Approximate Order-up-to Policies for Inventory Systems with Binomial Yield

Wanrong Ju, Adriana F. Gabor

Department of Econometrics, Erasmus University Rotterdam
Burgmeester Oudlaan 50, 3062 PA Rotterdam, the Netherlands

J.C.W. van Ommeren

Faculty of Mathematical Sciences, University of Twente
P.O.Box 217, 7500 AE Enschede, the Netherlands

Abstract

This paper studies an inventory policy for a retailer who orders his products from a supplier whose deliveries only partially satisfy the quality requirements. We model this situation by an infinite-horizon periodic-review model with binomial random yield and positive lead time. We propose an order-up-to policy based on approximating the inventory model with unreliable supplier by a model with a reliable supplier and suitably modified demand distribution. The performance of the order-up-to policy is verified by comparing it with both the optimal policy and the safety stock policy proposed in Inderfurth & Vogelgesang (2013). Further, we extend our approximation to a dual-sourcing model with two suppliers: the first slow and unreliable, and the other fast and fully reliable. Compared to the dual-index order-up-to policy for the model with full information on the yield, the proposed

Email addresses: `ju@ese.eur.nl` (Wanrong Ju), `gabor@ese.eur.nl` (Adriana F. Gabor), `J.C.W.vanOmmeren@math.utwente.nl` (J.C.W. van Ommeren)

Preprint submitted to European Journal of Operational Research *January 13, 2014*

approximation gives promising results.

Keywords: Inventory, Yield uncertainty, Order-up-to policy

1. Introduction

Rising 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 as yield uncertainty in the literature. Many factors may lead to yield uncertainty. When goods are transported from a global supplier or the transported goods are delicate parts, 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 the quality inspection of the buyers. For example, in the semiconductor industry, the yield rate may drop below 50% due to strict requirements on quality (Grasman et al., 2007). Yield uncertainty is also encountered in industries where production is influenced by exogenous factors, like weather and diseases. Kazaz (2004) reports that in agriculture, the yield rate can be as low as 30%.

Yield uncertainty significantly increases the difficulty of inventory management. Numerous papers have studied optimal or heuristic policies for inventory systems with uncertain yield. However, few have taken into account the effect of lead time. Lead time refers to the timespan between the moment an order is placed by the buyer and the moment when the ordered goods are delivered. Consisting of the order processing time, production time and transportation time, this period may sometimes be as long as several months. In practice, lead times can hardly afford to be neglected, especially in the case of global sourcing. This paper studies the inventory system of a retailer with positive lead time and yield uncertainty. The retailer has a global supplier whose deliveries only partially satisfy the quality requirements. We study the case in which failure of different units in an order is uncorrelated and each unit has the same probability of failing. This is often the situation if the uncertain yield is caused by damage during transportation or failure at quality inspection. The retailer checks his inventory level periodically and decides on the quantity to order based on his inventory control policy. Unsatisfied demand is fully backlogged. The number of usable units in an order becomes known only when the order physically arrives at the retailer. The total inventory costs of the retailer consist of the holding cost, penalty (backlogging) cost and ordering cost. Inventory holding costs are incurred for the items in inventory at the end of a period. On the other hand, penalty costs are incurred when there is not enough inventory to satisfy customer demand. For this model, we propose a simple order-up-to policy (OP) based on the optimal policy in an approximate model with a modified demand distribution and a reliable supplier. We call this 'the OPMD heuristic'. We then consider the case where the risk posed by the unreliable supplier is mitigated by ordering a part of the units from a more expensive and reliable supplier. To the best of our knowledge, this model has not been previously discussed in the OR literature. For this model, we propose a dual-index order-up-to policy (DOP) based on an approximate model with two reliable suppliers and modified demand distribution (called 'the DOPMD heuristic').

The remainder of the paper is organized as follows. Section 2 briefly reviews the related literature. Section 3 formulates the single-sourcing model with positive lead time and yield uncertainty. Subsequently, we propose a simple order-up-to heuristic and derive the optimal order-up-to level based on a reduction to a model with full returns. An extension of our heuristic to a dual-sourcing model with general lead times and yield uncertainty is presented in Section 4. Section 5 presents numerical results on the performance of the proposed heuristics. For the single-sourcing model, we compare our heuristic with the optimal policy and a recently proposed heuristic (Inderfurth & Vogelgesang, 2013). In the case of dual-sourcing, we compare the proposed heuristic and the optimal dual-index order-up-to policy for the studied model.

2. Literature Review

Yield uncertainty has drawn extensive attention in inventory management research in the past several decades. There are three types of random yield that have been considered in the literature: binomial yield (Inderfurth & Vogelgesang, 2013), stochastically proportional yield (Henig & Gerchak, 1990; Agrawal & Nahmias, 1997; Bollapragada & Morton, 1999; Inderfurth & Transchel, 2007; Li et al., 2008; Huh & Nagarajan, 2010; Inderfurth & Vogelgesang, 2013) and interrupted geometric yield (Inderfurth & Vogelgesang, 2013). Binomial yield is used when failures of different units in a batch are uncorrelated and occur with the same probability. Stochastically proportional yield, on the other hand, is used to characterize the situation in which a random process affects whole batches, and the proportion of usable units in an order is a random variable. Models using interrupted geometric yield assume that good items are generated independently with a fixed probability until a failure occurs, and thereafter all items are defective.

Most papers consider the effect of random yield under the assumption of zero lead time. Henig & Gerchak (1990) were among the first to study the structure of optimal policies in single-sourcing periodic review systems with random yield. They showed that, despite the existence of a reorder point, the optimal order quantity is not linear in the inventory position. Bollapragada & Morton (1999) and Inderfurth & Transchel (2007) revisited this problem and proved that the infinite-horizon periodic-review model can be reduced to a newsvendor problem. However, the distribution of the key variable in the newsvendor problem depends on the order quantity in each period. They therefore proposed several myopic heuristics. Li et al. (2008) found upper and lower bounds for the optimal reorder point and order quantity in an infinite-horizon model and provided valuable insights into the structure of the optimal policies.

Among the well-performing heuristics proposed for the inventory optimization problem with one unreliable supplier, many fall into the class of 'linear inflation rules'. 'Linear inflation rules' restrict the order quantity to a linear function of inventory position with two parameters, called the 'order-up-to level' and the 'inflation factor'. Some of the myopic heuristics proposed by Bollapragada & Morton (1999) fall into this class. Huh & Nagarajan (2010) found the optimal policy within this class and proved that the average total cost is convex in the order-up-to level for any given inflation factor. The study of Inderfurth & Vogelgesang (2013) was one of the few to consider the effect of positive lead time. The authors capture the two sources of uncertainty (i.e. yield and demand uncertainty) by the safety stock variable. Under the assumption that safety stock follows a normal distribution, they found the optimal safety stock levels for three different types of random yield. Inderfurth & Kiesmüller (2013) proposed two approaches to derive the optimal and near-optimal values for the order-up-to level for a given inflation factor. The first approach models the on-hand inventory by a Markov chain and is exact for zero lead time. For general lead time, the approximate approach is analyzed by assuming a standard or gamma distribution of the on-hand inventory.

Dual sourcing is often used for balancing cost and service level. Whittemore & Saunders (1977) proved that for periodic review models and a difference in lead time between the two suppliers equal to one, the optimal policy is a dual-index order-up-to policy. However, when the difference between lead times is larger than one, the optimal policy is hard to derive. Several heuristics have therefore been proposed in the literature. Veeraraghavan & Scheller-Wolf (2008) showed that the DOP performs well in dual-sourcing models with general lead times, and proved that for any given difference between the order-up-to levels, the optimal expedited order-up-to level can be found by solving a specific newsvendor problem. However, for finding the distribution of the demand in the newsvendor problem, they relied on simulation. Arts et al. (2011) proposed an approximation of this distribution, which is exact when the difference between the order-up-to levels is one or approaches infinity. Sheopuri et al. (2010) generalized the DOP and proposed three new policies for the same model: namely, the vector base-stock policy, the weighted DOP and the demand allocation policy. The first two policies use an order-up-to rule for the expedited supplier and the state information for deciding the regular order quantities. The last policy uses an order-up-to rule for the regular supplier and allocates demand between the two suppliers based on myopic costs. The authors show numerically that the three policies outperform on average the optimal DOP in either cost saving or computational time. Besides the DOP, other types of heuristics have also been proposed. Tagaras & Vlachos (2001) considered an order-up-to policy which places regular orders periodically to restore the inventory position to the target level and emergency orders only when the likelihood of a stock-out is very high. Allon & Van Mieghem (2010) studied a continuous review inventory model with two suppliers and proposed a tailored base-surge policy for this model. The cheap, offshore supplier is considered as the 'base' from which the buyer replenishes at a constant rate, while the responsive, nearshore supplier acts as the 'surge' from which the buyer replenishes only when on-hand inventory is below a certain level. They presented bounds on the optimal cost and an asymptotically optimal policy for a high volume system. A simple 'square-root' formula is presented which gives valuable insight into how to allocate orders between the two sources.

Statement of contribution: The contributions of this paper to the literature may be summarized as follows. First, we develop a simple order-up-to heuristic (the OPMD heuristic) for a single-sourcing model with positive lead time and binomial yield. The proposed order-up-to level is found based on an approximating inventory model with modified demand distribution and reliable supplier. We show that our heuristic performs well by comparing it with the optimal policy and the heuristic proposed in Inderfurth & Vogelgesang (2013). Second, we consider the model in which an expedited, reliable supplier is used for mitigating the risk posed by the unreliable supplier. To the best of our knowledge, this model has not been previously studied by the OR community. To solve it, we propose a dual-index order-up-to policy based on an approximate model with two reliable suppliers and modified demand distribution (the DOPMD heuristic). When compared to the optimal dual-index order-up-to policy, our heuristic gives promising results.

3. The Single-Sourcing Inventory Model with Unreliable Supplier

We consider an infinite-horizon periodic-review model with an unreliable supplier. For each order X placed with the supplier, only a binomial random portion $B(X, p)$ is returned, where $0 < p < 1$ is the long-run average fraction of orders being returned. We assume that p is known in advance. Demand in different periods, denoted as D_n, $n = 1, 2, \cdots$, is assumed to be independent and identically distributed, with $E(D) < \infty$. Revealed demand is fulfilled from on-hand inventory I, and unsatisfied demand is fully backlogged. Ordered items are delivered after a positive lead time l. The exact number of units returned remains unknown until delivery. The retailer pays a variable ordering cost c for each ordered unit. We assume zero fixed ordering cost. Backlogged demand is charged a penalty cost b per unit per period while inventory carried at the end of a period is charged a holding cost h per unit per period.

The sequence of events in each period is as follows. First, on-hand inventory is observed. Second, an order is placed according to the inventory control policy that is applied. Third, a binomial random portion of the order placed l periods in the past arrives. Fourth, demand of this period is revealed and fulfilled or backlogged.

We are interested in finding an efficient inventory control policy that minimizes the long-run average total cost given by $\lim_{N\mapsto\infty}\frac{\sum_{n=1}^{N}TC_n}{N}$, with

$$TC_n = cX_n + hI_n^+ + bI_n^-,$$

where X_n and I_n are the order placed and the on-hand inventory in period n respectively, $a^+ = \max(a, 0)$ and $a^- = \max(-a, 0)$.

Notations used in this paper are summarized in Table 1.

Table 1: Notations

Notations	Descriptions	Notations	Descriptions
n	Period index	c	Per unit ordering cost
I_n	On-hand inventory in period n	h	Inventory holding cost per unit per period
IP_n	Inventory position in period n	b	Penalty cost per unit per period
X_n	Order placed in period n	l	lead time
D_n	Demand in period n	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

An order-up-to policy with modified demand (the OPMD heuristic)

The optimal policy for the single-sourcing model with yield uncertainty can in principle be found by using a Markov decision process. Due to state space explosion of the underlying Markov chain, this approach is computationally intractable for large lead times. We therefore propose an order-up-to heuristic with optimal order-up-to levels determined on the basis of an approximate inventory model with full returns.

Without loss of generality, we assume that the system starts with zero items in transit; in other words, $X_0 = 0$.

To motivate our approximation, consider the single-sourcing inventory model described above with the order X_{n+1} in period $n+1$ defined by

$$X_{n+1} = B(X_{n-l}, 1-p) + D_n. \tag{1}$$

Lemma 1 The sequence of orders X_n, $n = 1, 2, 3, ...$ has a limiting distribution.

Proof By using iteratively (1), we obtain

$$\begin{aligned} X_{n+1} &= D_n + B(D_{n-l-1}, 1-p) + B(D_{n-2l-2}, (1-p)^2) + \cdots \\ &= \sum_{k=0}^{\lfloor \frac{n}{l+1} \rfloor} R_{n,k}, \end{aligned}$$

with $R_{n,k} = B(D_{n-k(l+1)}, (1-p)^k)$. 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 X_n converges almost surely.

The probability-generating function $\hat{R}_k$ of R_k is given by $\hat{R}_k(z) = \hat{D}(q_k z + (1 - q_k))$, where $q_k = (1-p)^k$ and $\hat{D}$ is the probability-generating function of D. 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. □

Let F_∞ be the limiting distribution of X_n. Consider a sequence of independent variables Y_n, $n = 1, 2, ...$, distributed according to F_∞. We approximate the model with uncertain yield with a model with full returns and demand in period n given by

$$D'_n = B(Y_n, 1-p) + D_n.$$

We call D'_n the virtual demand in the model with full returns. Observe that although the variables $B(X_n, 1-p) + D_n$, $n = 1, 2, ...$ are dependent, by our choice of Y_n, the variables D'_n, $n = 1, 2, ...$ are independent.

Remark In the model with full returns, the next recursion holds

$$I_{n+1} = I_n + D'_{n-l-1} - D'_n,$$

whereas in the model with binomial return, we have

$$\begin{aligned} I_{n+1} &= I_n + B(X_{n-l}, p) - D_n \\ &= I_n + X_{n-l} - [B(X_{n-l}, 1-p) + D_n]. \end{aligned}$$

When $n \mapsto \infty$, X_{n-l} has the same limiting distribution as D'_{n-l-1} and $B(X_{n-l}, 1-p) + D_n$ the same limiting distribution as D'_n.

It is well known that in the classic model with full returns, the order-up-to policy is optimal and that in each period, the order placed is equal to the demand in the previous period. Therefore the next equation holds:

$$I_n = z - (D'_{n-l-1} + D'_{n-l} + \cdots + D'_{n-1}),$$

where z is the order-up-to level. So the optimal order-up-to level in the approximate system can be found by solving a newsvendor problem; i.e.

$$z^* = F^{-1}_{D'(l+1)}(\frac{b}{b+h}),$$

where $F_{D'(l+1)}$ is the cumulative distribution function of $\sum_{k=0}^{l} D'_{n-k}$ for all n.

The performance of the proposed heuristic (the OPMD heuristic) in the original problem will be tested in Section 5 by comparing it with the optimal policy derived by dynamic programming and the safety stock policy proposed by Inderfurth & Vogelgesang (2013).

4. The Dual-Sourcing Inventory Model with Unreliable Supplier

This section considers the inventory system of a retailer who sources from two suppliers, a regular (r) and an expedited (e) supplier. The lead time l_r of the regular supplier is longer than the lead time l_e of the expedited supplier, while the the ordering cost c_r of the regular supplier is lower than the cost c_e of the expedited one. Moreover, the regular supplier has binomial random yield, which means that, out of an order X^r_n placed with him in period n, only a random portion $B(X^r_n, p)$ turns out to be usable when the order is delivered in period $n + l_r$. On the other hand, if an order X_n^e is placed with the expedited supplier in period n, the whole order will be delivered in period $n + l_e$. To the best of our knowledge, this model seems not to have been studied before in the literature.

For the case with two reliable suppliers, Veeraraghavan & Scheller-Wolf (2008) showed that the performance of the dual-index order-up-to policy(DOP) is close to that of the optimal policy. This section therefore focuses on finding the optimal DOP for the model with two suppliers (one of which is unreliable).

Each DOP is characterised by two order-up-to levels: one for the expedited supplier, z_e, and one for the regular supplier, z_r. In each period $n \geq l_r$, there are l_r regular and l_e expedited orders in the pipeline, denoted by $< X_{n-l_r}^r, ..., X_{n-1}^r >$, and $< X_{n-l_e}^e, ..., X_{n-1}^e >$, respectively. The expedited inventory position in period n, IP_n^e, is comprised of on-hand inventory and all of the orders due to arrive in the next l_e periods, while the regular inventory position IP_n^r is comprised of on-hand inventory and all the orders that will arrive in the next l_r periods. More precisely,

$$IP_n^e = I_n + (X_{n-l_e}^e + ... + X_{n-1}^e) + (X_{n-l_r}^r + ... + X_{n-l-1}^r)$$
$$IP_n^r = I_n + (X_{n-l_e}^e + ... + X_{n-1}^e) + (X_{n-l_r}^r + ... + X_{n-1}^r),$$

where $l = l_r - l_e$.

In each period n, the following sequence of events takes place. First, an expedited order X_n^e is placed, to restore the inventory position IP_n^e to the value z_e. Observe that when the size of X_n^e is decided, X_{n-l}^r enters the information horizon. Thus, one first checks if there is a surplus (i.e., whether $IP_n^e + X_{n-l}^r > z_e$). If this is the case, no expedited order is placed. Otherwise, an expedited order equal to the deficit $X_n^e = z_e - (IP_n^e + X_{n-l}^r)$ is placed. Then the expedited order X_n^e is added to the inventory position of the regular supplier, IP_n^r and a regular order $X_n^r = z_r - (IP_n^r + X_n^e)$ is placed. Finally, the orders due to arrive in this period, $X_{n-l_r}^r$ and $X_{n-l_e}^e$ arrive. Note that since the regular supplier is unreliable, only $B(X_n^r, p)$ units are usable. Finally, demand D_n is revealed and satisfied from the on-hand inventory, if available. Unsatisfied demand is back-ordered. The inventory level is then updated and holding or penalty costs are incurred.

In the literature, the quantity $O_n = (IP_n^e + X_{n-l}^r - z_e)^+$ is known as *the overshoot*. The overshoot and the inventory positions of the regular and expedited supplier satisfy the following equations:

$$IP_n^e + X_{n-l}^r + X_n^e = z_e + O_n \tag{2}$$

$$IP_n^r + X_n^e + X_n^r = z_r. \tag{3}$$

Subtracting (2) from (3), we obtain

$$\sum_{k=0}^{l-1} X_{n-k}^r = z_r - z_e - O_n$$

and

$$\sum_{k=0}^{l-1} E(X_{n-k}^r) = z_r - z_e - E(O_n).$$

The optimal DOP 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. The next section therefore proposes a dual-index order-up-to heuristic that can be used for large values of l_r.

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

As in the single-sourcing case, we propose approximating the dual-sourcing model with uncertain yield with a model with full returns, but with modified demand distribution.

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

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

If the variables $D_n + B(X^r_{n-l_r}, 1-p)$ were independent and their distribution easy to calculate, we could reduce the model with uncertain returns to a model with full returns and demand 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 the distribution of X^r_n difficult to find. We therefore propose using the following approximation.

Let Y_n be a random variable distributed according to F_∞, the limiting distribution of the orders in a system where the only supplier is the regular supplier. Observe that in the dual-sourcing model, X^r_n is usually smaller than Y_n, since part of the orders is delivered by the expedited supplier. We assume that $X^r_n = B(Y_n, \alpha)$, with $\alpha \in [0, 1]$. Thus, each unit that would be ordered from the regular supplier if he were the only supplier is now ordered with probability $1 - \alpha$ from the expedited supplier. To find an appropriate α, recall that $\sum_{k=0}^{l-1} E(X^r_{n-k}) = z_r - z_e - E(O_n)$. Since $E(O_n) \geq 0$ and $E(X^r_n) = \alpha E(Y_n)$, it holds that $\alpha l E(Y_n) \leq z_r - z_e$. We therefore propose choosing $\alpha = \min\{\frac{\Delta}{lE(Y_n)}, 1\}$, where $\Delta = z_r - z_e$. Since the cumulative distribution function of Y_n is $F_\infty(\cdot)$, $E(Y_n) = \frac{E(D_n)}{p}$ and $\alpha = \min\{\frac{\Delta p}{lE(D_n)}, 1\}$.

We are now able to describe the approximate dual-sourcing model with full returns. In the approximate model, both retailers are assumed to be reliable. Their costs and lead times are as in the initial model. We define the demand in period n as

$$D'_n = D_n + B(Y_n, \alpha(1-p)).,\tag{4}$$

where $\alpha = \min\{\frac{\Delta p}{lE(D_n)}, 1\}$. Since the variables Y_n are independent and identically distributed, so are the variables D'_n, $n = 1, 2, \cdots$.

It has been proven that for any fixed Δ, the optimal expedited order-up-to level in the dual-sourcing model with full returns can be found by solving a newsvendor problem (Veeraraghavan & Scheller-Wolf, 2008); i.e.

$$z_e^* = F^{-1}_{D'(l_e+1)-O}\left(\frac{b}{b+h}\right),$$

where $F_{D'(l+1)-O}$ is the cumulative distribution function of $\sum_{k=0}^{l} D'_{n-k} - O_{n-l}$ for all n. As in Veeraraghavan & Scheller-Wolf (2008), for each Δ, we derive the distribution of O_n by simulation and then determine the optimal expedited order-up-to level and the optimal total cost. Subsequently, we use one-dimensional search to find the optimal value for Δ. Note that, in order to reduce computation times, the distribution of O_n could also be approximated as described in Arts et al. (2011). This is, however, not the focus of this paper.

Section 5 testifies to the performance of DOPMD by comparing it with optimal DOP for the given model.

5. Numerical Results

This section presents numerical results on the performance of the proposed heuristics for the single- and dual-sourcing models.

5.1. Performance of the Heuristic for the Single-Sourcing Model (the OPMD Heuristic)

To study the influence of the parameters on the performance of the OPMD heuristic, we construct 74 different scenarios. We start with a base case in which the parameters take the values $h = 5$, $c = 150$, $l = 2$, $p = 0.8$, $b = 495$ and $D \sim U\{0, 1, \cdots, 4\}$[1]. Subsequently, we vary the values of one or two parameters and keep the others as in the base case. The optimal order-up-to level for the OPMD heuristic is found by solving the newsvendor problem in the approximate model with full returns. The average total cost for the given optimal order-up-to level is calculated as the long-run average cost of the underlying Markov chain. For small instances, we compare the OPMD heuristic with both the optimal policy and the safety stock policy proposed in Inderfurth & Vogelgesang (2013). The optimal policy is derived by using dynamic programming. For large instances, we only compare the OPMD heuristic with the safety stock policy.

To keep the dynamic program tractable, we focus on discrete demand distributions with bounded support. As $b, h > 0$, we restrict the backlogs and on-hand inventory to $[0, \lceil \frac{(l+1)D_{max}}{p} \rceil]$, where D_{max} denotes the maximum demand and $\lceil x \rceil$ denotes the minimum integer that is larger than or equal to x. Notice that the probability of the backlog being larger than $\lceil \frac{(l+1)D_{max}}{p} \rceil$ is

[1] $U\{0, 1, \cdots, n\}$ denotes the discrete uniform distribution on $\{0, 1, \cdots, n\}$

smaller than $(Pr(D = D_{max}))^{(l+1)}$ and that of the on-hand inventory being larger than $\lceil \frac{(l+1)D_{max}}{p} \rceil$ is smaller than $(Pr(D = 0))^{(l+1)}$. The order quantity is restricted to $[0, \lceil \frac{2D_{max}}{p} \rceil]$. Note that since every ordered unit is returned with probability p, the expected number of units that need to be ordered to get one unit returned is $\frac{1}{p}$. Hence, the probability of order quantity exceeding $\lceil \frac{2D_{max}}{p} \rceil$ is very small. Moreover, in all of our numerical experiments, the order quantities in the optimal policy did not exceed $\lceil \frac{2D_{max}}{p} \rceil$.

Impact of Yield Rate Next we examine the performance of the OPMD heuristic under different yield rates. We vary $p \in \{0.4, 0.6, 0.8, 1\}$ and $D \sim U\{0, 1, \cdots, n\}$, $n = 2, 4$ and compare the performance of the OPMD heuristic, the optimal policy and the safety stock policy. The results are shown in Table 2. The average relative difference between the OPMD heuristic and the optimal policy is 0.97% and the maximum difference is 2.35%. As shown in column 4 of Table 2, the performance of the OPMD heuristic improves when the yield rate increases. This is due to the fact that the OPMD heuristic assumes independent virtual demands, which holds if orders from different periods are independent. When the yield rate is high, the unreturned order quantities are relatively small, which leads to less correlation among orders.

On the other hand, the performance of the safety stock policy improves when the yield rate decreases, which can be seen in column 5 of Table 2. The average and maximum difference between the safety stock and the optimal policy are 1.36% and 3.86%, respectively. As the results in Table 2 indicate, when the yield rate is relatively high, our heuristic performs better than the safety stock policy. The reverse seems to hold for low yield rates. The

Table 2: Impact of yield rate (h=5, l=2, b=495, c=150)

		Optimal policy	OPMD	Safety stock policy
p	Demand dist.	Average total cost	% above optimal	% above optimal
0.4	U{0,1,2}	400.08	**2.35**	**0.31**
0.6	U{0,1,2}	273.01	**1.53**	**0.76**
0.8	U{0,1,2}	208.11	0.71	1.06
1	U{0,1,2}]	165.00	0.00	3.86
0.4	U{0,1,2,3,4}	789.94	**1.68**	**0.52**
0.6	U{0,1,2,3,4}	537.07	**1.11**	**0.85**
0.8	U{0,1,2,3,4}	408.87	0.40	1.50
1	U{0,1,2,3,4}	329.00	0.00	2.00

same patterns hold for the larger instances shown in Table 3, where $D \sim U\{0, 1, ..., 8\}$, $l \in \{2, 4, 8, 20\}$ and $p \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. Note that for these instances, since the state space of the dynamic program grows too large, we only compare the OPMD heuristic with the safety stock policy. *Impact of Lead Time* To study the impact of lead time on the performance of the OPMD heuristic, we first compare it with the optimal and the safety stock policy in small instances. For this, we modify the base case by first taking $D \sim U\{0, 1, 2\}$ and $l \in \{1, 2, 4, 6, 7\}$ and then $D \sim U\{0, 1, 2, 3, 4\}$ and $l \in \{1, 2, 4, 6\}$. The results appear in Table 4. For larger lead times, due to state space explosion, it is computationally intensive to find the optimal policy by dynamic programming. For these instances, we therefore only compare the OPMD heuristic with the safety stock policy. The results are summarized in Table 3.

The average and maximum difference (over all nine scenarios in Table 4)

Table 3: Impact of yield rate and lead time ($h = 5, b = 495, c = 150, D \sim U\{0, 1, \cdots, 8\}$)

		OPMD	Safety stock policy
l	p	Average total cost	% above proposed heuristic
2	0.1	6171.94	**-1.46**
2	0.3	2101.52	**-1.06**
2	0.5	1278.58	**-0.09**
2	0.7	922.49	1.07
2	0.9	723.02	2.21
4	0.1	6216.61	**-1.85**
4	0.3	2129.63	**-1.46**
4	0.5	1301.46	**-0.46**
4	0.7	942.59	0.72
4	0.9	743.11	1.73
8	0.1	6273.50	**-2.25**
8	0.3	2174.27	**-2.14**
8	0.5	1336.58	**-0.97**
8	0.7	973.78	0.20
8	0.9	771.22	1.14
20	0.1	6405.41	**-3.23**
20	0.3	2263.68	**-3.22**
20	0.5	1409.05	**-1.84**
20	0.7	1036.35	**-0.38**
20	0.9	827.79	0.55

between the OPMD heuristic and the optimal policy is 0.49% and 0.89%, respectively. We observe that the OPMD heuristic deviates slightly less from the optimal policy when lead time increases. To explain this, recall that the OPMD heuristic assumes independent virtual demands, and hence, independent order quantities in the original model. Since an order depends only on the orders placed $k(l+1)$ periods in the past, with $k \geq 1$, the larger the lead time, the less is the correlation among different orders. Moreover, we notice that the performance of the OPMD heuristic seems insensitive to changes in lead time. On the other hand, as column 5 in Table 4 shows, the safety stock policy performs significantly better for larger lead times.

To examine the performance of the OPMD heuristic for larger lead times, we refer to the rows corresponding to $l \in \{8, 20\}$ in Table 3. As column 4 in Table 3 indicates, the safety stock policy outperforms our heuristic for large lead times and relatively low yield rates. The reverse seems to hold for large lead times and high yield rates ($p = 0.9$).

Table 4: Impact of lead time (h=5, p=0.8, b=495 and c=150)

		Optimal policy	OPMD	Safety stock policy
l	Demand dist.	Average total cost	% above optimal	% above optimal
1	U{0,1,2}	203.56	0.89	2.00
2	U{0,1,2}	208.11	0.71	1.06
4	U{0,1,2}	214.76	0.60	0.98
6	U{0,1,2}	220.04	**0.47**	**0.29**
7	U{0,1,2}	222.37	**0.44**	**0.32**
1	U{0,1,2,3,4}	401.62	0.36	1.82
2	U{0,1,2,3,4}	408.87	0.40	1.50
4	U{0,1,2,3,4}	419.92	0.28	1.26
6	U{0,1,2,3,4}	428.68	0.24	0.51

Impact of Penalty Cost In order to study the influence of the penalty cost, we set $b \in \{5, 15, 95, 495\}$. Note that the penalty cost influences the optimal order-up-to level through the optimal fractile in the newsvendor

Table 5: Impact of penalty cost (h=5, l=2, p=0.8 and $D \sim U\{0,1,2,3,4\}$)

		Optimal policy	OPMD	Safety stock policy
b	c	Average total cost	% above optimal	% above optimal
5	5	23.29	1.15	6.06
15	5	29.65	1.39	8.11
95	5	39.62	1.78	13.01
495	5	46.37	3.46	14.70
5	10	35.79	0.85	2.76
15	10	42.15	1.08	5.65
95	10	52.12	1.44	9.88
495	10	58.87	2.69	12.00
5	50	65.00	109.38	111.12
15	50	142.15	0.32	0.89
95	50	152.12	0.49	3.85
495	50	158.87	0.90	4.40
5	150	65.00	493.05	498.63
15	150	195.00	101.01	102.66
95	150	402.12	0.03	1.43
495	150	408.87	0.40	1.50

problem in the model with full returns. For $h = 5$, the optimal fractile $\frac{b}{b+h} \in \{0.5, 0.75, 0.95, 0.99\}$. Moreover, we vary the value of the ordering cost in $c \in \{5, 10, 50, 150\}$. As can be seen in Table 5,the deviation of the OPMD heuristic from the optimal policy increases, in general, when the penalty cost (the optimal fractile) increases. However, when the penalty cost is much lower than the ordering cost (e.g. $b = 5, c = 50, 150$ and $b = 15, c = 150$), the OPMD heuristic leads to a large deviation from the optimal policy. This phenomenon can also be seen when the safety stock policy is applied. The reason is that the optimal policy is influenced by the ordering costs, while both the OPMD heuristic and the safety stock policy are not. When the ordering cost is much higher than the penalty cost, it is more cost-efficient to backlog demand and incur penalty cost than to order. Neither of the heuristics takes this aspect into account. If we exclude the three exceptional cases, the average deviation of the OPMD heuristic from the optimal policy is 1.20%, with the maximum being 3.46%, while the average deviation of the safety stock policy is 6.36%, with the maximum being 14.70%. The OPMD heuristic outperforms the safety stock policy in all cases shown in Table 5.

Impact of Mean, Variance and Skewness of Demand Distribution This section examines the influence of the demand distribution on the performance of the OPMD heuristic, by varying its mean, variance and skewness. In order to study the impact of mean, we choose demand distributions with the same variance and skewness but different means. For $k, n \in \mathbf{Z}_+$ and $k \leq n$, let $U\{n-k, n, n+k\}$ denote the distribution given by $Pr(D = n-k) = Pr(D = n) = Pr(D = n+k) = 1/3$. The skewness of this distribution is equal to 0. When $k = 1$, the variance of the distribution is $\frac{2}{3}$ and when $k = 2$, the variance of the distribution equals $\frac{8}{3}$. Table 6 contains the detailed results for this demand distribution. the OPMD heuristic seems robust under changes in mean demand, with an average deviation from the optimal policy of 0.41% and a maximum deviation of 0.67%. Moreover, the performance of the OPMD heuristic slightly improves when the mean demand increases.

Next we change the variance of the demand distribution while keeping constant the mean and the skewness. The results shown in Table 14 in the appendix, testify that the performance of the OPMD heuristic is also robust

Table 6: Impact of Mean Demand (h=5, b=495, c=150, l=2 and p=0.8)

			Optimal policy	OPMD	Safety stock policy
Demand dist.	Mean	Variance	Average total cost	% above optimal	% above optimal
U{0,1,2}	1	2/3	208.15	0.66	0.92
U{1,2,3}	2	2/3	400.77	0.49	1.79
U{2,3,4}	3	2/3	588.73	0.26	0.37
U{3,4,5}	4	2/3	778.60	0.23	0.27
U{0,2,4}	2	8/3	412.33	**0.59**	**0.17**
U{1,3,5}	3	8/3	601.52	**0.44**	**0.08**
U{2,4,6}	4	8/3	790.80	**0.34**	**0.05**
U{3,5,7}	5	8/3	979.96	**0.27**	**0.08**

against demand variability. The average deviation from the optimal policy is 0.45% and the maximum deviation is 0.88%. In our experiments, the OPMD heuristic outperforms the safety stock policy for small demand variances ($var(D) \in \{0, 2/3\}$), while the safety stock policy gives better results for larger demand variability ($var(D) \geq 1$).

In the end, we examine the influence of the skewness of demand distribution by choosing $D \sim NB(r, q)$, where $NB(r, q)$ denotes the negative binomial distribution with r being the number of failures until the experiment stops and q being the probability of success for each trial. In our experiments, we vary $r \in \{1, 2, 4, 6\}$ and $q \in \{0.2, 0.4\}$. In order to acquire distributions with different skewness, we truncate $NB(r, 0.2)$ to take values in $[0, r]$ and $NB(r, 0.4)$ to take values in $[0, 4r/3]$. The skewness for the truncated distributions is shown in column 2 of Table 7. As can be seen in column 4, the performance of the OPMD heuristic is robust against changes in the skewness of the demand distribution. The average deviation in average total costs from the optimal policy is 0.33%, while the maximum deviation is 0.70%. Compared with the safety stock policy, the OPMD heuristic performs better when the skewness is negative and has a large absolute value. When skewness is positive and has a small absolute value, the safety stock policy outperforms the OPMD heuristic.

Table 7: Impact of Skewness of Demand (h=5, b=495, c=150, l=2 and p=0.8)

		Optimal policy	OPMD	Safety stock policy
Demand dist.	skewness	Average total cost	% above optimal	% above optimal
NB(1,0.2)	-23.44	163.45	0.34	0.51
NB(2,0.2)	-34.91	364.90	0.39	0.81
NB(4,0.2)	-159.86	761.98	0.17	0.40
NB(6,0.2)	-934.14	1148.30	0.18	0.52
NB(1,0.4)	0.87	162.21	**0.70**	**0.86**
NB(2,0.4)	0.11	363.86	**0.32**	**0.26**
NB(4,0.4)	0.02	772.82	**0.29**	**0.08**
NB(6,0.4)	0.01	1177.3	**0.25**	**0.12**

5.2. Performance of the Heuristic for the Dual-Sourcing Model (the DOPMD heuristic)

This section studies the performance of the DOPMD heuristic by comparing it with the optimal dual-index order-up-to policy (the optimal DOP) for the studied model. The reason for using it as a benchmark is twofold: first, deriving the optimal policy for the dual-sourcing model with general lead times and random yield is computationally intensive even for small lead times and demand; second, the DOP has been proven to have a near optimal performance in dual-sourcing models with general lead times (Veeraraghavan & Scheller-Wolf, 2008). 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 either the 95% confidence intervals for both the expected on-hand inventory and the expected backlogged demand are smaller than 0.025 or the standard error is below 0.001 times the expected value for both the on-hand inventory and the backlogged demand. For our heuristic, the order-up-to levels are found by applying the solution procedure proposed in Veeraraghavan & Scheller-Wolf (2008) to a dual-sourcing model with full returns and modified demand defined by equation (4). When deriving the distribution of the overshoot, we run the simulation until either the 95% 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 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 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.

As in Section 5.1, we start with a base case and construct 30 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. We fix the values of

h, c_e and l_e in all instances, and then study the respective impact of c_r, p, l_r, b and demand on the performance of the DOPMD heuristic. The parameter values used in this section are summarised in Table 8.

Table 8: Parameter values in the dual-sourcing model

Parameter	Values
f_D	Poisson(λ), $\lambda \in \{2, 4, 6, 8\}$
l_r	2, 4, 6, 8
c_r	40, 70, 100, 130
b	10, 15, 95, 495
p	0.6, 0.7, 0.8, 0.9, 1

Impact of yield rate We begin by examining the impact of yield rate on the performance of the DOPMD heuristic by taking $p \in \{0.6, 0.7, 0.8, 0.9, 1\}$, $l_r \in \{2, 4\}$ and all of the other parameters as in the base case. The results appear in Table 9. As can be seen in column 8, the maximum deviation of the DOPMD heuristic from the optimal DOP is 1.64%, while the average deviation is 0.66%. Column 3, 4, 6 and 7 indicate the expected regular and expedited order quantities for the optimal DOP and the DOPMD heuristic. When $l_r = 2$, as can be seen in column 6 and 7, the DOPMD switches from single-sourcing from the expedited supplier (for $p = 0.6$) directly to single-sourcing from the regular supplier (for $p \geq 0.7$), whereas, as can be seen in column 3 and 4, the optimal DOP leads to dual-sourcing for the medium yield rate (i.e. $p = 0.7$). When $l_r = 4$, both the optimal DOP and the DOPMD heuristic switch from single-sourcing from the expedited supplier to dual-sourcing when the yield rate increases from 0.6 to 0.7 and then order almost exclusively from the regular supplier for high yield rates (for $p \in \{0.9, 1\}$).

Moreover, the DOPMD switches to ordering a larger portion from the regular supplier at a higher yield rate (at $p = 0.8$) compared with the optimal DOP (at $p = 0.7$).

Table 9: Impact of yield rate (c_r=100, $b = 495$ and $D \sim Pois(2)$)

		Optimal DOP			DOPMD		
l_r	p	E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	% above optimal DOP
2	0.6	0.00	2.00	329.98	0.00	2.00	0.03
2	0.7	2.18	2.48	324.02	2.86	0.00	1.64
2	0.8	2.42	0.06	289.19	2.50	0.00	0.78
2	0.9	2.21	0.00	259.53	2.23	0.00	0.39
2	1	1.98	0.00	233.89	2.00	0.00	1.41
4	0.6	0.00	2.00	328.71	0.00	2.00	0.82
4	0.7	1.74	0.78	325.77	0.63	1.56	1.04
4	0.8	2.31	0.15	294.67	2.44	0.05	0.52
4	0.9	2.13	0.09	266.96	2.12	0.09	-0.03
4	1	2.00	0.00	243.33	1.95	0.05	-0.01

Impact of regular lead time To analyze the influence of the regular lead time, we take $l_r \in \{2, 4, 6, 8\}$, $b \in \{95, 495\}$ and the other parameters as in the base case. The results are reported in Table 10. As can be seen in column 8, the maximum deviation of the DOPMD heuristic from the optimal DOP is 2.45%, while the average deviation is 0.74%. The regular lead time seems to have no significant effect on the performance of the DOPMD heuristic. As can be seen in column 3, 4, 6 and 7, as the regular lead time increases, the expected expedited order quantity increases slightly while the expected regular order quantity decreases.

Impact of regular ordering cost and penalty cost To examine the impact

Table 10: Impact of regular lead time (c_r=100, $p = 0.8$ and $D \sim Pois(2)$)

		Optimal DOP			DOPMD		
b	l_r	E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	% above optimal DOP
495	2	2.42	0.06	289.19	2.50	0.00	0.73
495	4	2.31	0.15	294.67	2.44	0.05	0.41
495	6	2.11	0.31	296.93	2.35	0.12	0.59
495	8	2.07	0.34	297.50	2.33	0.13	0.90
95	2	2.46	0.00	278.93	2.50	0.00	0.81
95	4	2.41	0.00	282.12	2.50	0.00	2.45
95	6	2.35	0.12	288.33	2.40	0.08	0.14
95	8	1.86	0.46	287.27	2.37	0.08	-0.09

of the regular ordering cost, we vary in the base case $c_r \in \{40, 70, 100, 130\}$ and $b \in \{95, 495\}$. The results appear in Table 11. The average deviation of the DOPMD heuristic from the optimal DOP is 0.77%, while the maximum deviation is 1.67%. As can be seen from column 6 and 7, the DOPMD heuristic changes from single-sourcing from the regular supplier (i.e. when $c_r \in \{40, 70, 100\}$) to single-sourcing from the expedited supplier (i.e. when $c_r = 130$) under both penalty costs. On the other hand, the optimal DOP places a small portion of the total orders with the expedited supplier even when the regular ordering cost is low but the penalty cost is high (i.e. when $c_r \in \{40, 100\}$ and $b = 495$).

Next we study the influence of the penalty cost as well as the optimal fractile on the performance of the DOPMD heuristic. For this we take $b \in \{10, 15, 95, 495\}$, which results in an optimal fractile $\frac{b}{b+h} \in \{0.67, 0.75, 0.95, 0.99\}$ for the newsvendor problem in the approximate model with full returns and

Table 11: Impact of regular ordering cost (l_r=2, $p = 0.8$ and $D \sim Pois(2)$)

		Optimal DOP			DOPMD		
b	c_r	E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	% above optimal DOP
95	40	2.47	0.00	129.70	2.49	0.00	0.49
95	70	2.48	0.00	203.81	2.50	0.00	0.93
95	100	2.47	0.00	278.93	2.49	0.00	0.48
95	130	0.00	2.00	319.17	0.00	1.98	0.52
495	40	2.51	0.01	139.75	2.50	0.00	1.25
495	70	2.46	0.00	213.39	2.50	0.00	1.67
495	100	2.42	0.06	289.19	2.51	0.00	1.05
495	130	0.00	2.00	332.24	0.00	2.00	-0.23

$c_r = 120$. Recall that the optimal fractile influences the expedited order-up-to level as $z_e^* = F^{-1}_{D'(l_e+1)-O}(\frac{b}{b+h})$. As column 7 of Table 12 shows, the DOPMD heuristic has an average deviation of 0.51% and a maximum deviation of 0.91% compared with the optimal DOP. Moreover, as can be seen from column 2, 3, 5 and 6, the DOPMD leads to single-sourcing from the expedited supplier for all values of the penalty cost, whereas the optimal DOP switches from single-sourcing from the regular supplier at low penalty cost (i.e. $b = 10$) to single-sourcing from the expedited supplier when the penalty cost increases (when $b \geq 15$).

Impact of demand size To examine the robustness of the DOPMD heuristic under different demand distributions, we change in the base case $D \sim Pois(\lambda)$, $\lambda \in \{2, 4, 6, 8\}$ and $c_r \in \{100, 120\}$. The results are shown in Table 13. As can be seen in column 8, the maximum and average deviation of the DOPMD heuristic from the optimal DOP are 0.68% and 2.39%, respectively.

Table 12: Impact of penalty cost (l_r=2, c_r=120, $p = 0.8$ and $D \sim Pois(2)$)

	Optimal DOP			DOPMD		
b	E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	% above optimal DOP
10	2.44	0.01	262.47	0.00	1.99	0.13
15	0.00	1.99	266.48	0.00	2.00	0.35
95	0.00	1.99	278.93	0.00	2.02	0.66
495	0.00	3.97	289.19	0.00	2.01	0.91

When the regular ordering cost is relatively low, i.e. $c_r = 100$, the DOP heuristic leads to single-sourcing from the regular supplier, whereas the optimal DOP occasionally places orders with the expedited supplier (i.e. when $\lambda \in \{2, 6\}$). When the regular ordering cost is high, the DOPMD heuristic switches to single-sourcing from the expedited supplier for all demand sizes, while the optimal DOP single sources from the expedited supplier when the demand size is small (i.e. when $\lambda \in \{2, 4\}$) and dual sources for large demand sizes (i.e. $\lambda \in \{6, 8\}$). Notice that when $c_r = 120$ and $\lambda = 8$, the optimal DOP seems to mainly rely on the expedited supplier. However, we found two other close to optimal solutions which rely more on the regular supplier. This is to be expected since $\frac{c_r}{p}$ and c_e are relatively close and so are the lead times.

6. Conclusions and Discussion

This paper studies both the single-sourcing and dual-sourcing inventory models with positive lead times and random yield. Yield uncertainty has rarely been considered in models with positive lead times and never in the dual-sourcing model with general lead times, which is the contribution of

Table 13: Impact of demand distribution (l_r=2, $p = 0.8$ and $b = 495$)

		Optimal DOP			DOPMD		
c_r	λ	E(X_r)	E(X_e)	Average total cost	E(X_r)	E(X_e)	% above optimal DOP
100	2	2.42	0.06	289.19	2.49	0.00	0.44
100	4	4.94	0.00	548.91	5.02	0.00	2.39
100	6	7.22	0.14	811.37	7.43	0.00	0.24
100	8	9.94	0.00	1069.25	10	0.00	0.81
120	2	0.00	1.99	327.25	0.00	2.01	0.91
120	4	0.00	3.97	636.73	0.00	3.99	0.57
120	6	2.95	3.57	942.08	0.00	5.93	-0.09
120	8	1.00	7.20	1257.82	0.00	8.01	0.14

this paper. For both models, we propose simple order-up-to heuristics. The optimal order-up-to levels are derived based on approximate models with full returns and modified demand distributions. Numerical results show that the performance of the proposed heuristic in the single-sourcing model is close to that of the optimal policy. Compared to the safety stock policy recently proposed by Inderfurth & Vogelgesang (2013), our heuristic seems to perform better than the safety stock policy when the yield rate is high or the lead time is small. For the dual-sourcing model, the numerical results indicate that the proposed heuristic gives, in most cases, results close to the optimal DOP. Moreover, the performance is robust with respect to changes in the main parameters.

References:

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

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

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

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

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

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

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

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

Inderfurth, K., & Kiesmüller, G. (2013). Exact and heuristic linear-inflation policies for an inventory model with random yield and arbitrary lead times. Working paper.

Inderfurth, K., & Transchel, S. (2007). Technical note on 'Myopic heuristics for the random yield problem'. *Operations Research*, *55*, 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*, 293–301.

Kazaz, B. (2004). Production planning under yield and demand uncertainty with yield-dependent cost and price. *Manufacturing & Service Operations Management*, *6*, 209–224.

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

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

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

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

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

Table .14: Impact of Variance of Demand ($h = 5$, $b = 495$, $c = 150$, $l = 2$ and $p = 0.8$)

			Optimal policy	OPMD	Safety stock policy
Demand dist.	Mean	Variance	Average total cost	% above optimal	% above optimal
U{1}	1	0	198.11	0.88	4.35
U{0,1,2}	1	2/3	208.15	0.66	0.92
U{0,2}	1	1	210.79	**0.75**	**0.21**
U{2}	2	0	391.01	0.49	0.85
U{1,2,3}	2	2/3	400.77	0.49	1.79
U{0,1,2,3,4}	2	2	408.87	**0.37**	**0.15**
U{0,2,4}	2	8/3	412.33	**0.44**	**0.08**
U{0,4}	2	4	417.00	**0.46**	**0.28**
U{3}	3	0	582.12	0.23	1.32
U{2,3,4}	3	2/3	588.73	0.26	0.37
U{1,2,3,4,5}	3	2	598.22	**0.29**	**0.15**
U{1,3,5}	3	8/3	601.52	**0.44**	**0.07**
U{1,5}	3	4	605.45	**0.65**	**0.27**
U{0,6}	3	9	621.57	**0.76**	**0.72**
U{4}	4	0	772.56	0.16	0.76
U{3,4,5}	4	2/3	778.60	0.23	0.27
U{2,3,4,5,6}	4	2	787.56	**0.28**	**0.14**
U{2,4,6}	4	8/3	790.80	**0.35**	**0.03**
U{2,6}	4	4	795.01	**0.44**	**0.13**