European Journal of Operational Research 244 (2015) 490–497

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

An approximate policy for a dual-sourcing inventory model with positive lead times and binomial yield

Wanrong Ju[a,*], Adriana F. Gabor[a], J.C.W. van Ommeren[b]

[a] Department of Econometrics, Erasmus University Rotterdam, Burgmeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands
[b] Faculty of Mathematical Sciences, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

ARTICLE INFO

Article history:
Received 14 January 2014
Accepted 26 January 2015
Available online 31 January 2015

Keywords:
Inventory
Supply chain management
Applied probability
Yield uncertainty
Dual-index order-up-to policy

ABSTRACT

This paper studies the inventory system of a retailer who orders his products from two supply sources, a local one that is responsive and reliable, but expensive, and a global one that is low-cost but less reliable. The deliveries from the global source only partially satisfy the quality requirements. We model this situation with a dual-sourcing inventory model with positive lead times and random yield. We propose a dual-index order-up-to policy (DOP) based on approximating the inventory model with an unreliable supplier by a sequence of dual-sourcing models with reliable suppliers and suitably modified demand distributions. Numerical results show that the performance of this heuristic is close to that of the optimal DOP. Moreover, we extend the heuristic to models with advance yield information and study its impact on the total inventory costs.

1. Introduction

Rising in tandem with the prevalence of outsourcing activities, supply risk has recently attracted a great deal of attention from the OR research community. One important type of risk in outsourcing processes is the uncertainty regarding the order quantities that turn out to be usable at the buyer companies. This uncertainty is often referred to in the literature as yield uncertainty. Many factors may lead to yield uncertainty. When goods are transported from a global supplier, yield uncertainty is often related to damage that occurs during transportation due to humidity, collision and other reasons. Part of the goods received may also fail to pass the quality inspection of the buyers. For example, in the semiconductor industry, the yield rate may drop below 50 percent due to strict requirements on quality (Grasman, Sari, & Sari, 2007).

Yield uncertainty significantly increases the difficulty of inventory management. Often, big OEMs and retailers in Europe and the US, use responsive but more expensive local suppliers to mitigate the yield uncertainty caused by the low-cost global suppliers.

This paper focuses on the inventory system of a retailer who sources from two suppliers: one global and one local. Both suppliers have positive lead times. The global supplier is unreliable in the sense that his deliveries only partially satisfy the quality requirements. Thus, the goods need to pass quality inspection in order to be identified as usable items. Due to delays in quality inspection, the usable portion of an order is not known when the order is placed. In this paper we assume that failures of different units in an order are independent and that the failure probability is the same for all units. We are interested in the sourcing strategy that minimizes the average total inventory costs and in the impact of advance information about the quantity of the usable items on this strategy. To the best of our knowledge, this model has not been previously discussed in the OR literature.

Since the optimal policy for the simpler dual-sourcing model with positive lead times and full returns does not have a simple form (Whittemore & Saunders, 1977), we focus on finding a simple and efficient heuristic for our model. According to Veeraraghavan and Scheller-Wolf (2008), the dual-index order-up-to policy (DOP) performs close to the optimal policy in a dual-sourcing model with full returns. We therefore propose a DOP, establishing the order-up-to levels by using a sequence of dual-sourcing models with reliable suppliers and suitably modified demand distributions. We then show how to extend this heuristic to incorporate advanced information on the yield quantities.

1.1. Related literature

Yield uncertainty has attracted a great deal of attention in inventory management research in the past several decades. Three types of random yield have been considered in the literature: binomial yield (Inderfurth & Vogelgesang, 2013), stochastically proportional yield (Agrawal & Nahmias, 1997; Bollapragada & Morton, 1999; Henig & Gerchak, 1990; Huh & Nagarajan, 2010; Inderfurth & Transchel, 2007; Inderfurth & Vogelgesang, 2013; Li, Xu, & Zheng, 2008) and

* Corresponding author. Tel.: +31 646612983.
E-mail addresses: ju@ese.eur.nl, wodeaiweier@gmail.com (W. Ju), gabor@ese.eur.nl (A. F. Gabor), J.C.W.vanOmmeren@math.utwente.nl (J.C.W. van Ommeren).

http://dx.doi.org/10.1016/j.ejor.2015.01.052
0377-2217/

interrupted geometric yield (Inderfurth & Vogelgesang, 2013). Binomial yield is used when failures of different units in a batch are independent and occur with the same probability. In situations where a random process affects whole batches, stochastically proportional yield is used instead. Models using interrupted geometric yield assume that good items are generated independently with a fixed probability until a failure occurs, and that thereafter all items are defective.

Most papers consider the effect of random yield in single-sourcing models with zero lead time. Henig and Gerchak (1990) studied the optimal policy and showed that it has an order point, but does not have the order-up-to structure. Bollapragada and Morton (1999) and Inderfurth and Transchel (2007) proved that the infinite-horizon periodic-review model can be reduced to a newsvendor problem. However, a closed-form solution cannot easily be found, since the 'demand distribution' in the newsvendor problem depends on the order quantities. Li et al. (2008) gave upper and lower bounds for the optimal reorder point and order quantity in an infinite-horizon model. Cheong and Song (2013) studied the value of the yield information in a newsvendor problem with stochastically proportional yield. They compared the ordering decisions and overall profits when different levels of yield information are available (i.e. no information at all, known expectation and variance of the yield factor and known distribution of the yield factor).

Since the optimal policy is difficult to find even when the lead time is neglected, numerous heuristics have been proposed. Bollapragada and Morton (1999) studied several myopic heuristics. Huh and Nagarajan (2010) found the optimal policy within the class of 'linear inflation rules' and proved that the average total cost is convex in the order-up-to level for any given inflation factor.

For positive lead times, Inderfurth and Vogelgesang (2013) proposed a linear inflation rule for which the optimal critical stock level is derived based on a normal approximation of the difference between the lead time demand and the yield of the pipeline orders. Inderfurth and Kiesmüller (2013) proposed two methods to derive optimal or near-optimal critical stock levels: one method based on modeling the on-hand inventory by a Markov chain and the other method based on fitting a normal or gamma distribution to the on-hand inventory.

Dual sourcing is often used in practice for balancing cost and service level or for mitigating yield uncertainty. Due to the complexity of the problem, however, the literature focusses on models with full returns. Whittemore and Saunders (1977) proved that when the difference between the lead times is larger than one, the optimal policy does not have a simple structure. For general lead times, Veeraraghavan and Scheller-Wolf (2008) proposed a dual-index order-up-to policy (DOP), and showed that the DOP performs close to the optimal policy. The main difficulty in finding the optimal order-up-to levels in a DOP is that, due to different lead times, the expedited inventory position may exceed its order-up-to level. The distribution of this excess, called the overshoot, is difficult to find. Veeraraghavan and Scheller-Wolf (2008) showed that for any given difference between the order-up-to levels, if the distribution of the overshoot were known, the optimal expedited order-up-to level could be found by solving a newsvendor problem. In order to find the distribution of the overshoot, they relied on simulations. Arts, van Vuuren, and Kiesmüller (2011) gave an approximation of the distribution of the overshoot, which is exact when the difference between the order-up-to levels is one or approaches infinity. Sheopuri, Janakiraman, and Seshadri (2010) generalized the DOP and studied three new policies which outperform the optimal DOP in their numerical experiments. Tagaras and Vlachos (2001) considered a heuristic which uses the order-up-to rule for the regular supplier and places orders with the expedited supplier only when the likelihood of a stockout is very high. Allon and Van Mieghem (2010) studied a continuous review model and proposed a tailored base-surge policy (TBS). This policy sources from the cheap, offshore supplier at a constant rate (to meet a *base* level of demand) and from the responsive, nearshore supplier only when on-hand inventory is below a certain level (to manage demand *surges*). They presented bounds on the optimal cost and an asymptotically optimal policy for a high-volume system. Janakiraman, Seshadri, and Sheopuri (2014) proved that the TBS policy is optimal for periodic review inventory systems in which demand follows a two-point distribution and the probability of a high demand is sufficiently small. Moreover, if demand can be represented as a sum of two random variables, the base and the surge demand, and the last one occurs with a small probability, the TBS policy performs close to the optimal. Chen, Feng, and Seshadri (2013) incorporated price-dependent demand in multi-period models with more than one unreliable suppliers and negligible lead times. They proved that for general demand models, there exists a time-dependent reorder point for each supplier such that a positive order is placed for almost every inventory level below the reorder point. They also studied sufficient conditions under which the optimal policies have a strict reorder point structure or the optimal order quantities are decreasing in the inventory level.

Other papers related to ours studied the optimal policies for the single-sourcing inventory systems in which the realized demand is known to the decision maker only after a certain number of periods (Bensoussan, Çakanyildirim, & Sethi, 2006; Bensoussan, Cakanyildirim, & Sethi, 2007a; Bensoussan, Çakanyildirim, & Sethi, 2007b; Bensoussan, Çakanyildirim, Feng, & Sethi, 2009; Bensoussan, Cakanyildirim, Sethi, Wang, & Zhang, 2011). In contrast, we focus on dual-sourcing models in which the usable quantities of the orders placed with the offshore supplier are not immediately known when the orders are placed.

1.2. Statement of contribution

We consider yield uncertainty in a dual-sourcing model with positive lead times, which to the best of our knowledge, has not been studied previously in the literature. For this model we propose a dual-index order-up-to heuristic. In order to account for yield uncertainty, we find the order-up-to levels based on a sequence of dual-sourcing models with full returns and suitably modified demand distributions. When compared to the optimal DOP, our heuristic gives promising results. We further extend our heuristic to models with advance yield information, and study its impact on the total costs.

The remainder of the paper is organized as follows. Section 2 formulates the dual-sourcing model with positive lead times and yield uncertainty, and gives some preliminary results on the DOP. Section 3 describes in detail our heuristic for the case where the yield becomes known only at the moment of delivery. Section 4 extends the heuristic to models with advanced information regarding the yield quantities. Section 5 presents numerical results on the performance of the proposed heuristic and on the importance of taking into account the yield uncertainty when designing inventory policies. This section also examines the impact of the advance yield information on the total inventory costs. Section 6 contains a summary of our results and concluding remarks.

2. Model and preliminaries on dual-index order-up-to policies

We consider an infinite-horizon periodic-review inventory model with two suppliers, one regular (denoted as 'r') and one expedited supplier (denoted as 'e'). The lead time l_r of the regular supplier is larger than the lead time l_e of the expedited supplier, while the per-unit ordering cost c_r of the regular supplier is lower than the cost c_e of the expedited one. Assume that there is no fixed ordering cost for either supplier and that c_r is paid for every ordered unit from the regular supplier. Moreover, the regular supplier has binomial random yield, which means that, out of an order X_n^r placed with him, only a random portion $B(X_n^r, p)$ turns out to be usable upon delivery, where p is the long-run average fraction of usable items. In Sections 2 and 3 we assume that the yield quantities are known only when the orders physically arrive. In Section 4 we loosen this assumption and allow advance yield information. Demand in different periods, denoted as D_n, $n = 1, 2, \ldots$, is considered to be independent and identically distributed, with $E(D) < \infty$. Revealed demand is fulfilled from on-hand inventory and unsatisfied demand is fully backlogged. Backlogged demand is charged a penalty cost b per unit per period, and inventory carried to the next period is charged a holding cost h per unit.

The sequence of events in each period n is as follows:

- The physical arrival of both the expedited order placed l_e periods in the past (i.e. $X^e_{n-l_e}$) and a binomial portion of the regular order placed l_r periods in the past (i.e. $X^r_{n-l_r}$).
- Placement of new orders with the suppliers according to the inventory policy applied.
- Realization of demand of this period, which is then either fulfilled or backlogged.
- On-hand inventory is charged a holding cost and backlogged demand is charged a penalty cost.

We are interested in finding an efficient dual-index order-up-to policy (DOP) that minimizes the long-run average total cost given by $\lim_{N\mapsto\infty} \frac{1}{N}\sum_{n=1}^{N} TC_n$, with

$$TC_n = c_e X^e_n + c_r X^r_n + hI^+_n + bI^-_n,$$

where $a^+ = \max(a, 0)$ and $a^- = \max(-a, 0)$.

A DOP is characterized by two order-up-to levels: one for the expedited supplier, z_e, and one for the regular supplier, z_r. In each period $n \geq l_r - 1$, there are $l_r - 1$ regular and $l_e - 1$ expedited orders in the pipeline, denoted by $\langle X^r_{n-l_r+1}, \ldots, X^r_{n-1}\rangle$ and $\langle X^e_{n-l_e+1}, \ldots, X^e_{n-1}\rangle$, respectively.

We define the expedited inventory position in period n, IP^e_n, as the sum of the net inventory level when deciding the order quantities and all of the orders due to arrive in the next l_e periods. We define the regular inventory position in period n, IP^r_n, as the sum of the net inventory level when deciding the order quantities and all of the orders due to arrive in the next l_r periods. More precisely,

$$IP^e_n = I^b_n + (X^e_{n-l_e+1} + \cdots + X^e_{n-1}) + (X^r_{n-l_r+1} + \cdots + X^r_{n-l-1}) \tag{1}$$

$$IP^r_n = I^b_n + (X^e_{n-l_e+1} + \cdots + X^e_{n-1}) + (X^r_{n-l_r+1} + \cdots + X^r_{n-1}), \tag{2}$$

where $l = l_r - l_e$ and I^b_n is the net-inventory level after orders arrive and before demand is realized in period n. Clearly, $I^b_n - D_n = I_n$.

Note that (1) and (2) express the exact inventory positions in the case with full returns. When there is yield uncertainty, however, the exact order quantities that will turn out to be usable in the next l_r periods are unknown. In this case, the inventory positions defined above overestimate the real ones. In order to compensate for this, we will adjust in our heuristic the demand in each period.

When the DOP is applied, in each period n, first an expedited order X^e_n is placed in order to restore the inventory position IP^e_n to the target level z_e. Observe that when X^e_n is decided, X^r_{n-l} enters the information horizon. Thus, one first checks if there is a surplus (i.e. whether $IP^e_n + X^r_{n-l} > z_e$). If this is the case, no expedited order is placed. Otherwise, an expedited order equal to the deficit $X^e_n = z_e - (IP^e_n + X^r_{n-l})$ is placed. Then the expedited order X^e_n is added to the regular inventory position and a regular order $X^r_n = z_r - (IP^r_n + X^e_n)$ is placed.

Similar to the literature on dual-sourcing models with full returns, we call the quantity $O_n = (IP^e_n + X^r_{n-l} - z_e)^+$ the *overshoot*. The overshoot and the inventory positions of the regular and expedited supplier satisfy the following equations:

$$IP^e_n + X^r_{n-l} + X^e_n = z_e + O_n \tag{3}$$

$$IP^r_n + X^e_n + X^r_n = z_r. \tag{4}$$

Table 1
Notations and descriptions.

	Descriptions		Descriptions
n	Period index	D_n	Demand of period n
I_n	Net inventory level at the end of period n	I^b_n	Net inventory level after orders arrive and before demand is realized in period n
c_e	Per-unit ordering cost of the expedited supplier	c_r	Per-unit ordering cost of the regular supplier
l_e	Lead time of the expedited supplier	l_r	Lead time of the regular supplier
h	Holding cost per unit per period	b	Penalty cost per unit per period
X^e_n	Expedited order placed in period n	X^r_n	Regular order placed in period n
z_e	Expedited order-up-to level	z_r	Regular order-up-to level
Δ	Difference between the order-up-to levels	p	Success rate of the binomial yield distribution
f_U	Probability density function of random variable U	F_U	Cumulative distribution function of random variable U
$\hat{U}$	Probability generating function of U		

From (3) and (4), we obtain

$$\sum_{k=0}^{l-1} X^r_{n-k} + O_n = \Delta, \tag{5}$$

where $\Delta = z_r - z_e$.

The following result can be proven for a model with yield uncertainty.

Proposition 1. *The distribution of the overshoot is a function of Δ, independent of z_e.*

The proof is similar to the proof of Lemma 5.1 in Veeraraghavan and Scheller-Wolf (2008) and is therefore omitted.

The optimal DOP for the model with yield uncertainty can be found by formulating the problem as a Markov decision process. However, since a state contains all of the pipeline information, the optimization problem becomes intractable for large l_r.

Notations used in this paper are summarized in Table 1.

3. A dual-index order-up-to policy with modified demand (the DOPMD heuristic)

We propose a DOP in which the order-up-to levels are found based on a sequence of dual-sourcing models with full returns and modified demand distributions. We call this heuristic the DOPMD heuristic. The main steps of the DOPMD can be described as follows:

1. For each Δ, find the optimal $z_e(\Delta)$ and $z_r(\Delta)$ in an approximate dual-sourcing model with full returns and modified demand.
2. Find the pair $(z_e(\Delta), z_r(\Delta))$ with the lowest total costs.

Next we explain in more detail Step 1 of the heuristic.

3.1. Approximate dual-sourcing model

Note that in a dual-sourcing model with uncertain returns, the following recursion holds:

$$\begin{aligned} I_n &= I_{n-1} + X^e_{n-l_e} + B(X^r_{n-l_r}, p) - D_n \\ &= I_{n-1} + X^e_{n-l_e} + X^r_{n-l_r} - (D_n + B(X^r_{n-l_r}, 1-p)). \end{aligned}$$

If the variables $D_n + B(X^r_{n-l_r}, 1-p)$ were independent and their distribution was easy to calculate, the order-up-to levels could be found based on a model with *full returns* and demand in each period defined as $D'_n = D_n + B(X^r_{n-l_r}, 1-p)$. However, a regular order depends on the orders placed in the previous l_r periods, thus making it difficult to find the distribution of X^r_n. Therefore, we propose approximating X_n^r based on the order quantities in a *single*-sourcing model with yield uncertainty. Before describing the heuristic in more detail, we first analyze the distribution of the order quantities in a single-sourcing model with yield uncertainty. Let $\hat{D}$ be the probability generating function of the demand.

Lemma 1. *Under the order-up-to policy, in a single-sourcing periodic-review inventory model with binomial random yield, the sequence of orders Q_n, $n = 1, 2, \ldots$ converges almost surely to a random variable Q with probability generating function $\hat{Q}(z) = \prod_{k=0}^{\infty} \hat{D}(q_k z + 1 - q_k)$ where $q_k = (1-p)^k$, and mean $E(Q) = \frac{E(D)}{p}$.*

Proof. When the order-up-to policy is applied, the following relation holds for the order quantities Q_n, $n = 1, 2 \ldots$,

$$Q_{n+1} = B(Q_{n-l_r+1}, 1-p) + D_n. \tag{6}$$

By using iteratively (6), we obtain

$$Q_{n+1} = \sum_{k=0}^{\lfloor \frac{n}{l_r} \rfloor} R_{n,k}, \tag{7}$$

with $R_{n,k} = B(D_{n-kl_r}, q_k)$, where $q_k = (1-p)^k$. Note that since demand in different periods is i.i.d., the distribution of $R_{n,k}$ does not depend on n. For simplicity, we will hereafter omit the index n and refer to $R_{n,k}$ as R_k. We will show that $S_m = \sum_{k=0}^{m} R_k$ converges almost surely, which implies that Q_n converges almost surely.

The probability-generating function $\hat{R}_k$ of R_k is given by $\hat{R}_k(z) = \hat{D}(q_k z + (1-q_k))$. Since

$$\begin{aligned} P(R_{n+1} \geq \frac{1}{n^2}) &= 1 - P(R_{n+1} = 0) \\ &= 1 - \hat{D}(1 - (1-p)^{n+1}) \\ &= (1-p)^{n+1} E(D) + o((1-p)^{n+1}), \end{aligned}$$

$E(D) < \infty$ and $0 < p < 1$, based on Borel Cantelli lemma (Proposition 2.8, Çinlar, 2011), we can conclude that S_n converges almost surely, and thus Q_n converges almost surely to a random variable Q.

The probability-generating function of Q is given by

$$\hat{Q}(z) = \prod_{k=0}^{\infty} \hat{R}_k(z) = \prod_{k=0}^{\infty} \hat{D}(q_k z + 1 - q_k).$$

Using the monotone convergence theorem, we conclude that

$$E(Q) = E(\sum_{k=0}^{\infty} B(D_k, (1-p)^k) = \sum_{k=0}^{\infty} (1-p)^k E(D) = \frac{E(D)}{p}. \tag{8}$$

□

Remark. Since $E(Q_n) = \frac{1-(1-p)^{\lfloor \frac{n}{l_r} \rfloor + 1}}{p} E(D)$, it follows that Q_n also converges in mean to Q.

Let F_∞ be the limiting distribution function defined by $\hat{Q}$ and let $\{U_n, n = 1, 2, \ldots\}$ be a sequence of independent random variables distributed according to F_∞. Note that X_n^r is usually smaller than U_n, since part of the orders in the single-sourcing system is delivered by the expedited supplier in the dual-sourcing case. For a given Δ, we propose approximating X_n^r by $\tilde{X}_n^r$ where

$$\tilde{X}_n^r = B(U_n, \alpha(\Delta)), \text{ for } \alpha(\Delta) \in [0, 1]. \tag{9}$$

Intuitively, according to this approximation, each unit that would be ordered from the regular supplier if he were the only supplier is now ordered with probability $1 - \alpha(\Delta)$ from the expedited supplier. The portion we order from the regular supplier depends in general on Δ, as indicated by Proposition 1 and Eq. (5). For the moment, assume that a suitable $\alpha(\Delta)$ has been found. Then the variables D'_n defined by

$$D'_n = D_n + B(U_n, \alpha(\Delta)(1-p)) \tag{10}$$

are independent and identically distributed. Hence, we can construct a dual-sourcing model with full returns and i.i.d. demand given by (10). We will call this model *the approximate model associated with* Δ.

Observe that in the approximate model we increase the actual demand D_n in each period by a quantity that approximates the unreturned part of the regular order that arrives in the current period. This allows us to compensate for the fact that the real inventory positions are overestimated.

Next we discuss a procedure to find a suitable $\alpha(\Delta)$. According to Eqs. (8) and (9), the following holds: $E(\tilde{X}_n^r) = \alpha(\Delta)E(U_n) = \alpha(\Delta)\frac{E(D)}{p}$. This gives us an approximation of the expected order quantity in the original model, namely, $EXR_{app} = \alpha(\Delta)\frac{E(D)}{p}$. In our heuristic, we modify $\alpha(\Delta)$ iteratively until either EXR_{app} converges to the average regular order quantity derived by simulating a dual-sourcing model with full returns and modified demand D'_n or a certain number of iterations is reached. In order to find a suitable initial value of $\alpha(\Delta)$, we use Eq. (5) to derive that $\sum_{k=0}^{l-1} E(X_{n-k}^r) = \Delta - E(O_n)$. If X_n^r is approximated by $\tilde{X}_n^r$, we get $\frac{\alpha(\Delta) l E(D)}{p} \approx \Delta - E(O_n)$. Since $E(O_n) \geq 0$, we set the initial value of $\alpha(\Delta)$ at $min\{1, \frac{\Delta p}{lE(D)}\}$. The iterative procedure is described in detail in Algorithm 1. In our experiments, convergence is always reached in a few iterations.

Algorithm 1

Initialization: $\alpha(\Delta) = min\{1, \frac{\Delta p}{lE(D)}\}$, $EXR_{sim} = 0$, $EXR_{app} = \Delta$,
K=Nr. of iterations, $k = 0$
while $|EXR_{sim} - EXR_{app}| > \epsilon$ and $k \leq K$ **do**
$EXR_{sim} \leftarrow$ the average regular order quantity derived by simulation in a dual-sourcing model with full returns and demand defined as (10)
$EXR_{app} \leftarrow \alpha(\Delta)\frac{E(D)}{p}$
$\alpha(\Delta) \leftarrow min\{1, \frac{pEXR_{sim}}{E(D)}\}$
$k \leftarrow k + 1$
end while

3.2. Finding the order-up-to levels in the approximate dual-sourcing model

According to Veeraraghavan and Scheller-Wolf (2008), in a dual-sourcing model with full returns, for any given Δ, the optimal $z_e^*(\Delta)$ is equal to the newsvendor fractile of the sum of lead-time demand and the stationary overshoot. In our case, the lead-time demand is equal to $\sum_{k=0}^{l} D_{n+k} + \sum_{k=1}^{l} B(U_{n+k-l_r}^r, \alpha(\Delta)(1-p))$. Observe that $B(U_{n-l_r}^r, \alpha(\Delta)(1-p))$ is not part of the lead-time demand, since we place new orders only after the exact value of $B(U_{n-l_r}^r, \alpha(\Delta)(1-p))$ is known. Similar to Veeraraghavan and Scheller-Wolf (2008), the stationary (cumulative) distribution of the overshoot is defined as $P(O \leq x) = \lim_{N \mapsto \infty} \frac{1}{N} \sum_{n=1}^{N} \mathbf{I}(O_n \leq x)$, where $\mathbf{I}$ denotes the indicator function, and is found by using simulation. Note that in order to reduce computation times, the distribution of overshoot could also be approximated as described in Arts et al. (2011). This, however, is not the focus of this paper.

Section 5.1 verifies the performance of the DOPMD by comparing it with the optimal DOP for the given model.

4. Models with advance information on the yield quantities

In the above model, the exact yield quantity of each regular order is assumed to be known only when the order physically arrives. This situation arises when the quality inspection takes place at the ordering company. In this section we examine models with advance yield information, where the exact yield quantities are known before the orders physically arrive. Advance yield information can be obtained by investing in the suppliers' facilities and authorizing them to perform the quality inspections.

Assume that when there is advance yield information, the yield quantity of a regular order X_n^r is known in period $n + l_i$ and $l_i < l_r$. In this case, the recursion of the net-inventory level can be rewritten as

$$I_n = I_{n-1} + X_{n-l_e}^e + Y_{n-l_r}^r - D_n, \tag{11}$$

where $Y_{n-l_r}^r$ is the realized yield quantity of $X_{n-l_r}^r$. Note that with advanced information, the exact yield quantities of the pipeline orders $\{X_{n-l_r+1}^r, \dots, X_{n-l_i}^r\}$ are known. Therefore, in the definition of the inventory positions given by (1) and (2), the regular orders (of which the exact yields are known) are replaced by the yield quantities.

In adapting the DOPMD heuristics to incorporate advance information, we distinguish two cases, namely $l_i \le l$ and $l_i > l$.

4.1. The information lead time is smaller than the difference between the lead times (i.e. $l_i \le l$)

When $l_i \le l$, IP_n^e given by (1) represents the real expedited inventory position. Define O_n as $(IP_n^e + Y_{n-l}^r - z_e)^+$. Based on the recursion of I_n in (11) and the definition of O_n, we get

$$\begin{aligned} I_{n+l_e} &= I_n + \sum_{k=0}^{l_e-1} X_{n-k}^e + \sum_{k=l}^{l_r-1} Y_{n-k}^r - \sum_{k=1}^{l_e} D_{n+k} \\ &= I_n^b + \sum_{k=0}^{l_e-1} X_{n-k}^e + \sum_{k=l}^{l_r-1} Y_{n-k}^r - \sum_{k=0}^{l_e} D_{n+k} \\ &= IP_n^e + X_n^e + Y_{n-l}^r - \sum_{k=0}^{l_e} D_{n+k} \\ &= z_e + O_n - \sum_{k=0}^{l_e} D_{n+k} \end{aligned} \tag{12}$$

Moreover, we can write the recursions of the inventory positions as below:

$$\begin{aligned} IP_{n+1}^e &= IP_n^e + X_n^e + Y_{n-l}^r - D_n = z_e + O_n - D_n \\ IP_{n+1}^r &= IP_n^r + X_n^e + X_n^r - B(X_{n-l_i+1}^r, 1-p) - D_n \\ &= z_r - (D_n + B(X_{n-l_i+1}^r, 1-p)). \end{aligned}$$

It then follows that

$$X_{n+1}^e = (D_n - O_n - Y_{n-l+1}^r)^+ \tag{13}$$

$$X_{n+1}^r = D_n + B(X_{n-l_i+1}^r, 1-p) - X_{n+1}^e \tag{14}$$

Based on Eqs. (13) and (14), we can prove the following proposition.

Proposition 2. *The distributions of X_n^r, X_n^e and O_n are independent of z_e for given Δ.*

The proof follows the same line of reasoning as that of Lemma 5.1 in Veeraraghavan and Scheller-Wolf (2008), and is therefore omitted here.

Using Eq. (12) and Proposition 2, we propose finding the optimal $z_e^*(\Delta)$ in a newsvendor problem as in Veeraraghavan and Scheller-Wolf (2008). More precisely, $z_e^*(\Delta) = F_{D^{(l_e+1)}-O}^{-1}(\frac{b}{b+h})$ where $F_{D^{(l_e+1)}-O}(\cdot)$ is the cumulative distribution function of $\sum_{k=0}^{l_e} D_{n+k} - O_n$. The distribution of O_n needs to be estimated or approximated by one of the methods proposed in dual-sourcing models with full returns. As we did in the previous section, we rely on simulation to estimate this distribution. The pair $(\Delta, z_e^*(\Delta))$ with the lowest cost can now be found by a one-dimensional search.

4.2. The information lead time is larger than the difference between the lead times (i.e. $l_i > l$)

When $l_i > l$, both IP_n^e and IP_n^r overestimate the real inventory positions. Recall that in the DOPMD, the modified demand distributions are found based on the distribution of the order quantities in a single-sourcing model with yield uncertainty. When yield information is available, the order quantities in the single-sourcing model satisfy the following relation:

$$\begin{aligned} Q_{n+1} &= D_n + B(Q_{n-l_i+1}, 1-p) \\ &= \sum_{k=0}^{\lfloor \frac{n}{l_i} \rfloor} B(D_{n-kl_i}, (1-p)^k) \end{aligned}$$

Based on Lemma 1, we conclude that the limiting distribution of Q_n exists and coincides with $F_\infty(\cdot)$ derived in Section 3. Similar to the case without advance yield information, for any given Δ, $\alpha(\Delta)$ is found by using Algorithm 1; then the optimal $z_e^*(\Delta)$ is derived as the newsvendor fractile of the sum of lead-time demand and the stationary overshoot, where the lead-time demand equals $\sum_{k=0}^{l_e} D_{n+k} + \sum_{k=l_r-l_i+1}^{l_e} B(U_{n-l_r+k}, \alpha(\Delta)(1-p))$. Since in the case with advance yield information the exact yield quantities of the regular orders $\{X_{n-l_r+1}^r, \dots, X_{n-l_i}^r\}$ are known, the lead-time demand does not include $\sum_{k=l_i}^{l_r-1} B(U_{n-k}, \alpha(\Delta)(1-p))$ (as it did in the case without yield information). The distribution of the overshoot is found by simulation. The lowest-cost pair $(\Delta, z_e^*(\Delta))$ is found by a one-dimensional search.

As can easily be seen, $\sum_{k=0}^{l_e} D_{n+k} + \sum_{k=l_r-l_i+1}^{l_e} B(U_{n-l_r+k}, \alpha(\Delta)(1-p)) - O_n$ is stochastically smaller than $\sum_{k=0}^{l_e} D_{n+k} + \sum_{k=1}^{l_e} B(U_{n-l_r+k}, \alpha(1-p)) - O_n$. As a consequence, the optimal $z_e^*(\Delta)$ derived here are smaller than their counterparts in the case without advance yield information.

5. Numerical results

This section first presents numerical results on the performance of the proposed DOPMD heuristic. It is shown that the average total costs of the DOPMD are on average only 0.19 percent higher than those of the optimal DOP and in 39 out of the 40 instances tested, the average total costs of the DOPMD are less than 1 percent higher than those of the optimal DOP. This section then goes on to examine the importance of taking into account the yield uncertainty when deciding the inventory policies. Numerical results show that the inventory costs that occur when the DOP derived in a model without yield uncertainty is applied to the model studied here are considerably higher than the costs that occur when the DOPMD is applied. Last but not least, we present results on the impact of the advance yield information. According to our experiments, advance yield information can lead to cost reductions of up to 1.11 percent. The smaller the information lead time, the larger the cost reduction is.

5.1. Performance of the DOPMD heuristic

This section studies the performance of the DOPMD heuristic by comparing it with the optimal DOP for the studied model. The optimal DOP is derived by using the two-dimensional search on both the expedited and the regular order-up-to levels. For each pair of the order-up-to levels, we run the simulation until the 95 percent confidence interval for the average per-period total cost is smaller than 0.1 percent of its expected value. For the DOPMD, the order-up-to levels are found by applying the solution procedure described in Section 3. When deriving the distribution of the overshoot, we run the simulation until either the 95 percent confidence interval for the expected overshoot is smaller than 0.01 or the standard error is less than 0.001 times the expected value for the overshoot. The average total costs corresponding to these order-up-to levels are also derived by simulation. The stopping criterion is the same as that described above for the optimal DOP. One could also derive the average costs from the underlying Markov process; however, since the state space includes information

Table 2
Impact of yield rate ($c_r = 100, b = 495$ and $D \sim Pois(2)$).

l_r	p	Optimal DOP			DOPMD		
		E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent above optimal DOP
2	0.6	0.00	1.99	328.11	0.00	1.99	−0.06
2	0.7	2.18	0.46	320.76	2.53	0.23	0.10
2	0.8	2.42	0.05	286.24	2.43	0.05	0.35
2	0.9	2.18	0.03	257.47	2.22	0.00	0.26
2	1	1.99	0.00	234.13	1.99	0.00	0.25
4	0.6	0.00	2.00	328.30	0.00	1.99	−0.09
4	0.7	1.97	0.62	322.46	1.97	0.62	0.05
4	0.8	2.31	0.14	291.40	2.31	0.14	−0.01
4	0.9	2.17	0.04	264.56	2.13	0.08	−0.12
4	1	1.95	0.04	241.23	1.95	0.04	0.03

Table 3
Impact of regular lead time ($c_r = 100, p = 0.8$ and $D \sim Pois(2)$).

b	l_r	Optimal DOP			DOPMD		
		E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent above optimal DOP
95	2	2.43	0.05	275.78	2.43	0.05	0.80
95	4	2.38	0.08	282.10	2.31	0.14	0.33
95	6	2.29	0.16	285.28	2.29	0.16	0.22
95	8	2.23	0.21	286.65	2.16	0.27	0.20
495	2	2.42	0.05	286.24	2.43	0.05	0.35
495	4	2.31	0.14	291.50	2.31	0.14	−0.01
495	6	2.29	0.16	293.23	2.12	0.30	0.34
495	8	2.21	0.20	292.34	2.08	0.33	1.07

on both regular and expedited orders in transit, the dynamic program becomes computationally intractable. Since we rely on simulation, the average total costs obtained by the heuristic may occasionally be slightly smaller than those obtained by the optimal DOP.

We start with a base case and construct 40 scenarios by modifying one or two of its parameters. In the base case, we choose $l_e = 1, l_r = 2$, $c_r = 100$, $c_e = 150$, $h = 5$, $b = 495$, $p = 0.8$ and $D \sim Pois(2)$, where $Pois(\lambda)$ denotes the Poisson distribution with mean λ. Notice that when demand follows a Poisson distribution, F_∞ is a Poisson distribution with rate $\frac{\lambda}{p}$. All demand distributions are truncated to be between 0 and RD, where RD is the smallest integer that satisfies $F_D(RD) \geq 0.99$. We fix the values of h, c_e and l_e in all instances, and study the respective impact of c_r, p, l_r, b and demand on the performance of the DOPMD heuristic. The results are summarized in Tables 2–6. The layout of each table is as follows: Columns 3 and 4 show the average regular- and expedited-order quantities for the optimal DOP and columns 6 and 7 give those for the DOPMD heuristic; Column 5 gives the average total costs of the optimal DOP and column 8 reports the percentages by which the average total costs of the DOPMD are higher than those of the optimal DOP.

The performance of the DOPMD is robust in all of the cases we tested. The average difference between the average total costs of the DOPMD and of the optimal DOP is only 0.19 percent. In almost all of the cases tested, the difference is smaller than 1 percent.

5.1.1. Impact of yield rate

Table 2 shows that the DOPMD performs well for both high- and low- yield rates. As can be seen from columns 3, 4, 6 and 7, when the yield rate is higher than 0.7, the buyer mainly sources from the regular supplier and uses the expedited supplier only for emergencies. When the yield rate is lower than 0.7, the buyer exclusively sources from the expedited supplier. The reason is that when $p < 0.7$, the actual ordering cost for each usable unit from the regular supplier (i.e. c_r/p) is larger than that from the expedited one, thereby rendering the regular supplier less attractive.

Table 4
Impact of demand size ($l_r = 2, p = 0.8$ and $b = 495$).

c_r	λ	Optimal DOP			DOPMD		
		E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent above optimal DOP
100	2	2.42	0.05	286.24	2.43	0.05	0.35
100	4	4.93	0.00	547.12	4.98	0.00	0.52
100	6	7.46	0.00	809.53	7.31	0.14	0.03
100	8	9.94	0.00	1068.79	9.88	0.09	0.01
110	2	2.30	0.15	310.46	2.42	0.05	−0.03
110	4	4.70	0.22	597.30	4.70	0.00	0.06
110	6	7.27	0.13	879.16	7.30	0.14	0.34
110	8	9.93	0.04	1167.93	9.59	0.09	−0.48
120	2	0.00	2.00	328.57	2.42	0.05	−0.12
120	4	0.00	3.99	636.93	4.70	0.23	0.00
120	6	0.00	5.98	944.31	7.30	0.14	0.04
120	8	4.92	4.04	1251.53	9.59	0.28	0.04

Table 5
Impact of penalty cost ($l_r = 2, p = 0.8$ and $D \sim Pois(2)$).

c_r	Optimal DOP				DOPMD		
	b	E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent above optimal DOP
100	10	2.47	0.01	263.82	2.49	0.00	−0.05
100	15	2.49	0.00	266.01	2.49	0.00	0.14
100	95	2.43	0.05	275.78	2.43	0.05	0.80
100	495	2.42	0.05	286.24	2.43	0.05	0.35
120	10	0.00	1.99	310.27	0.00	1.99	−0.06
120	15	0.00	1.99	312.19	0.00	2.00	0.04
120	95	0.89	1.28	320.96	0.00	1.99	0.24
120	495	0.00	2.00	328.57	0.00	2.00	−0.12

Table 6
Impact of regular ordering cost ($l_r = 2, p = 0.8$ and $D \sim Pois(2)$).

b	c_r	Optimal DOP			DOPMD		
		E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent above optimal DOP
95	100	2.43	0.05	275.78	2.43	0.05	0.80
95	105	2.49	0.00	290.36	2.43	0.05	0.00
95	110	2.30	0.15	301.92	2.30	0.15	0.07
95	115	2.40	0.05	312.65	2.29	0.15	−0.14
95	120	0.00	1.98	319.70	0.00	1.99	0.52
95	125	0.00	1.99	321.27	0.00	2.01	0.59
495	100	2.42	0.05	286.24	2.43	0.05	0.35
495	105	2.47	0.00	296.91	2.42	0.05	0.50
495	110	2.42	0.05	310.16	2.43	0.05	0.12
495	115	2.04	0.36	321.28	2.04	0.36	−0.07
495	120	0.00	1.98	325.95	0.00	1.99	0.61
495	125	0.00	2.00	328.28	0.00	1.99	−0.13

5.1.2. Impact of regular lead time

The results in Table 3 reveal that the performance of the DOPMD is robust for different regular lead times. For both policies, when l_r increases from 2 to 8, the expected expedited order quantities increase. However, in all the tested instances the regular supplier is the main source of supply even when l_r is large.

5.1.3. Impact of demand size

As can be seen in Table 4, the performance of the heuristic is not influenced by the demand size. When the average demand rate increases from 2 to 8, the allocation of orders between the suppliers does not change significantly. When $c_r = 100$ or 110, both policies mainly rely on the regular supplier, and when $c_r = 120$, both policies order from the expedited supplier.

Table 7
Impact of ignoring the yield uncertainty.

c_r	b	DOPMD			DOP ignoring yield uncertainty					
		λ	l_r	p	E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent above DOPMD
100	495	2	2	0.7	2.53	0.22	319.60	2.85	0.00	9.92
100	495	2	2	0.8	2.43	0.05	286.29	2.49	0.00	3.13
100	495	2	2	0.9	2.22	0.00	258.26	2.22	0.00	0.45
100	495	2	2	0.8	2.43	0.05	286.29	2.49	0.00	3.13
100	495	4	2	0.8	4.97	0.01	549.46	4.98	0.00	4.81
100	495	6	2	0.8	7.40	0.06	809.55	7.48	0.00	4.94
100	495	2	2	0.8	2.43	0.05	286.29	2.49	0.00	3.13
100	495	2	4	0.8	2.31	0.14	291.38	2.31	0.14	1.57
100	495	2	6	0.8	2.12	0.30	294.09	2.21	0.22	1.87
80	495	2	2	0.8	2.46	0.01	235.98	2.49	0.00	1.79
100	495	2	2	0.8	2.43	0.05	286.29	2.49	0.00	3.13
120	495	2	2	0.8	0.00	2.00	328.10	2.49	0.00	3.57
100	15	2	2	0.8	2.49	0.00	266.47	2.49	0.00	0.10
100	95	2	2	0.8	2.49	0.00	278.45	2.49	0.00	1.67
100	495	2	2	0.8	2.43	0.05	286.29	2.49	0.00	3.13

5.1.4. *Impact of regular ordering cost and penalty cost*

Tables 5 and 6 show that the DOPMD performs well for different values of the backlogging costs as well as the regular ordering costs. As the backlogging cost increases, both the optimal DOP and the DOPMD slightly increase the portion sourced from the expedited supplier. When the regular ordering cost increases from 90 to 125, both policies change from single sourcing from the regular supplier to dual sourcing, and then to single sourcing from the expedited supplier.

5.2. *Impact of ignoring the yield uncertainty*

This section compares the total costs of the DOPMD and the optimal DOP when the yield uncertainty is ignored. In the latter case, we derive the order-up-to levels by using the solution procedure proposed in Veeraraghavan and Scheller-Wolf (2008). We then apply them in a dual-sourcing model with random yield and derive the average costs by simulation.

Table 7 shows the results of 11 instances with different values of l_r, c_r, λ, b and p. Our experiments suggest that ignoring the yield uncertainty can cause significant increases in costs (on average, 3.07 percent in the tested instances). Moreover, the results in columns 3, 4, 6 and 7 suggest that the cost increase is mainly caused by an inefficient allocation of orders to the suppliers. For example, when $c_r = 120$ and $p = 0.8$ (the fourth to the last row), the optimal DOP when the yield uncertainty is ignored, chooses to single source from the regular supplier since the per-unit ordering cost of the regular supplier is much lower than that of the expedited one. However, when the random yield is taken into account, the expected ordering cost for each usable unit from the regular supplier is 150, which is equal to the expedited ordering cost. Since the regular supplier has a larger lead time, it is inefficient to source from him.

5.3. *Impact of advance yield information*

This section examines the influence of the advance yield information on the total inventory costs. Table 8 shows the cases where $l_i \le l$ and Table 9 those where $l_i > l$.

As can be seen in both tables, the advance yield information leads to cost reduction, which is larger when the yield rates are lower. In our experiment, the average decrease of the total inventory costs is 1.11 percent (1.27 percent for the cases where $l_i \le l$ and 0.95 percent for those where $l_i > l$). When the yield rate is low and $l_i \le l$ (for example, $p = 0.4$ and $l_i \le l$), knowledge of the yield quantities even one period before the orders arrive could lead to 1.65 percent cost reduction (row 3 of Table 8). As expected, the inventory costs decrease when the information lead time decreases. These numerical results are helpful when an assessment must be made as to the value of investing in quality inspection facilities at the supplier's site.

Table 8
Impact of advance yield information when $l_i \le l$ ($b = 495$, $l_r = 4$ and $D \sim Pois(2)$).

p	l_i	c_r	DOPMD			DOP with advance information		
			E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent lower DOPMD
0.4	1	50	4.45	0.21	298.29	4.53	0.18	2.23
0.4	2	50	4.45	0.21	298.29	4.36	0.25	2.02
0.4	3	50	4.45	0.21	298.29	4.38	0.24	1.65
0.6	1	80	2.68	0.38	308.97	2.99	0.20	1.29
0.6	2	80	2.68	0.38	308.97	2.93	0.22	1.80
0.6	3	80	2.68	0.38	308.97	2.89	0.26	0.95
0.8	1	100	2.31	0.14	291.39	2.22	0.21	0.66
0.8	2	100	2.31	0.14	291.39	2.28	0.17	0.42
0.8	3	100	2.31	0.14	291.39	2.22	0.22	0.40

Table 9
Impact of advance yield information when $l_i > l$ ($b = 495$, $l_e = 4$, $l_r = 5$ and $D \sim Pois(2)$).

p	l_i	c_r	DOPMD			DOP with advance information		
			E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	Percent lower DOPMD
0.4	2	55	3.30	0.67	336.71	4.30	0.28	2.21
0.4	3	55	3.30	0.67	336.71	4.59	0.15	1.53
0.4	4	55	3.30	0.67	336.71	4.30	0.28	0.49
0.6	2	85	2.33	0.60	338.41	3.06	0.16	1.30
0.6	3	85	2.33	0.60	338.41	2.78	0.33	0.78
0.6	4	85	2.33	0.60	338.41	3.04	0.15	0.79
0.8	2	115	2.30	0.15	338.47	2.30	0.15	0.74
0.8	3	115	2.30	0.15	338.47	2.04	0.36	0.47
0.8	4	115	2.30	0.15	338.47	2.04	0.36	0.24

6. Conclusions

This paper studies a dual-sourcing inventory system with positive lead times and binomial random yield for the slower supplier. Although commonly seen in practice, this system has never been studied in the literature due to its complexity. We strive to be the first to propose an efficient and easy-to-implement policy for this inventory system. The heuristic we proposed (i.e. the DOPMD) has the structure of a dual-index order-up-to policy. To find the order-up-to levels of the DOPMD, we propose using the full returns in the inventory positions and increasing demand by a quantity that approximates the unreturned regular order. The regular order quantities are approximated by binomial portions of the order quantities in a single-sourcing model with yield uncertainty. The success rates of the binomial random variables are found by using an iterative algorithm. In all of our numerical experiments, the average difference between the costs of the DOPMD heuristic and the optimal DOP is below 1 percent. Moreover, we extend our heuristic to models with advance yield information and study the impact of the advance yield information on the total inventory costs.

References

Agrawal, N., & Nahmias, S. (1997). Rationalization of the supplier base in the presence of yield uncertainty. *Production and Operations Management*, *6*(3), 291–308.

Allon, G., & Van Mieghem, J. (2010). Global dual sourcing: Tailored base-surge allocation to near- and offshore production. *Management Science*, *56*(1), 110–124.

Arts, J., van Vuuren, M., & Kiesmüller, G. (2011). Efficient optimization of the dual-index policy using Markov chains. *IIE Transactions*, *43*(8), 604–620.

Bensoussan, A., Çakanyildirim, M., & Sethi, S. P. (2006). Optimality of base-stock and (s, s) policies for inventory problems with information delays. *Journal of optimization theory and applications*, *130*(2), 153–172.

Bensoussan, A., Cakanyildirim, M., & Sethi, S. P. (2007a). Economic evaluation of systems that expedite inventory information. *Production and Operations Management*, *16*(3), 360–368.

Bensoussan, A., Çakanyildirim, M., & Sethi, S. P. (2007b). Optimal ordering policies for inventory problems with dynamic information delays. *Production and Operations Management*, *16*(2), 241–256.

Bensoussan, A., Çakanyildirim, M., Feng, Q., & Sethi, S. P. (2009). Optimal ordering policies for stochastic inventory problems with observed information delays. *Production and Operations Management*, *18*(5), 546–559.

Bensoussan, A., Cakanyildirim, M., Sethi, S. P., Wang, M., & Zhang, H. (2011). Average cost optimality in inventory models with dynamic information delays. *Automatic Control, IEEE Transactions on*, *56*(12), 2869–2882.

Bollapragada, S., & Morton, T. E. (1999). Myopic heuristics for the random yield problem. *Operations Research*, *47*(5), 713–722.

Çinlar, E. (2011). *Probability and stochastics*. Springer.

Chen, W., Feng, Q., & Seshadri, S. (2013). Sourcing from suppliers with random yield for price-dependent demand. *Annals of Operations Research*, *208*(1), 557–579.

Cheong, T., & Song, S. H. (2013). The value of information on supply risk under random yields. *Transportation Research Part E*, *60*, 27–38.

Grasman, S. E., Sari, Z., & Sari, T. (2007). Newsvendor solutions with general random yield distributions. *RAIRO-Operations Research*, *41*(04), 455–464.

Henig, M., & Gerchak, Y. (1990). The structure of periodic review policies in the presence of random yield. *Operations Research*, *38*(4), 634–643.

Huh, W. T., & Nagarajan, M. (2010). Linear inflation rules for the random yield problem: Analysis and computations. *Operations Research*, *58*(1), 244–251.

Inderfurth, K., & Kiesmüller, G. P. (2013). *Exact and heuristic linear-inflation policies for an inventory model with random yield and arbitrary lead times* (Working Paper No. 7). Otto van Guericke Universität Magdeburg.

Inderfurth, K., & Transchel, S. (2007). Technical note on 'Myopic heuristics for the random yield problem'. *Operations Research*, *55*(6), 1183–1186.

Inderfurth, K., & Vogelgesang, S. (2013). Concepts for safety stock determination under stochastic demand and different types of random production yield. *European Journal of Operational Research*, *224*(2), 293–301.

Janakiraman, G., Seshadri, S., & Sheopuri, A. (2014). Analysis of tailored base-surge policies in dual sourcing inventory systems. *Management Science*, doi: 10.1287/mnsc.2014.1971

Li, Q., Xu, H., & Zheng, S. (2008). Periodic-review inventory systems with random yield and demand: Bounds and heuristics. *IIE Transactions*, *40*(4), 434–444.

Sheopuri, A., Janakiraman, G., & Seshadri, S. (2010). New policies for the stochastic inventory control problem with two supply sources. *Operations Research*, *58*(3), 734–745.

Tagaras, G., & Vlachos, D. (2001). A periodic review inventory system with emergency replenishments. *Management Science*, *47*(3), 415–429.

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

Whittemore, A. S., & Saunders, S. (1977). Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics*, *32*(2), 293–305.