hard drive production. Accessed October 20, 2011, `http://www.ft.com/cms/s/2/f0f9a234-fb33-11e0-8756-00144feab49a.html#axzz3k4JlYlJk`.

Moinzadeh K, Nahmias S (1988) A continuous review model for an inventory system with two supply modes. *Management Sci.* 34(6): 761-773.

Moinzadeh K, Schmidt C (1991) An (S-1, S) inventory system with emergency orders. *Oper. Res.* 39(2): 308-321.

Muharremoglu A, Yang N (2010) Inventory management with an exogenous supply process. *Oper. Res.* 58(1): 111-129.

Murota, K (2003) *Discrete Convex Analysis,* SIAM, Philadelphia.

Puterman M (1994) *Markov Decision Processes*: *Discrete Stochastic Dynamic Programming.* John Wiley & Sons, New York.

Rosenshine M, Obee D (1976) Analysis of a standing order inventory system with emergency orders. *Oper. Res.* 24(6):1143–1155.

Schectman J (2013) LOreal Looks to Collaboration for More Agile Manufacturing Process. Accessed September 23, 2013, `http://blogs.wsj.com/cio/2013/09/23/loreal-looks-to-collaboration-for-more-agile-manufacturing-process/`

Scheller-Wolf A, Veeraraghavan S, van Houtum G (2007) Effective dual sourcing with a single index policy. Working paper, Carnegie-Mellon University, Pittsburgh.

Sethi S P, Yan H, Zhang H (2003) Inventory models with fixed costs, forcast updates, and two delivery modes. *Oper. Res.* 51(2): 321-328.

Shang Y, Dunson D, Song J (2014) Exploiting big data in logistics risk assessment via Bayesian nonparametrics. Working Paper, Duke University.

Sheopuri A, Janakiraman G, Seshadri S (2010) New policies for the stochastic inventory control problem with two supply sources. *Oper. Res.* 58(3): 734-745.

Song J, Zipkin P (1996) Inventory control with information about supply conditions. *Management Sci.* 42(10): 1409-1419.

Song J, Zipkin P (2009) Inventories with multiple supply sources and networks of queues with overflow bypasses. *Management Sci.* 55(3): 362-372.

Veatch M, Wein L M (1992) Monotone control of queueing networks. *Queueing Systems* 12(3-4): 391-408.

Veatch M, Wein L M (1994) Optimal control of a two-station tandem production/inventory system. *Oper. Res.* 42(2): 337-350.

Veeraraghavan S, Scheller-Wolf A (2008) Now or later: A simple policy for effective dual sourcing in capacitated systems. *Oper. Res.* 56(4): 850-864.

Weber R, Stidham S (1987) Optimal control of service rates in networks of queues. *Adv. Appl. Probab.* 19(1): 202-218.

Whittemore A, Saunders S (1977) Optimal inventory under stochastic demand with two supply options. *SIAM J. Appl. Math.* 32(2): 293-3057.

Wikipedia (2011) 2011 Thailand floods. `https://en.wikipedia.org/wiki/2011_Thailand_floods`.

Xin L, Goldberg D (2015) Asymptotic optimality of Tailored Base-Surge policies in dual-sourcing inventory systems. Working paper, arXiv:1503.01071v1.

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

Jing-Sheng (Jeannette) Song ("Optimal Policies for a Dual-Sourcing Inventory Problem with Endogenous Stochastic Lead Times") is a professor of operations management at the Fuqua School of Business, Duke University. Her main research interests lie in operations and supply chain management, including developing effective mechanisms to deal with uncertainties.

Li Xiao ("Optimal Policies for a Dual-Sourcing Inventory Problem with Endogenous Stochastic Lead Times") is a postdoc fellow in CUHK Business School, The Chinese University of Hong Kong. Her main research interests lie in operations management and service management.

Hanqin Zhang ("Optimal Policies for a Dual-Sourcing Inventory Problem with Endogenous Stochastic Lead Times") is a professor at School of Business, the National University of Singapore. His current research interests include stochastic models and inventory management.

Paul Zipkin ("Optimal Policies for a Dual-Sourcing Inventory Problem with Endogenous Stochastic Lead Times") is a professor emeritus of operations management at the Fuqua School of Business, Duke University. He is an INFORMS Fellow.