

Production and distribution policy in a two-stage stochastic push-pull supply chain

HYUN-SOO AHN¹ and PHILIP KAMINSKY²

¹*Operations and Management Science, Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI 48109, USA*
E-mail: hsahn@bus.umich.edu

²*Department of Industrial Engineering and Operations Research, University of California, Berkeley, CA 94720, USA*
E-mail: kaminsky@ieor.berkeley.edu

Received September 2002 and accepted September 2004

We consider a model of a two-stage push-pull production-distribution supply chain. The orders arrive at the final stage according to a Poisson process. Two separate operations, which take place at different locations with exponential service times, are required to convert the raw materials into finished goods. When the first operation is completed the intermediate inventory is held at the first stage and then transported to the second stage where the items are produced to order. The objective is to minimize the average sum of the production, transportation, and holding costs. We consider the optimal policy for a version of this model. Our experimental analysis demonstrates that this optimal policy is counter-intuitive. We develop a heuristic based on a deterministic version of this model, and computationally test the heuristic.

1. Introduction

Traditional supply chain strategies are often categorized as either push or pull strategies. Probably, this stems from the manufacturing revolution of the 1980s, in which manufacturing systems were divided into these categories. Interestingly, in the last few years a number of progressive companies have started to employ a hybrid approach, the push-pull supply chain paradigm (Simchi-Levi *et al.*, 2003). In the most traditional type of supply chain, a *push-based supply chain*, production and distribution decisions are based on long-term forecasts, leading to slow reaction to the changing marketplace. In a *pull-based supply chain*, production and distribution are demand driven so that they are coordinated with true customer demand rather than forecast demand. Pull systems lead to: a decrease in lead times, a decrease in system inventory, a decrease in system variability, and an increase in customer service levels; however, they make it more difficult to take advantage of economies of scale in manufacturing and transportation since systems are not planned far ahead in time.

These advantages and disadvantages of push and pull supply chains have led companies to look for a new supply chain strategy that takes advantage of the best of both of the existing approaches. With increasing frequency, this is a push-pull supply chain strategy. In a push-pull strategy, some stages of the supply chain, typically the initial stages, are operated in a push-based manner whereas the remaining stages employ a pull-based strategy. The interface between

the push-based stages and the pull-based stages is known as the push-pull boundary.

In addition, for many firms, the costs associated with transportation and distribution and the inventory holding costs are a large percentage of the total product costs. Indeed, US industry spends more than \$350 billion on transportation and more than \$250 billion on inventory holding costs annually (Lambert and Stock, 1993). In many industries, parts and subcomponents are manufactured across a variety of sites. For example, in the personal computer industry, where lowering the product cost is of paramount importance, the location of final assembly is often different from the location of component assembly. Similarly, in the pharmaceutical industry, manufacturing facilities are expensive to build, and pharmaceutical products have a limited profitable life span (since after patent protection expires, generic manufacturers can manufacture the same product). Thus, multi-purpose plants, which can perform several different manufacturing steps for many different products, are typically built. Once a network of these plants is constructed, new products are manufactured sequentially at several different plants, depending on the particular processes required for manufacture.

Thus, it is important to coordinate both production and inventory policies in multi-stage supply chains. However, although there have been tremendous academic and practical efforts focused on minimizing the inventory level for components and parts within a facility while maintaining efficient production, there has been significantly less work on

coordinating production and distribution simultaneously, particularly when production faces capacity constraints. Efforts in this direction are complicated by the non-linear nature of transportation costs. Indeed, over the last 20 years, driven by the adoption of JIT, CONWIP, and flexible manufacturing systems, manufacturers have devoted significant effort to setup reduction, and thus manufacturing setup costs play a smaller and smaller role in manufacturing decision-making. Most transportation, however, exhibits natural economies of scale; in many cases, an empty truck does not cost that much less to operate than a full one, so it is crucial for successful decision-making approaches for multi-stage manufacturing supply chains to explicitly account for these non-linear transportation costs.

The increasing prevalence of push-pull strategies, coupled with the importance of coordinating production and inventory policies, leads to the need to create models, and develop approaches to manage these systems. In this paper, we initiate this research by considering a simple two-stage push-pull supply chain with non-linear transportation costs. In our model, production at the first stage is completed and “pushed” to the second stage. At this second stage, production is completed only once specific orders have arrived: this is the “pull” stage.

2. Literature review

There has been a variety of literature published on multi-stage inventory systems. Clark and Scarf (1960) show that a base-stock policy is optimal in a finite-horizon periodic problem with no fixed cost and no capacity constraint. Federgruen and Zipkin (1984a, 1984b) show that the optimality of the order-up-to policy is still valid in the case of discounted and average infinite horizon problems. Chen (1994, 1998) demonstrates the effectiveness of the (R, nQ) policy in a similar setting, and Chen and Song (2001) characterize the structure of the optimal policy in a multi-stage inventory model with Markov-modulated demand. A few authors have considered the control of capacitated production systems. Federgruen and Zipkin (1986a, 1986b) consider an inventory policy in a capacitated production system under stationary demand with no fixed cost. Glasserman and Tayur (1994, 1995, 1996) study a multi-stage stationary production-inventory model with capacitated production and demonstrate the effectiveness of infinitesimal perturbation analysis as a tool to obtain an optimal policy. Kapuscinski and Tayur (1998) extend the model to periodic demand. Parker and Kapuscinsky (2004) consider a periodic review problem of a two-stage capacitated echelon inventory system with zero or deterministic lead time and show that the optimal policy is a modified base-stock policy. However, in their model, the replenishment decision is free of economies of scale in shipping quantity so that it is reasonable to conjecture that a variant of an order-up-to policy should be optimal. In most of these papers,

the optimal policy is either a base-stock policy, or a policy with a monotone structure. An exception to this is Duenyas *et al.* (2003), who consider a periodic review model with a fixed cost and lost sales when the system has “must-meet” deterministic demand and random demand at each period. They show that the optimal shipping quantity and rationing policy are not necessarily monotone, because the on-hand inventory may not satisfy existing demands, and can be conserved for deterministic demand in subsequent periods. Our analysis indicates that the optimal policy can be very complex, as well as counter-intuitive, when both production and distribution are considered in a system with non-linear distribution costs and finite production rate. We are not aware of any research which considers similar systems.

There has been another stream of literature on the control of production systems. Ha (1997) characterizes the structure of the optimal policy in make-to-stock systems as a switching-curve-type policy. Carr and Duenyas (2000) consider a make-to-stock/make-to-order system. Duenyas and Patane-Anake (1998), Ahn *et al.* (1999) and Iravani *et al.* (2003) characterize the optimal policy in a serial production line with single or multiple resources. Duenyas and Tsai (2002) consider a two-stage production system where stage 1 can sell its finished goods and provide the components to the downstream stage at the same time. However, in contrast to our proposed research, very little research considers multiple units shipped at a fixed cost.

Also, for relevant *deterministic* models, see Kaminsky and Simchi-Levi (2003) and the references therein. In the next section, we formally introduce our model.

3. The model

We consider a two-stage push-pull supply chain. The orders arrive at stage 2, the final stage, according to a Poisson process with rate λ . Two separate operations, which take place at different locations, are required to convert the raw materials into finished goods. We assume that an infinite supply of raw material is available at stage 1. A single server whose processing time follows an i.i.d exponential distribution with rate μ_i , $i = 1, 2$ is available at each of the two stages. Items are produced at stage 1, the push stage, and can then be either held as inventory at stage 1, or shipped to stage 2, in which case shipping costs are incurred. Items will be processed at stage 2, the pull stage, *only to meet outstanding orders*. Non-negative holding costs are incurred on any intermediate inventory, and penalty costs are incurred for orders waiting to be filled. The holding cost is incurred at a rate h_i ($i = 1, 2$) per unit time per item, whereas a penalty is incurred at rate b per unit time per outstanding order. We assume that holding costs are non-decreasing (i.e., $h_1 \leq h_2 \leq b$). Shipping is assumed to be immediate, and a fixed cost, K , is incurred for each shipment to reflect shipping economies of scale. This model is illustrated in Fig. 1.




Fig. 1. Two-stage push-pull system.

We model this problem as a Markov Decision Process (MDP) and attempt to minimize the long-run average cost per unit time. Without loss of generality, we uniformize λ , μ_1 and μ_2 such that $\lambda + \mu_1 + \mu_2 = 1$. The states of the MDP are $S(t) = (n_1(t), n_2(t), n_3(t))$ where $n_1(t)$ denotes the number of units held in inventory at stage 1 at time t , n_2 denotes the number of units held in inventory at stage 2 (before the completion of the stage 2 operation) at time t and n_3 denotes the number of outstanding orders at time t . We assume that only Markovian stationary deterministic policies are under consideration (and call the class of these policies Π). In other words, the policies in Π specify the decision solely as a function of the current state. Since processing times and interarrival times are exponential, it is easy to see that $\{S(t)\}$ is a continuous-time Markov chain where the transition rate at any state is bounded by unity, and therefore $\{S(t)\}$ is uniformizable. By using the results of Lippman (1975), we can translate the original continuous-time optimization problem into an equivalent (discrete-time) MDP with state $s = (n_1, n_2, n_3)$. Decisions can be made upon the arrival of a new order or at any service completion. It is easy to see that stage 2 will always be busy as long as it has product to process and outstanding orders to fill (i.e., $n_2 > 0$ and $n_3 > 0$). At any decision epoch, the production decision at stage 1 must be made (whether or not to initiate production of a unit at stage 1), and if $n_1 > 0$, the distribution decision must be made (i.e., the quantity greater than or equal to zero which should be shipped must be decided). Using the uniformization technique, we write the equivalent discrete-time dynamic programming. For example, the optimality equation for a state (n_1, n_2, n_3) , $n_1 > 0$, $n_2 = 0$, $n_3 > 0$ must satisfy:

$$v(n_1, n_2, n_3) + g = \sum_{i=1}^2 h_i n_i + b n_3 + \min_{0 \leq s \leq n_1} \left[\begin{aligned} &K \mathbf{1}_{(s>0)} + (h_2 - h_1)s + \lambda v(n_1 - s, s, n_3 + 1) \\ &+ \mu_2 v(n_1 - s, s - 1, n_3 - 1) \\ &+ \mu_1 \min[v(n_1 - s + 1, s, n_3), v(n_1 - s, s, n_3)] \end{aligned} \right],$$

where g is the average cost per unit time and $v(n_1, n_2, n_3)$ is the relative value function from state (n_1, n_2, n_3) . The

optimality equation for other states can be written in a similar way. In order to ensure the existence of a solution for the average cost optimality equation, we focus on the systems with $(\lambda/(\mu_1 + \mu_2)) < 1$. To see that this is indeed a sufficient condition, consider a stationary policy that produces at stage 1 only if $n_2 = 0$ and $n_3 > 0$, and ships one unit at a time when the inventory level at stage 2 reaches zero and there is at least one outstanding order. It is easy to see that, in the long-run, the stochastic process induced by this policy is identical to that of a M/G/1 queue whose interarrival time is exponentially distributed with rate λ , and whose service time is a convolution of two exponential distributions with rates μ_1 and μ_2 . Under these conditions, the continuous-time Markov chain induced by this policy has only one recurrent class containing $(0, 0, 0)$, whereas both the expected first passage time and the expected total cost incurred from any arbitrary state to $(0, 0, 0)$ are finite. Therefore, this policy induces a unichain (i.e., a Markov chain with one positive recurrent class and one transient class), and the average expected cost associated with this policy is finite. The existence of the optimal stationary policy follows directly from the results of Meyn (1997).

To develop insight into the structure of optimal solutions for this model, we performed computational testing using the value iteration algorithm with a variety of parameters. For the value iteration, we truncated the state space such that $n_i \leq 100$ by redirecting any transition to a state outside the truncated state space to the nearest state. Through an extensive computational study, we observed the following. In all cases, it is optimal to not ship if $n_2 > 0$. (We prove a related result below.) For fixed values of n_1 , the shipping decision increases monotonically in n_3 . That is, there is some level of n_3 , which is a function of n_2 , below which it is optimal not to ship, and above which it is optimal to ship. We note that in all observed cases, this level was greater than zero. However, for fixed values of n_1 , the shipping quantity fluctuates in increasing n_3 . For fixed values of n_3 , the shipping quantity increases in n_1 , and then cycles. The maximum and minimum values of the optimal shipping quantities increase with increasing n_3 . This result is surprising, since it implies that in some cases it is optimal

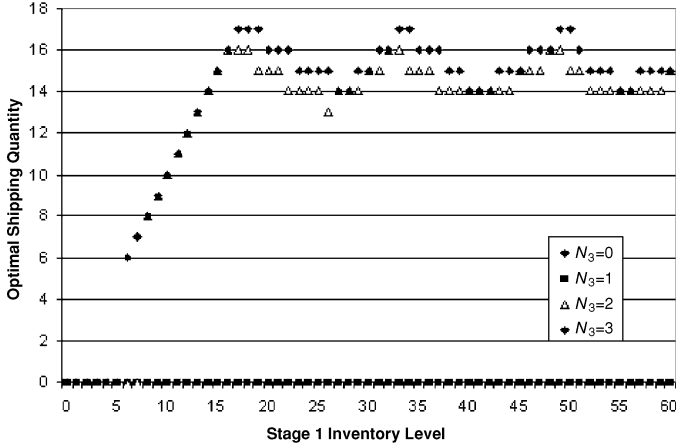


Fig. 2. An example of the optimal shipping quantity when $n_3 = 0, 1, 2$ and 3 .

to ship a smaller quantity as the number of outstanding orders increases. Note that we expended considerable effort in eliminating truncation of countable state space and round-off error as potential reasons for this striking behavior. For example, we tried conducting a value iteration on a considerably larger state space, and extrapolating the relative value function for states outside the boundary, but the optimal policy was still non-monotonic.

Figure 2 illustrates that the optimal shipping quantity is not always monotone in the amount of on-hand inventory at stage 1 or the number of outstanding orders. In this example, it is optimal to ship 17 units at state $(19,0,3)$, but 14 units at state $(27,0,3)$, even though this second state has more units to ship from stage 1.

We characterize a sufficient condition for *not* shipping in this model:

Lemma 1. *For any state (n_1, n_2, n_3) , it is optimal not to ship as long as $n_2 > 0$.*

This result can be easily shown using an interchange argument.

Unfortunately, the lack of monotonicity in the shipping quantities makes additional structural properties of this model difficult to determine. Thus, we are motivated to consider various restrictions of this model. In Section 4, we consider a version of this model for which the possible set of shipping quantities is limited. In Section 5, we propose a heuristic for the model discussed in Section 4. Finally, in Section 6, we apply this heuristic to the original model, and computationally test its effectiveness. In Section 7, we conclude, and discuss potential future research.

4. Restricting the possible shipping quantities

As we observed above, the optimal policy of the original formulation does not possess the hoped for monotonicity

property. In addition, our original formulation, which considers shipping every quantity between zero and the current inventory level at each decision epoch, is computationally inefficient. We therefore consider a simplified version of the problem, in which the firm ships the minimum of the maximum transportation capacity and the current inventory level whenever the decision to ship is made. We note that this may be a reasonable restriction in practice, since firms may experience a physical limit on the quantity that can be shipped at a time (one truckload, for example). In addition, we assume that shipments only occur when the inventory at stage 2 is zero, which we proved above is optimal in our original model, although not necessarily in this model. As we report in Section 6, computational testing shows that this restricted model performs about as well as the original model if the shipping quantity is selected appropriately.

At state $s = (n_1, n_2, n_3)$, the set of feasible policies, A_s is given by:

$$A_s = \{(i, q) | i = 1 \text{ if produce, } 0 \text{ otherwise, } q \in \{0, \mathbf{1}_{\{n_2=0\}} \min[n_1, Q]\}\}.$$

Using the uniformization technique, the optimal average cost, g and the relative value function, $v(\cdot)$, must satisfy the following equations:

$$v(n_1, n_2, n_3) + g = \sum_{i=1}^2 h_i n_i + b n_3 + \min \left[\begin{aligned} &K + (h_2 - h_1)Q + \lambda v(n_1 - Q, Q, n_3 + 1) \\ &+ \mu_2 v(n_1 - Q, Q - 1, n_3 - 1) \\ &+ \mu_1 \min[v(n_1 - Q + 1, Q, n_3), v(n_1 - Q, Q, n_3)] \\ &\lambda v(n_1, 0, n_3 + 1) + \mu_2 v(n_1, 0, n_3) \\ &+ \mu_1 \min[v(n_1 + 1, 0, n_3), v(n_1, 0, n_3)] \end{aligned} \right] \text{ for } n_1 \geq Q, n_2 = 0, n_3 > 0.$$

$$v(n_1, n_2, n_3) + g = \sum_{i=1}^2 h_i n_i + b n_3 + \min \left[\begin{aligned} &K + (h_2 - h_1)n_1 + \lambda v(0, n_1, n_3 + 1) \\ &+ \mu_2 v(0, n_1 - 1, n_3 - 1) \\ &+ \mu_1 \min[v(1, n_1, n_3), v(0, n_1, n_3)] \\ &\lambda v(n_1, 0, n_3 + 1) + \mu_2 v(n_1, 0, n_3) \\ &+ \mu_1 \min[v(n_1 + 1, 0, n_3), v(n_1, 0, n_3)] \end{aligned} \right] \text{ for } n_1 < Q, n_2 = 0, n_3 > 0.$$

$$v(n_1, n_2, n_3) + g = \sum_{i=1}^2 h_i n_i + b n_3 + \lambda v(n_1, n_2, n_3 + 1) + \mu_2 v(n_1, n_2 - 1, n_3 - 1) + \mu_1 \min[v(n_1 + 1, n_2, n_3), v(n_1, n_2, n_3)] \text{ for } n_2 > 0, n_3 > 0.$$

$$v(n_1, n_2, n_3) + g = \sum_{i=1}^2 h_i n_i + b n_3 + \lambda v(n_1, n_2, 1) + \mu_2 v(n_1, n_2, 0) + \mu_1 \min[v(n_1 + 1, n_2, 0), v(n_1, n_2, 0)] \text{ for } n_2 \geq 0, n_3 = 0.$$

In the next subsection, we characterize the optimal policy of this model.

4.1. The optimal policy for the capacitated shipping problem: Counter-intuitive observations

As we mentioned above, the restriction of the shipping quantity simplifies the computational efforts. Furthermore, as the shipping decision itself is monotone in the original problem and the shipping quantity initially increases then fluctuates, we had anticipated that the monotonicity would continue to hold when we imposed a reasonable monotone shipping quantity which resembled the spirit of the optimal shipping quantity, such as $\min(n_1, Q)$. However, even this simplification does not result in a simple optimal policy. In fact, in many examples, the optimal policy is not only complex, but also counter-intuitive, and very different from results found in the literature for more simple make-to-order models. This counter-intuitive behavior seems to be due to the non-linear shipping costs. Later in the paper, we computationally compare this model to an analogous model with linear shipping costs. In this subsection, we discuss the structure (or lack of structure) of the optimal policy when the shipping quantity is bounded by Q . Clearly, the overall performance of the system will be significantly affected by the choice of Q . In the next section, we discuss a heuristic approach to selecting a good Q value.

First, by studying a variety of computational examples, we observe that neither production nor shipping decisions are monotone in most of the state variables. In particular:

1. Even if it is optimal to ship at state $(n_1, 0, n_3)$, it is not necessarily optimal to ship at state $(n_1, 0, n_3 + 1)$.
2. Even if it is optimal to ship at state $(n_1, 0, n_3)$, it is not necessarily optimal to ship at state $(n_1 + 1, 0, n_3)$, even in the case when $n_1 > Q$ so the shipping quantity remains the same.
3. The production policy at stage 1 is not necessarily monotone in the inventory level at stage 1, that is, even if it is optimal not to produce at state (n_1, n_2, n_3) , it may be optimal to produce at state $(n_1 + 1, n_2, n_3)$.
4. The production policy at stage 1 is not necessarily monotone in the number of outstanding orders. Even if it is optimal to produce at state (n_1, n_2, n_3) , it is not necessarily optimal to produce at state $(n_1, n_2, n_3 + 1)$.

Whereas in some examples the optimal policy is monotone (mostly when the fixed shipping cost is close to zero or small), there are many other examples showing that such monotonicity does not hold in general. As can be seen in Fig. 3, the optimal policy cannot be characterized by a simple switching curve. First, consider the lack of monotonicity in n_3 . In the example, when $n_1 = 10$, it is optimal not to ship when $n_3 \in \{0, 1, 5, 6, 7\}$, but optimal when n_3 takes on any other value. Although this seems counter-intuitive, consider the following possible explanation. In the optimal policy, units shipped from stage 1 cover current and

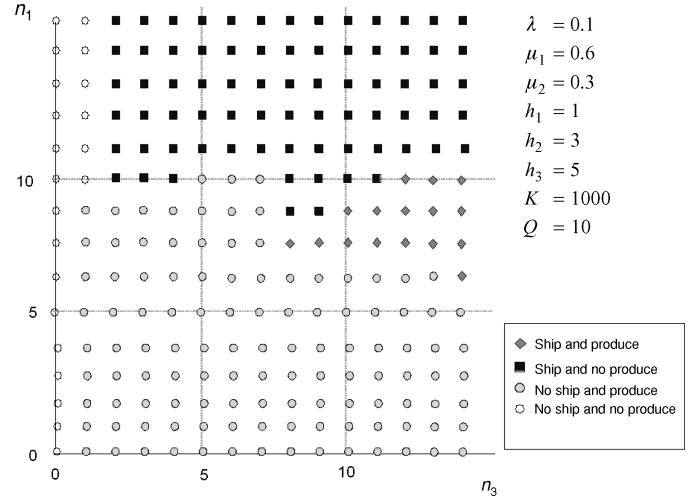


Fig. 3. An optimal policy at $n_2 = 0$ in a capacitated shipping problem.

(possibly) future outstanding orders. Any shipment decision balances the increased holding cost in stage 2 with the backorder cost. In addition, any shipment less than Q also accounts for increased shipping costs. Also, once a shipment is made, the next shipment cannot be made until the inventory at stage 2 is once again zero, and the time until the inventory at stage 2 is once again zero depends on the backorder level during the first shipment, and the amount shipped. For very low backorder levels ($n_3 = 1$ in this example), increased holding and transportation costs offset the increased backorder costs, so it is optimal not to ship. For certain low backorder levels ($n_3 = 2$ in this example), the benefit from shipping, decreased backorder costs, may outweigh the benefits of waiting and producing more before shipping. This effect is enhanced by the fact that the backorder level is likely to be relatively low at the time of the next shipment, since the amount being shipped is relatively small, and much of it is already allocated to existing backorders. However, for other small backorder levels ($n_3 = 5, 6, 7$ in this example), the benefit of delaying shipment (consequently shipping more in a later time) dominates the benefit of shipping immediately. Finally, as the backorder level gets larger, it again makes sense to ship immediately, since the backorder level at the next shipment will be large regardless of whether or not some quantity is immediately shipped.

The non-monotonicity of shipping in n_1 can be explained similarly. When there are few units in stage 1, shipping immediately is suboptimal since doing so fails to take advantage of the economies of scale in shipping. However, as the number of units in stage 1 increases, the shipping cost per unit decreases. However, at some point, the increasing stage 1 holding cost may begin to dominate, so that it therefore becomes optimal to wait until additional orders arrive before shipping.




Fig. 4. The non-monotonicity of the production policy in n_3 when $n_2 = 0$.

Figure 4 shows an example where the optimal production policy is not monotone in backorder levels. For instance, it is optimal to produce for $n_3 \leq 3$ or $n_3 \geq 8$, but optimal not to produce for $4 \leq n_3 \leq 7$ when $n_1 = 7$. For low backorder levels, it is optimal to produce, but not to ship since the backorder cost is not high enough to justify immediate shipping (with resultant increased per unit shipping charges). However, as the backorder level gets larger, it makes sense to ship immediately for $n_3 \geq 4$ as in Fig. 4. Whereas it is optimal not to produce after a shipment for medium backorder levels ($n_3 = 4, 5, 6, 7$ in this example), it is optimal to produce after a shipment for high backorder levels. This is true because with a high backorder, the time until the next time that the inventory is again zero at stage 2 decreases. Thus, production must continue to ensure that there is enough material to constitute the next shipment. On the other hand, with a lower backorder, it will be longer until the next shipment is required. Indeed, the time at which the inventory at stage 2 is zero after a shipment stochastically decreases in the number of outstanding orders.

This phenomenon is affected by the relative production and arrival rates. When the production rate at stage 1 is higher than the rate at stage 2, it is more likely to be optimal to wait before producing at stage 1, as the time until another shipment is required is greater. As backorder costs increase, the value of waiting until the next shipment decreases, so production is more likely, although only if relative processing rates suggest that this production is likely to be shipped relatively quickly.

This example contradicts the intuition established by the optimal solution to many make-to-order models. Conventional intuition suggests that the optimal production rate should increase as the number of outstanding orders increases so that the system can clear orders as soon as possible. Therefore, as the backorder cost becomes high, it be-

comes urgent to reduce outstanding orders. However, at the push-pull boundary, the finite capacity at stage 2 dampens such urgency. When there is enough time to produce units at stage 1 while stage 2 is busy, the production can be delayed. On the other hand, the benefit of increasing shipping quantities increases when the backorder cost is small and shipping cost is high. Therefore, it becomes optimal not to ship, but to produce when there are few backorders.

The non-monotonicity of the optimal production policy in n_1 is even more intriguing. Observe in Fig. 5 that it is optimal to produce at state $(9, 0, 2)$ although it is optimal not to produce at state $(8, 0, 2)$. Most previous analyses of related models imply that it is optimal to produce until the inventory level reaches a certain point (although this point may be state dependent). This example contradicts that intuition, since it is optimal to produce at a particular state, $(9, 0, 2)$, but not at a state with a lower inventory, $(8, 0, 2)$. In general, this counter-intuitive behavior seems to occur when n_2 is low. In fact, when $n_2 = 0$ or 2, the production policy is characterized by a monotone switching curve in the opposite direction to what intuition may suggest.

Although this seems perplexing, consider the following possible explanation. When the fixed shipping cost is extremely high as in this example, it is optimal to ship Q units when a shipment occurs. Therefore, stage 1 will always produce Q units before a shipment occurs. Now, if holding costs are similar at stage 1 and stage 2, it makes little difference if inventory is held at stage 1 or at stage 2. Thus, when the inventory level at stage 1 is close to Q , it becomes optimal to produce more units (up to Q), then ship to stage 2, rather than storing inventory at stage 1. However, when the inventory level at stage 1 is low, it is optimal to wait for more outstanding orders before producing additional units to ship. Thus, in this case, the optimal policy is monotone in the counter-intuitive direction (i.e., the more inventory there is, the more valuable it is to produce more). However, as the level of inventory at stage 2 increases, this effect diminishes, as the time until the next shipment increases. In Fig. 5, for example, this “opposite direction monotonicity” disappears when $n_2 \geq 6$.

Although the optimal policy to this problem is clearly complex, we are able to partially analytically characterize its structure. We show that it is always optimal to ship if there are at least Q units of inventory and at least Q units of backorder.

Lemma 2. For any state $(n_1, 0, n_3)$ such that $n_1 \geq Q$ and $n_3 \geq Q$, it is optimal to ship independent of production policy.

Proof. We prove the claim by a sample path argument that any policy with delayed shipping can be improved by a policy with immediate shipping. Without loss of generality, we assume that $t = 0$. Let w be a sample path in a probability space \mathbf{P} large enough to contain all future arrivals and service times. Suppose that in the optimal policy Π^* , shipping is delayed by $t_s > 0$. Let T denote the first time that $n_2(t)$




Fig. 5. The non-monotonicity of the production policy in n_1 ($\lambda = 0.05$, $\mu_1 = 0.5$, $\mu_2 = 0.45$, $h_1 = 1.0$, $h_2 = 1.05$, $b = 1.10$, $K = 1000$ and $Q = 12$).

becomes zero after shipping Q units at time t_s on a given sample path, w . Also, let $p^{\Pi^*}(t)$ be the cumulative units produced in stage 1 by time t , $D^{\Pi^*}(t)$ be the cumulative units produced in stage 2 at time t under policy Π^* , and $A(t)$ be the cumulative number of arrivals to the system by time t . It can be easily shown that the stage 2 remains idle until time t_s (that is when the first shipment occurs) and works to produce Q consecutive units without idling. The process will reach state $(n_1 + P^{\Pi}(T) - Q, 0, n_3 + A(T) - Q)$ under policy Π^* at time T and the total cost accumulated over this interval $[0, T]$, denoted as $C^{\Pi^*}((0, T]; w)$, can be expressed as the following:

$$\begin{aligned} C^{\Pi^*}((0, T]; w) &= \int_0^{t_s} [h_1(n_1 + p^{\Pi^*}(t)) + b(n_3 + A(t))]dt \\ &\quad + \int_{t_s}^T [h_1(n_1 + p^{\Pi^*}(t) - Q) \\ &\quad + h_2(Q - D^{\Pi^*}(t)) + b(n_3 + A(t) \\ &\quad - D^{\Pi^*}(t))]dt. \end{aligned}$$

Consider another policy Π on the same sample path, that mimics the production of policy Π , but ships Q units at time 0 instead of t_s . Under this policy, Q units are produced without idling and idling occurs in the interval $(T - t_s, T]$ at stage 2 (i.e., $D^{\Pi}(t) = D^{\Pi^*}(t + t_s)$ for $t \in [0, T - t_s]$ and $D^{\Pi}(t) = D^{\Pi^*}(T) = Q$ for $t \in [T - t_s, T]$.) After T , policy Π mimics the production and shipping decisions of policy Π^* . For every sample path, it can be easily shown that the sample paths under both policies coincide after T . The total cost accrued over this interval $[0, T]$, denoted as $C^{\Pi}((0, T]; w)$, is:

$$\begin{aligned} C^{\Pi}((0, T]; w) &= \int_0^T [h_1(n_1 + p^{\Pi}(t) - Q) + b(n_3 \\ &\quad + A(t) - D^{\Pi}(t))]dt \\ &\quad + \int_0^{T-t_s} [h_2(Q - D^{\Pi}(t))]dt \\ &= \int_0^T [h_1(n_1 + p^{\Pi^*}(t) - Q)]dt \\ &\quad + \int_0^{T-t_s} [b(n_3 + A(t) - D^{\Pi^*}(t + t_s)) \\ &\quad + h_2(Q - D^{\Pi^*}(t + t_s))]dt \\ &\quad + \int_{T-t_s}^T [b(n_3 + A(t) - Q)]dt. \end{aligned}$$

Comparing the cost associated with policy Π^* and the cost associated with policy Π , we have:

$$C^{\Pi^*}((0, T]; w) - C^{\Pi}((0, T]; w) = \int_0^{t_s} Q(h_1 + b)dt \geq 0.$$

The result holds for every sample path. Therefore, for any policy under which the shipping is delayed, it is possible to construct an immediate shipping policy which is better for

every sample path. This contradicts the optimality of policy Π^* . ■

4.2. Selecting a good transportation quantity

Of course, the performance of this restricted model relative to the original model depends on the selection of a good Q value. In Section 6, we demonstrate that if we select the best possible Q value, the gap between the average cost of this policy and the average optimal cost is very small. In this section, we suggest a *heuristic approach* to selecting a good transportation quantity Q . We note that if we have an infinite supply and no holding cost at stage 1, and an infinite service rate at stage 2, it is well known that the stage 2 problem is solved by a (Q, r) policy, in which Q units are ordered when the inventory level falls below r (see, for example, Zipkin (2000)). In addition, researchers have found that the deterministic Economic Order Quantity (EOQ) is an effective heuristic choice for Q in this (Q, r) model. Thus, we are motivated to solve a deterministic EOQ-style problem in order to determine the Q value to apply to our heuristic in our stochastic model.

In particular, we consider a two-stage deterministic model. At the first stage, items are manufactured at rate μ_1 . Orders are transferred from the first to the second stage instantaneously. Demand arrives at the second stage at rate λ , and demand is either backordered, or met by converting intermediates to finished goods at rate μ_2 . As with our continuous model, there is a fixed cost of K each time items are transferred from stage 1 to stage 2, a holding cost of h_1 per unit time per item at stage 1, a holding cost of h_2 per unit time per item at stage 2, and a penalty cost of b per waiting order per unit time. We focus on policies with repeating cycles, and observe that it will be optimal for the backorder to go to zero during each cycle.

Following standard EOQ derivations, we assume that there is no inventory at time 0, and that an order of size Q is transferred from stage 1 to stage 2 at time 0. We also assume that there are B units of backorder at time 0. These backorders cannot be met until the raw materials are processed. Figure 6(a-c) demonstrates the evolution of backorder level and inventory levels in this system.

Figure 6(a) shows the evolution of stage 1 inventory levels over two inventory cycles and Fig. 6(b) shows the evolution of stage 2 inventory levels over two inventory cycles. Figure 6(c) shows the evolution of the backorder levels over 2 cycles. At the start of the cycle, Q inventory units arrive at stage 2. The B backorder level goes to zero at rate $\mu_2 - \lambda$, as newly arriving orders must also be filled. Thus, $T_1 = B/(\mu_2 - \lambda)$. Once the backorder reaches zero, new orders are processed at rate λ , since they can only be processed as fast as they arrive, and so the stage 2 inventory goes to zero at rate λ during interval T_2 . Finally, during interval T_3 , there is no raw material, so backorders once again accumulate. Thus, $T_3 = B/\lambda$, and since $T = T_1 + T_2 + T_3 = Q/\lambda$, $T_2 = T - T_1 - T_3$. Meanwhile, at stage 1, production starts

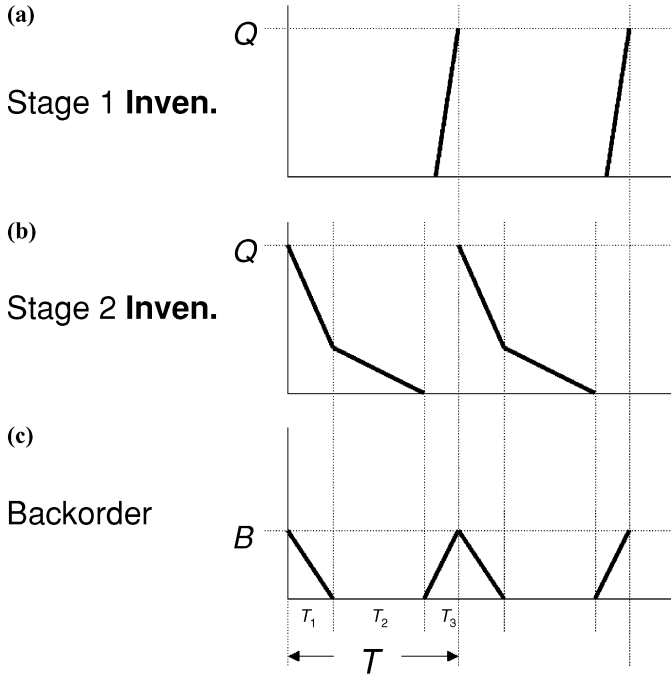


Fig. 6. Inventory and backorder levels: (a) the stage 1 inventory level; (b) the stage 2 inventory level; and (c) the backorder level.

so that Q units will be ready at the end of the cycle, so production starts $T_4 = Q/\mu_1$ before the end of the cycle. The average cost per unit time, $TC(Q, B)$ is equal to:

$$TC(Q, B) = \frac{B^2}{2Q} \frac{\lambda}{\mu_2 - \lambda} b + \frac{B^2}{2Q} b + \frac{2Q(\mu_2 - \lambda)B\lambda - \mu_2\lambda B^2}{2Q(\mu_2 - \lambda)^2} h_2 + \frac{(Q(\mu_2 - \lambda) - \mu_2 B)^2}{2Q(\mu_2 - \lambda)^2} h_2 + \frac{\lambda Q}{2\mu_1} h_1 + \frac{K\lambda}{Q},$$

where the first term represents the average backorder cost during T_1 , the second term represents the average backorder cost during T_3 , the third term represents the average stage 2 inventory holding cost during T_1 , the fourth term represents the average stage 2 inventory holding cost during T_2 , the fifth term represents the stage 1 inventory cost, and the final term represents the average fixed order cost.

First, we observe that each of the terms of this total cost equation are increasing in B , so it is always optimal to avoid backorders in this deterministic model. Thus, the total cost equation simplifies to:

$$TC'(Q, B) = \frac{Q}{2} h_2 + \frac{\lambda Q}{2\mu_1} h_1 + \frac{K\lambda}{Q}.$$

The total cost is minimized when Q is equal to:

$$Q^* = \sqrt{\frac{2\lambda\mu_1 K}{\lambda h_1 + h_2\mu_1}}. \quad (1)$$

We note that this number is likely to be a fraction. In Section 6, we compare the effectiveness of $Q = \lceil Q^* \rceil$ to that of the best possible Q value, where we use the ceiling of Q^* as the conservative choice, since the backorder cost is much higher than the holding cost.

Below, we propose a heuristic for this model, and use the observations above to set the parameters for this heuristic.

5. The (L, Q, α) policy

Based on the analysis of the deterministic system above, we investigate a heuristic we call the (L, Q, α) policy. Under this policy, when inventory at stage 2 minus the backorder level falls below L , the firm produces to raise the inventory to Q . Also, if $n_2 = 0$ and the inventory at stage 1 plus α times the backorder level is greater than or equal to Q , the firm ships a minimum of Q units of the available inventory at stage 1 to stage 2. We implement this policy within the same discrete-time dynamic programming framework as before. For example, when $n_2 - n_3 < L$, $0 < n_1 < Q$, $n_2 = 0$, and $n_1 + \alpha n_3 \geq Q$, the following equation holds:

$$v(n_1, 0, n_3) + g = K + \sum_{i=1}^2 h_i n_i + b n_3 + (h_2 - h_1) \min\{Q, n_1\} + \lambda v(n_1 - \min\{Q, n_1\}, n_2) + \min\{Q, n_1\}, n_3 + 1 + \mu_1 v(n_1 - \min\{Q, n_1\} + 1, n_2) + \min\{Q, n_1\}, n_3 + \mu_2 v(n_1 - \min\{Q, n_1\}, n_2 + \min\{Q, n_1\}) - \mathbf{1}_{(n_3 > 0)}, n_3 - \mathbf{1}_{(n_3 > 0)}.$$

5.1. Estimating policy parameters

To implement this policy, we need values for Q , L , and α . For Q we use the ceiling of Q^* determined in Equation (1).

For our estimate of L , we observe that since backorders are more expensive than inventory holding, we would like to be reasonably sure that no more than L additional orders arrive during the time that Q units are being manufactured at stage 1. The expected value of this time is Q/μ_1 . The number of arrivals A during this time has a Poisson distribution with mean $\lambda Q/\mu_1 Q$. To be reasonably sure that no more than Q/μ_1 customers arrive during that time, we select L as follows:

$$\text{Prob}\{A < L\} < 90\%, \quad \text{Prob}\{A < L + 1\} > 90\%.$$

We elected to use these criteria because 90% appears to be an effective level for our heuristic, but requiring $\text{Prob}\{A > L\}$ to be greater than 10% often required L to be too big, and gave us a probability closer to 100% than to 90%.

Table 1. Performance when the shipping quantity is restricted and also when the shipping costs are linear

h_1	h_2	b	K	λ	μ_1	μ_2	<i>Optimal</i>	<i>Restrict. ship</i>	<i>Linear</i>
1	2	5	250	0.2	0.4	0.4	22.2961	22.2966	10.5471
1	4	15	250	0.2	0.4	0.4	38.6485	38.6502	29.4737
1	2	5	1000	0.2	0.4	0.4	37.7050	37.7056	14.5638
1	4	15	1000	0.2	0.4	0.4	59.7015	59.7025	34.9888
1	2	5	4000	0.2	0.4	0.4	68.0791	68.0804	26.5500
1	4	15	4000	0.2	0.4	0.4	101.1060	101.1076	50.8105
1	2	5	250	0.15	0.6	0.25	21.7331	21.7344	14.4781
1	4	15	250	0.15	0.6	0.25	42.1668	42.1690	33.8936
1	2	5	1000	0.15	0.6	0.25	34.5979	34.599 685	18.6450
1	4	15	1000	0.15	0.6	0.25	60.1309	60.133 752	40.1438
1	2	5	4000	0.15	0.6	0.25	59.9318	59.934 515	26.9787
1	4	15	4000	0.15	0.6	0.25	95.6139	95.6182	52.6442
1	2	5	250	0.15	0.25	0.6	16.9865	16.9876	9.2840
1	4	15	250	0.15	0.25	0.6	26.3363	26.3505	21.8587
1	2	5	1000	0.15	0.25	0.6	29.8138	29.8154	11.1591
1	4	15	1000	0.15	0.25	0.6	43.6287	43.6309	24.3597
1	2	5	4000	0.15	0.25	0.6	55.4480	55.4504	26.1596
1	4	15	4000	0.15	0.25	0.6	78.3368	78.3406	41.6330

Finally, since α needs to reflect the relationship between the stage 2 holding cost and the backorder cost, we set:

$$\alpha = \frac{b}{h_1}.$$

In the next section, we computationally test the effectiveness of this heuristic.

6. Computational analysis

Our computational experiments are designed to characterize: (i) the loss associated with restricting the shipping quantity to the minimum of some Q value and the available inventory; (ii) the impact of non-linear shipping costs; (iii) the effectiveness of selecting a heuristic Q level; and (iv) the performance of our (L, Q, α) heuristic.




Fig. 7. The objective value as a function of Q for the restricted shipping case.

We tested a variety of combinations of system parameters:

- $h_1, h_2, b \in \{(1, 2, 5), (1, 4, 15)\}$;
- $K \in \{250, 1000, 5000\}$;
- $\lambda, \mu_1, \mu_2 \in \{(0.2, 0.4, 0.4), (0.15, 0.6, 0.25), (0.15, 0.25, 0.6)\}$.

We vary the costs so that they increase either a relatively small amount or a relatively large amount between stages, we consider a variety of fixed costs, and we consider processing rates which are much faster at stage 1, or much faster at stage 2. All of the examples are solved through a value iteration on a sufficiently large truncated state space to alleviate any boundary effect.

6.1. The effectiveness of restricting the shipping quantities

For each of the combinations of system parameters listed above, we tested a variety of Q values $Q = 1, 2, \dots, 70$. Table 1 displays these results. Note that we essentially lose nothing by making this assumption. Indeed, in all cases the objective value for the best possible Q is almost identical to the optimal objective value for the original model. We have to consider four decimal digits to see a difference in most cases. In addition, this result is relatively insensitive to Q . Figure 7 is a plot of the objective value against the Q value for one sample problem (the problem instance is: $h_1 = 1, h_2 = 2, b = 5, K = 250, \lambda = 0.2, \mu_1 = 0.4$ and $\mu_2 = 0.4$). Observe that the objective value is very close to optimal for a large range of Q values. The performance of all the sample problems we considered was similar.

6.2. The impact of non-linear shipping costs

To determine the impact of non-linear shipping costs on the total system cost, in the last column of Table 1, we present the system cost if the shipping cost is linear, and equal to K/Q where Q is the optimal Q value for the restricted shipping problem. In all cases, this system is significantly less expensive than the system with fixed shipping costs.

6.3. Effectiveness of the heuristic Q selection

Of course, it takes considerable time to evaluate a range of Q values for a particular problem instance. One alternative is to use the heuristic Q value derived in Section 4.2. Table 2 compares the best Q value with this heuristic Q value, and the resulting objective function value of these Q selections. In all cases, the heuristic Q is less than the optimal Q . However, in all cases, the objective value with the heuristic Q value is within 2% of the optimal objective value for the original model, and in most cases it is within 1%.

6.4. Effectiveness of the (L, Q, α) heuristic

We calculated $L, Q,$ and α values using the techniques described in this paper, and computationally compared the performance of the heuristic to the optimal solution. As can be seen in Table 3, the heuristic in general performed very well. For all of these cases, the average cost of implementing the heuristic policy is within 15% of the optimal cost, and frequently much better. The heuristic performs worst when the stage 1 production rate is relatively low. This suggests that in these cases, there is little “margin for error” in determining when to start production, and the heuristic is not starting production at the best possible time.

Table 2. Performance of the restricted shipping case using heuristic Q values

h_1	h_2	b	K	λ	μ_1	μ_2	<i>Optimal</i>	<i>Restricted shipping</i>	<i>Best Q</i>	<i>Heuristic</i>		<i>Performance vs. optimal (%)</i>
										<i>Q value</i>	<i>Obj.</i>	
1	2	5	250	0.2	0.4	0.4	22.30	22.30	36	7	22.42	0.53
1	4	15	250	0.2	0.4	0.4	38.65	38.65	8	5	38.82	0.44
1	2	5	1000	0.2	0.4	0.4	37.71	37.71	37	13	38.01	0.82
1	4	15	1000	0.2	0.4	0.4	59.70	59.70	17	10	59.70	0.00
1	2	5	4000	0.2	0.4	0.4	68.08	68.08	46	26	68.50	0.62
1	4	15	4000	0.2	0.4	0.4	101.11	101.11	29	19	101.27	0.16
1	2	5	250	0.15	0.6	0.25	21.73	21.73	9	6	21.86	0.60
1	4	15	250	0.15	0.6	0.25	42.17	42.17	6	5	42.21	0.11
1	2	5	1000	0.15	0.6	0.25	34.60	34.60	18	12	34.72	0.36
1	4	15	1000	0.15	0.6	0.25	60.13	60.13	12	9	60.26	0.22
1	2	5	4000	0.15	0.6	0.25	59.93	59.93	36	24	60.04	0.18
1	4	15	4000	0.15	0.6	0.25	95.61	95.62	24	17	95.88	0.28
1	2	5	250	0.15	0.25	0.6	16.99	16.99	12	6	17.11	0.71
1	4	15	250	0.15	0.25	0.6	26.34	26.35	5	5	26.35	0.05
1	2	5	1000	0.15	0.25	0.6	29.81	29.82	30	11	30.23	1.40
1	4	15	1000	0.15	0.25	0.6	43.63	43.63	15	9	43.88	0.57
1	2	5	4000	0.15	0.25	0.6	55.45	55.45	30	22	56.04	1.08
1	4	15	4000	0.15	0.25	0.6	78.34	78.34	22	17	78.88	0.69

Table 3. Performance of the heuristic

h_1	h_2	b	K	λ	μ_1	μ_2	L	Q	α	Z	Perf. (%)
1	2	5	250	0.2	0.4	0.4	6	7	5.00	24.2809	8.90
1	4	15	250	0.2	0.4	0.4	4	5	15.00	40.9793	6.03
1	2	5	1000	0.2	0.4	0.4	9	13	5.00	40.4565	7.30
1	4	15	1000	0.2	0.4	0.4	7	10	15.00	62.1764	4.15
1	2	5	4000	0.2	0.4	0.4	17	26	5.00	73.622	8.14
1	4	15	4000	0.2	0.4	0.4	13	19	15.00	105.798	4.64
1	2	5	250	0.15	0.6	0.25	3	6	2.08	23.286	7.15
1	4	15	250	0.15	0.6	0.25	3	5	6.25	44.4807	5.49
1	2	5	1000	0.15	0.6	0.25	4	12	2.08	36.2035	4.64
1	4	15	1000	0.15	0.6	0.25	4	9	6.25	62.5075	3.95
1	2	5	4000	0.15	0.6	0.25	8	24	2.08	63.1557	5.38
1	4	15	4000	0.15	0.6	0.25	7	17	6.25	99.2182	3.77
1	2	5	250	0.15	0.25	0.6	6	6	12.00	19.2791	13.50
1	4	15	250	0.15	0.25	0.6	4	5	36.00	29.6771	12.69
1	2	5	1000	0.15	0.25	0.6	9	11	12.00	33.717	13.09
1	4	15	1000	0.15	0.25	0.6	8	9	36.00	47.7994	9.56
1	2	5	4000	0.15	0.25	0.6	18	22	12.00	63.0055	13.63
1	4	15	4000	0.15	0.25	0.6	14	17	36.00	85.1396	8.68

7. Conclusions and future research

We considered a simple two-stage model of a push-pull production-distribution system. Even a simplified variation of this model, with a reduced shipping action space, has an extremely counter-intuitive optimal policy. This policy is very different from the majority of related results found in the literature. For example, there are cases where it is optimal to halt the production at the upstream stage as the number of outstanding orders increases, even though traditional intuition suggests that the production is typically a monotone decision with respect to the number of outstanding orders. We have also shown that the optimal production policy at the upstream stage is not monotone in its own inventory level, even though results from similar models without fixed shipping cost indicate otherwise. The optimal shipping decision also exhibits highly counter-intuitive behavior, as the benefit of immediate versus delayed shipping varies non-monotonically over states and parameters. Based on some structural results and extensive numerical testing, we developed a heuristic for our model, and computationally tested the performance of the heuristic. The results show that our heuristic is quite robust with respect to the changes in the maximum capacity of shipping (i.e., the size of the maximum shipment). Although we have only considered the case of a zero shipping time, we suspect that the qualitative structure of the optimal policy will be preserved even in the case of a random or deterministic shipping time.

On the other hand, the model considered in this paper has many limitations and restrictions, some of which we hope to relax in future research. For example, we only consider a two-stage system, rather than a multi-stage system which is more common in practice. We assume that production

and interarrival times are exponentially distributed, when other distributions are frequently more appropriate in practice. Also, we assume a fixed cost of transportation, when in many cases transportation costs are better captured by more complicated non-linear functions. Nevertheless, it is encouraging to see that this relatively complex two-stage problem is amenable to analysis, and that simple but effective heuristics can be developed. We hope to build on this insight in future research as we alter our models to address some of these issues.

Acknowledgements

We would like to thank the Associate Editor and the anonymous referees for their valuable comments. This research was partially supported by NSF grant DMI-0200439.

References

- Ahn, H.-S., Duenyas, I. and Zhang, R. (1999) Optimal stochastic scheduling of a 2-stage tandem queue with parallel servers. *Advances in Applied Probability*, **38**, 1095–1117.
- Carr, S. and Duenyas, I. (2000) Optimal admission control and sequencing in a make-to-stock/make-to-order production system. *Operations Research*, **48**, 709–720.
- Chen, F. (1998) Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Science*, **44**(12), S221–S234.
- Chen, F. and Song, J.-S. (2001) Optimal policies for multi-echelon inventory problems with Markov modulated demand. *Operations Research*, **49**, 226–234.
- Chen, F. and Zheng, Y.-S. (1994) Evaluating echelon stock (R, nQ) policies in serial production/inventory systems with stochastic demand. *Management Science*, **40**, 1262–1275.

- Clark, A.J. and Scarf, H. (1960) Optimal policies for a multi-echelon inventory problem. *Management Science*, **6**, 475–490.
- Duenyas, I. and Patane-Anake, P. (1998) Base-stock control for single product tandem make-to-stock systems. *IIE Transactions*, **30**, 31–39.
- Duenyas, I. and Tsai, C. (2002) Centralized and decentralized control of a two-stage tandem manufacturing system with demand for intermediate and end products. Working paper, The University of Michigan, Ann Arbor, MI 48109, USA.
- Duenyas, I., Frank, K. and Zhang, R. (2003) Optimal inventory policies in systems with priority demand classes. *Operations Research*, **51**, 993–1002.
- Federgruen, A. and Zipkin, P. (1984a) Computational issues in an infinite-horizon, multi-echelon inventory model. *Operations Research*, **32**, 818–836.
- Federgruen, A. and Zipkin, P. (1984b) An efficient algorithm for computing optimal (s,S) policies. *Operations Research*, **32**, 1268–1285.
- Federgruen, A. and Zipkin, P. (1986a) An inventory model with limited production capacity and uncertain demands I: the average-cost criterion. *Mathematics of Operations Research*, **11**, 193–207.
- Federgruen, A. and Zipkin, P. (1986b) An inventory model with limited production capacity and uncertain demands II: The discounted-cost criterion. *Mathematics of Operations Research*, **11**, 208–215.
- Glasserman, P. and Tayur, S. (1994) The stability of a capacitated, multi-echelon production-inventory system. *Operations Research*, **42**, 913–925.
- Glasserman, P. and Tayur, S. (1995) Sensitivity analysis for base-stock levels in multi-echelon production-inventory system. *Management Science*, **41**, 263–281.
- Glasserman, P. and Tayur, S. (1996) A simple approximation for multi-stage capacitated production-inventory system. *Naval Research Logistics*, **43**, 41–58.
- Ha, A. (1997) Optimal dynamic scheduling policy for a make-to-stock production system. *Operations Research*, **43**, 1093–1103.
- Iravani, S.M.R., Posner, M.J.M. and Buzacott, J.A. (2003) Scheduling of a moving server in a serial production process. Working paper. Northwestern University, Evanston, IL, 60208.
- Kaminsky, P. and Simchi-Levi, D. (2003) Production and distribution lot sizing in a two-stage supply chain. *IIE Transactions*, **35**, 1065–1074.
- Kapuscinski, R. and Tayur, S. (1998) A capacitated production-inventory model with periodic demand. *Operations Research*, **46**, 899–911.
- Lambert, D. M. and Stock, J.R. (1993) *Strategic Logistics Management*, McGraw Hill, New York, NY.
- Lippman, S.A. (1975) Applying a new device in the optimization of exponential queuing systems. *Operations Research*, **23**, 686–710.
- Meyn, S. (1997) The policy improvement algorithm for Markov decision processes. *IEEE Transactions on Automatic Control*, **42**(12), 1663–1680.
- Parker, R. and Kapuscinski, R. (2004) Optimal inventory policies for a capacitated two-echelon system. *Operations Research* (to appear).
- Simchi-Levi, D., Kaminsky, P. and Simchi-Levi, E. (2003) *Designing and Managing the Supply Chain, 2nd edn.*, McGraw-Hill, New York, NY.

Biographies

Hyun-soo Ahn is an Assistant Professor of Operations and Management Science at the Stephen M. Ross School of Business at the University of Michigan. His research interests include production and inventory systems, dynamic pricing and revenue management, workforce agility, and resource allocation under uncertainty. He received a BA in Industrial Engineering from KAIST and a Ph.D. in Industrial and Operations Engineering from the University of Michigan.

Philip Kaminsky is an Associate Professor in the Industrial Engineering and Operations Research Department at the University of California, Berkeley. He received his Ph.D. in Industrial Engineering and Management Science from Northwestern University in 1997. Prior to that, he worked in production engineering and control at Merck and Co. His current research focuses on the analysis and development of robust and efficient techniques for the design and operation of logistics systems and supply chains. He is a co-author of *Designing and Managing the Supply Chain: Concepts, Strategies and Case Studies*, published by McGraw Hill in 1999, winner of the Book-of-the-Year Award and Outstanding IIE Publication Award given in 2000. He has consulted in the areas of production and logistics system control.

Contributed by the Engineering Statistics and Applied Probability Department.