

LINEAR INFLATION RULES FOR THE RANDOM YIELD PRODUCTION CONTROL PROBLEM WITH UNCERTAIN DEMAND: ANALYSIS AND COMPUTATIONS

Woonghee Tim Huh

Department of Industrial Engineering & Operations Research
Columbia University
New York, N.Y. 10027, U.S.A.

Mahesh Nagarajan

Sauder School of Business
University of British Columbia
Vancouver, BC V6T-1Z2, CANADA.

ABSTRACT

Since the dawn of wafer fabrication and the production of microelectronic parts a fundamental characteristic of this environment has been uncertainty in production yields and in demand for product. The impact of the uncertainty is so prevalent that even deterministic models in practice have incorporated some allowance for uncertainty through features such as date effective yields, moving average capacity, etc. In this paper, we propose a simple heuristic approach for the inventory control problem with stochastic demand and multiplicative random yield. Our heuristic tries to find the best candidate within a class of policies which are referred to in the literature as the linear inflation rule (LIR) policies. Our approach is computationally fast, easy to implement and intuitive to understand. Moreover, we find that in a significant number of instances our heuristic performs better than several other well-known heuristics that are available in the literature.

1 INTRODUCTION

Random yield has been recognized as an important phenomenon in inventory management. It is well known that optimally managing periodic-review inventory control systems with random yield and stochastic demand is difficult. (See (Yano and Lee 1995) and (Zipkin 2000)). Thus the literature has seen several heuristic approaches to this problem. Perhaps the best known heuristic approaches were proposed by (Bollapragada and Morton 1999). These heuristics had several attractive features. First, they are myopic, i.e., in each period, the quantity ordered is based on a solution to a single-period problem. Second, many of these heuristic procedures fall under the class of what (Zipkin 2000) refers to as the *linear inflation rule* (LIR) policies. This policy class is easy to understand and is a generalization of the well-known base-stock policy that is commonly used in inventory management. Finally, as shown in (Bollapragada and Morton 1999), these heuristics perform extremely well

in several reasonable problem instances. (Bollapragada and Morton 1999) as in our paper, consider the infinite-horizon average cost criterion.

In a recent paper, (Inderfurth and Transchel 2007) revisit the same solution procedures used in (Bollapragada and Morton 1999), and test their performances for a wide range of problem settings. They point out certain flaws in the calculations in (Bollapragada and Morton 1999), which explain the steady deterioration in the myopic heuristics of (Bollapragada and Morton 1999). Further, they show that the performance of the heuristics proposed by (Bollapragada and Morton 1999) is worse than that reported in the original study, especially when the ratio of the backorder penalty cost to the holding cost is high. However, they do not provide any alternate heuristic approach. (Li, Xu, and Zheng 2006) also examine the heuristics of (Bollapragada and Morton 1999), but under the *discounted-cost criterion*. They demonstrate scenarios where the heuristics do not perform well and point out certain difficulties associated with implementing some of the heuristics in (Bollapragada and Morton 1999). They propose a new heuristic that is dependent on the discount factor and cannot be adopted to the average cost model. Thus, finding a good heuristic method for the random yield problem is an open area of research.

In this paper, we study the periodic-review random yield problem under the average cost criterion with emphasis on the class of LIR policies. Our contribution in this paper is both analytical and computational, and is as follows: (1) We show certain useful properties of the cost function when the class of LIR policies are used; in particular, we demonstrate a convexity property, which is useful in finding the best parameters for this policy class. (2) We propose the best candidate in the class of LIR policies as a heuristic for the random yield problem. This heuristic is intuitive to implement and computationally fast. Numerical experiments indicates that, for a wide class of problem instances, our heuristic performs better than other approaches available in the literature.

The rest of the paper is organized as follows: In Section 2, we state the problem and the notations used in our analysis. In Section 3, we demonstrate certain simple properties of the random yield problem and use these to construct our heuristic. In Section 4, we present computational results that demonstrate the performance of our approach. We also present comparisons with other approaches that exist in the literature.

2 MODEL

The description of the model is as follows. Let $t = 1, 2, \dots$ index time periods in a forward manner. Let $D_t \geq 0$ be a nonnegative random variable representing demand in period t . We assume that the sequence $\{D_t | t \geq 1\}$ is independent and identically distributed (IID), and we denote by D a generic demand random variable. Let $Y_t \geq 0$ denote a nonnegative random variable representing the stochastically proportional yield factor in period t . The sequence $\{Y_t | t \geq 1\}$ is also IID, and we denote by Y a generic yield factor random variable.

The sequence of events in each period is as follows. At the beginning of each period $t \geq 1$, the manager observes the beginning inventory level I_t , and orders Q_t . The replenishment is instantaneous, and the manager receives $Y_t \cdot Q_t$ units. Then, demand D_t is realized. Any excess inventory incurs the holding cost of $h \geq 0$ per unit, and any excess demand is backlogged at the rate of $b \geq 0$ per unit-period. Thus, the cost incurred in period t is given by $h \cdot IN_t^+ + b \cdot IN_t^-$, where $IN_t = I_t + Y_t \cdot Q_t - D_t$ is the ending inventory in period t . Let $I_{t+1} = IN_t$.

In this paper, we assume that the manager’s decision in each period is the order quantity Q_t , and his objective is to minimize the infinite-horizon average-cost of the system. As mentioned, the optimal order quantity in each period t depends on the starting inventory level I_t and does not have a simple structure ((Henig and Gerchak 1990)). Consequently, we restrict our attention to a heuristic where orders are placed in every period by following a linear-control policy which in the literature is referred to as the *linear inflation rule* (LIR) policy. This policy has two parameters θ and β . Under this policy, the order quantity in each period depends on the beginning inventory I_t and is given by

$$Q_t = Q(I_t) = \begin{cases} \beta \cdot (\theta - I_t) & \text{if } I_t \leq \theta, \\ 0 & \text{otherwise.} \end{cases}$$

We say θ is the *target inventory level*, and $\beta > 0$ is the *inflation factor*. (Note that if $\beta = 1$, then the LIR policy is the familiar base-stock policy with an order-up-to level of θ .)

Several control methods proposed as heuristic approaches to the random yield problem in the literature belong to the class of LIR policies. In (Bollapragada and Morton

1999), the following two LIR policies have been considered: (i) MULT METHOD where θ is the $b/(b+h)$ -fractile of the demand distribution D and $\beta = 1/E[Y]$, and (ii) NLH1 METHOD where θ is $E[D]$ plus the $b/(b+h)$ -fractile of $D - (E[D]/E[Y]) \cdot Y$ and $\beta = 1/E[Y]$. (We remark that the MULT method was originally proposed for the single-period problem by (Ehrhardt and Taube 1987).) Note that in both of the above methods, the inflation factor β is set to the reciprocal of the expected yield factor $E[Y]$. Computation results in (Bollapragada and Morton 1999) and (Inderfurth and Transchel 2007) show that NLH1 performs reasonably well when $b/(b+h) \leq 0.95$. (Zipkin 2000), similar to our approach, discusses the question of picking the best θ and β . (To our knowledge, he coined the term “linear inflation rule”, and first considered the possibility of β being values other than $1/E[Y]$.) He analyzes the behavior of the ordering quantity by calculating (approximately) its first two moments, and then uses a Normal approximation to propose reasonable values for β and θ . Our approach, as will be evident, is different.

Let $\mathcal{S}(\theta, \beta)$ be the inventory system managed by the LIR policy with parameters θ and β assuming that the starting inventory level is θ . (Such an assumption on the starting inventory level has been used in the inventory literature without random yield, e.g., (Glasserman and Tayur 1995) and (Janakiraman and Roudy 2004). It facilitates the derivation of analytic results. The analytic results of this note are based on and enabled by this specific assumption of the starting inventory level.) Under the infinite-horizon average cost criterion, the starting inventory level does not affect the objective function if the inventory levels in the $\mathcal{S}(\theta, \beta)$ system, for a given pair of θ and β , converge to a steady-state distribution. Unfortunately, this condition does not hold in general for the random yield problem. However, let us suppose that demands are integer-valued, and consider a *discretized* approximate problem where the yield quantity $Y_t \cdot Q_t$ is probabilistically rounded to an integer. Then, for this approximate problem, if $P[D = 0], P[D = 1] > 0$ and $P[\beta \cdot Y \geq 1] > 0$, then the time to reach the state of *having θ units in inventory after production* is finite, and the cost incurred in these initial periods does not affect the infinite-horizon average cost.) We sometimes use the superscript (θ, β) on the variables to denote the inventory system. Thus for instance, $Q_t^{(\theta, \beta)}$ denotes the order quantity in period t in the $\mathcal{S}(\theta, \beta)$ system. Let $C(\theta, \beta)$ be the expected long-run average of the $\mathcal{S}(\theta, \beta)$ system, i.e.,

$$C(\theta, \beta) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T C_t(\theta, \beta) \quad \text{where}$$

$$C_t(\theta, \beta) = \left\{ h \cdot E[IN_t^{(\theta, \beta)}]^+ + b \cdot E[IN_t^{(\theta, \beta)}]^- \right\} .$$

Our heuristic approach seeks to find the optimal choice of the parameters in the LIR policy such that $C(\theta, \beta)$ is minimized.

3 ANALYSIS

In this section, we discuss how to find values of θ and β in the LIR policy that yield good approximations to the random yield problem. In Section 3.1, we show that $C(\theta, \beta)$ is convex with respect to θ ; thus, it is easy to find the optimal choice of θ for any given β . In Section 3.2, we discuss several candidates for the choice of β .

3.1 Convexity of $C(\theta, \beta)$ with Respect to θ : Fixed β

In the following discussion, we first fix the value of $\beta > 0$. (Later, we explore properties of the cost function $C(\theta, \beta)$ with respect to both its arguments. Moreover, we discuss our procedure behind choosing β optimally.) The main result of this section is to show that $C(\theta, \beta)$ is convex with respect to θ (Lemma 1). The proof of this result is based on an observation that two inventory systems with the same β parameter but different θ parameters are translations of each other. We note that (Henig and Gerchak 1990) demonstrate convexity of the dynamic programming cost-to-go function with respect to the order quantity. This however does not imply our result.

Proposition 1 For any pair of LIR parameters θ and β , the following statements hold for any sequence of demand and random yield factor realizations: for each $t \geq 1$, (i) $I_t^{(\theta, \beta)} = I_t^{(0, \beta)} + \theta$, (ii) $Q_t^{(\theta, \beta)} = Q_t^{(0, \beta)}$, and (iii) $IN_t^{(\theta, \beta)} = IN_t^{(0, \beta)} + \theta$.

Proof. Recall from the definition of the $\mathcal{S}^{(\theta, \beta)}$ system that $I_1^{(\theta, \beta)} = \theta$ and $I_1^{(0, \beta)} = 0$. In the first period $t = 1$, we have $I_1^{(\theta, \beta)} - I_1^{(0, \beta)} = \theta$. Since $\theta - I_1^{(\theta, \beta)} = 0 - I_1^{(0, \beta)}$, it follows that $Q_1^{(\theta, \beta)} = Q_1^{(0, \beta)}$ from the definition of the LIR policy. Thus,

$$\begin{aligned} IN_1^{(\theta, \beta)} &= I_1^{(\theta, \beta)} + Y_1 \cdot Q_1^{(\theta, \beta)} - D_1 \\ &= I_1^{(0, \beta)} + \theta + Y_1 \cdot Q_1^{(0, \beta)} - D_1 = IN_1^{(0, \beta)} + \theta . \end{aligned}$$

Thus, we have proved the required result for $t = 1$. Since $I_{t+1}^{(\theta, \beta)} = IN_t^{(\theta, \beta)}$ and $I_{t+1}^{(0, \beta)} = IN_t^{(0, \beta)}$, the remainder of the proof follows easily by induction using a similar argument. \square

Lemma 1 For any fixed β , $C(\theta, \beta)$ is convex with respect to θ .

Proof. For any θ , compare the two inventory systems $\mathcal{S}(\theta, \beta)$ and $\mathcal{S}(0, \beta)$. From Proposition 1, we have $IN_t^{(\theta, \beta)} = IN_t^{(0, \beta)} + \theta$ for each $t \geq 1$. The single-period

costs in period t satisfy

$$\begin{aligned} &h \cdot [IN_t^{(\theta, \beta)}]^+ + b \cdot [IN_t^{(\theta, \beta)}]^- \\ &= (h + b) \cdot [IN_t^{(\theta, \beta)}]^+ - b \cdot IN_t^{(\theta, \beta)} \\ &= (h + b) \cdot [IN_t^{(0, \beta)} + \theta]^+ - b \cdot IN_t^{(0, \beta)} - b \cdot \theta , \end{aligned}$$

which is convex with respect to θ . Thus, $C(S, \beta)$ is convex with respect to S . \square

Lemma 1 shows that finding the optimal choice of θ for a given value of β is a convex function minimization problem. It can be shown that $C(\theta, \beta) \rightarrow \infty$ as either $\theta \rightarrow \infty$ or $\theta \rightarrow -\infty$. Thus, $\min_{\theta} C(\theta, \beta)$ is attained, and we denote its minimizer by $\theta^*(\beta)$. Minimizing this convex function can be done efficiently, for instance, using bisection methods, given that there is a fast oracle for evaluating the cost function $C(\theta, \beta)$. Though it is difficult to derive an analytic expression for $C(\theta, \beta)$, one can compute it using simulation. Thus, one approach to minimize $C(\cdot, \beta)$ is to choose the candidates of θ iteratively using a convex programming technique, and evaluate the simulated cost for each θ . We note that extensive numerical simulations seem to indicate that $C(\theta, \beta)$ is in general not jointly convex but $C(\theta^*(\beta), \beta)$ is convex in β .

However, the following observation enables us to design a faster simulation optimization method for obtaining $\theta^*(\beta)$ without any explicit search. It shows that the optimal choice of θ must satisfy a newsvendor-like condition. Thus, one needs to simulate $C(0, \beta)$ exactly once and then choose θ to satisfy an optimality condition as in the following lemma. Recall $C(\theta, \beta) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T C_t(\theta, \beta)$.

Lemma 2 For any β and T , a minimizer of $\frac{1}{T} \sum_{t=1}^T C_t(\cdot, \beta)$ is given by

$$\inf \left\{ \theta : \frac{1}{T} \sum_{t=1}^T P[IN_t^{(0, \beta)} + \theta \leq 0] \leq h/(b+h) \right\}$$

Proof. Note that

$$\begin{aligned} &\sum_{t=1}^T C_t(\theta, \beta) \\ &= \sum_{t=1}^T h \cdot E[IN_t^{(\theta, \beta)}]^+ + b \cdot E[IN_t^{(\theta, \beta)}]^- \\ &= \sum_{t=1}^T h \cdot E[IN_t^{(0, \beta)}] + (b+h) \cdot E[IN_t^{(\theta, \beta)}]^- \\ &= \sum_{t=1}^T h \cdot E[IN_t^{(0, \beta)} + \theta] + (b+h) \cdot E[IN_t^{(0, \beta)} + \theta]^- . \end{aligned}$$

By differentiating it with respect to θ , we obtain

$$\frac{d}{d\theta} \sum_{t=1}^T C_t(\theta, \beta) = T \cdot h - (b+h) \cdot \sum_{t=1}^T P[IN_t^{(0,\beta)} + \theta \leq 0].$$

By setting the above expression to zero, we obtain the result. \square

We note that in both MULT and NLH1, β is fixed at $1/E[Y]$, and in each case θ assumes a pre-defined ad-hoc value. Instead, the results of this section show that the optimal choice of θ can be obtained easily for any fixed β without much computational effort. We show computationally that this simple calculation sometimes leads to significant improvements as compared to MULT and NLH1.

Our approach and analysis in this section are quite general and are applicable to other models of random yield. For example, the results in Lemmas 1 and 2 can be generalized to the case of a positive replenishment lead-time, regardless of whether the yield is realized as soon as the order is placed or at the time of delivery (or sometime in between). They can be also extended to the case with a finite capacity on the order quantities, to a serial production system, and to any yield model where the distribution of the quantity delivered in period t depends only on the order quantities $\{Q_1, Q_2, \dots, Q_t\}$. In many such extensions, finding the optimal replenishment policy is prohibitively difficult as the number of states in the dynamic programming formulation explodes quickly; thus, our approach of optimizing within the class of LIR policies may be useful in finding an efficient heuristic.

3.2 Optimizing $C(\theta^*(\beta), \beta)$

In Section 3.1, we have shown that it is easy to compute $\theta^*(\beta)$ for a given value of β . Minimizing $C(\theta, \beta)$ is equivalent to minimizing $C(\theta^*(\beta), \beta)$, a single-dimensional problem with respect to β . While computational results in Section 4 indicate the convexity of $C(\theta^*(\beta), \beta)$, we are unable to prove structural results (such as the convexity or unimodality of $C(\theta^*(\beta), \beta)$ with respect to β).

In this section, we suggest some reasonable choices for the inflation factor β . Such choices can be either directly employed, or used as initial solutions for a single-dimension minimization procedure.

1. CHOICE A: $\beta_A = 1/E[Y]$. (We use the subscript A to distinguish it from other choices.) Both MULT and NLH1 use this value of β . It is the optimal choice if the yield factor distribution is deterministic. Note that β_A does not depend on the variability of Y or on the demand distribution D .

2. CHOICE B: Define

$$\beta_B = \sup\{\beta : E[Y \cdot \mathbb{I}[\frac{1}{\beta} \leq Y]] \leq \frac{b}{b+h} \cdot E[Y]\}.$$

Note that while β_A depends solely on the expected value of the yield distribution Y , β_B depends on the distribution of Y , as well as the ratio between b and h . The choice of β we introduce here is motivated by the deterministic version of the demand distribution. Consider a single-period problem with the starting inventory level I and the known demand D . The following proposition, which we state without proof, shows that the optimal policy for this problem is an LIR policy, where the target inventory level is D and the inflation factor is given by β_B .

Proposition 2 Consider a single-period problem where D is the known deterministic demand and I is the starting inventory level. Then, the order quantity minimizing the expected holding and backlogging cost is given by $\max\{\beta_B \cdot (D - I), 0\}$.

The following proposition establishes an upper bound and a lower bound on β_B .

Proposition 3

- (i) $1/\beta_B$ is bounded below by the $h/(b+h)$ -fractile of Y .
- (ii) Suppose that the probability density function of Y is symmetric and quasi-concave. If $b/(b+h) \geq 0.75$, then $\beta_B \geq \beta_A$.

Proof. (i) Let $\bar{Q} = \inf\{Q : E[\mathbb{I}[Q \leq Y]] \leq b/(b+h)\}$. For any $Q < \bar{Q}$, we have $E[\mathbb{I}[Q \leq Y]] > b/(b+h)$, and thus

$$\begin{aligned} \frac{b}{b+h} \cdot E[Y] &< E[\mathbb{I}[Q \leq Y]] \cdot E[Y] \\ &\leq \frac{E[Y \cdot \mathbb{I}[Q \leq Y]]}{E[Y]} \cdot E[Y] = E[Y \cdot \mathbb{I}[Q \leq Y]], \end{aligned}$$

implying $1/\beta_B \geq Q$. Since this result holds for any $Q < \bar{Q}$, it follows $1/\beta_B \geq \bar{Q}$.

(ii) From the definition of β_A and β_B , it suffices to show $E[Y \cdot \mathbb{I}[E[Y] \leq Y]] \leq b/(b+h) \cdot E[Y]$. Let f_Y denote the density function of Y . Since Y is nonnegative and has a symmetric and quasi-concave probability density, the support of Y is bounded above by $2 \cdot E[Y]$. Let U be the uniform distribution on $[0, 2 \cdot E[Y]]$. It is easy to verify that $E[Y] = E[U]$, and that the CDF of Y coincides with the CDF of U at $E[Y]$ and $2 \cdot E[Y]$. It can also be shown that the CDF of Y is bounded below by the CDF of U in the interval $(E[Y], 2 \cdot E[Y])$. (To see this, suppose by the way of contradiction that there exists $y \in (E[Y], 2 \cdot E[Y])$ such that $F_Y(y) < F_U(y)$, where F_Y and F_U denote the CDF's of Y and U , respectively. Then, since the derivative of F_Y is the density of Y , there exists, by the Mean Value Theorem,

$y' \in (E[Y], y)$ such that the density f_Y of Y satisfies

$$\begin{aligned} f_Y(y') &= \frac{F_Y(y) - F_Y(E[Y])}{y - E[Y]} \\ &< \frac{F_U(y) - F_U(E[Y])}{y - E[Y]} = \frac{1}{2 \cdot E[Y]} . \end{aligned}$$

Similarly, there exists $y'' \in (y, E[Y])$ such that $f_Y(y'') > 1/(2 \cdot E[Y]) > f_Y(y')$. This contradicts the condition that Y is symmetric and quasi-concave.) Then,

$$\begin{aligned} E[Y \cdot \mathbb{I}[E[Y] \leq Y]] &\leq E[U \cdot \mathbb{I}[E[U] \leq U]] \\ &= 0.75 \cdot E[U] = 0.75 \cdot E[Y] . \end{aligned}$$

The required result follows since $b/(b+h) \geq 0.75$. \square

3. CHOICE AB: $\beta_{AB} = (\beta_A + \beta_B)/2$. We combine β_A and β_B by taking their average. Numerical results suggest that this choice of β performs well; see Section 4.

4. CHOICE Z: Based on the analysis of the first two moments of the Y , (Zipkin 2000) proposes the following choice of β value: $\beta_Z = E[Y]/E[Y^2]$. Since $E[Y^2] = \text{Var}(Y) + E[Y]^2$ and the coefficient of variation for Y satisfies $\rho_Y^2 = \text{Var}(Y)/E[Y]^2$, it follows that $\beta_Z = [(1 + \rho_Y^2) \cdot E[Y]]^{-1} \leq 1/E[Y] = \beta_A$. Therefore, under the condition of Proposition 3 (ii), $\beta_Z \leq \beta_A \leq \beta_{AB} \leq \beta_B$.

5. CHOICE C: One of the heuristic methods introduced by (Bollapragada and Morton 1999) is the NLH2 method, which is *not* an LIR policy. The ordering quantity as a function of the beginning inventory level is piece-wise linear. (We describe this function in Section 4.) One of the non-zero slopes is $1/E[Y]$, and the other one is

$$\beta_C = E[Y]^{-1} \cdot [1 - s^2 \cdot \rho_Y^2 / (\rho_D^2 + \rho_Y^2)]^{-1/2}$$

where s is the $b/(b+h)$ -fractile of $D/E[D] - Y/E[Y]$, and ρ_D and ρ_Y are the coefficients of variation for D and Y , respectively. We take this slope as the inflation factor in an LIR policy.

We note that (Bollapragada and Morton 1999) do not provide any theoretical justification for this choice except that their computational experiments show good performance of NLH2. If β_C exists, then the second factor in the definition of β_C is at least 1, implying $\beta_C \geq \beta_A$. However, β_C may not be well-defined since the expression inside the square root operator in the definition of β_C may be negative, as pointed out by (Li, Xu, and Zheng 2006).

4 COMPUTATIONAL RESULTS

In this section, we summarize the computational results on the performance of LIR policies for the random yield problem. For several candidates for the inflation factor proposed in Section 3, we find the best target level. We

also find, numerically, the optimal parameters (θ^*, β^*) , and report the performance of the corresponding LIR policy. We also compare the LIR policies to other heuristics in the literature.

In our experiments, we fix $h = 1$ and choose b such that $b/(b+h)$ is one of six possible values from $\{0.85, 0.90, 0.95, 0.97, 0.99, 0.995\}$. We use two possible demand and yield distribution cases: D and Y are either both normal or both uniform. (When the distributions are normal, we truncate the random variables in our simulations such that they are nonnegative.) Since we can replace Y with $Y/E[Y]$ without changing the nature of the problem ((Henig and Levin 1992)), we assume $E[Y] = 1$. Similarly, we fix $E[D]$ at 20. We let $\rho_Y \in \{0.1, 0.2, 0.3, 0.4\}$ and $\rho_D \in \{0.2, 0.4\}$. (Since the variance of Y satisfies $\sigma_Y^2 = E[Y]^2 \cdot \rho_Y^2$, the standard deviation of Y is given by $\sigma_Y = E[Y] \cdot \rho_Y$. In particular, if Y is uniform, then it can be shown that $Y \sim \text{Uniform}[E[Y] - \Delta, E[Y] + \Delta]$ where $\Delta = \sqrt{3} \cdot \rho_Y \cdot E[Y]$ since $\sigma_Y^2 = \Delta^2/3$. A similar result holds for D .) In total, we consider $6 \cdot 2 \cdot 4 \cdot 2 = 96$ combinations. Our test scenarios are based on (Bollapragada and Morton 1999) and (Inderfurth and Transchel 2007); (Bollapragada and Morton 1999) use the $b/(b+h)$ ratios of $\{0.85, 0.90, 0.95\}$ only, and (Inderfurth and Transchel 2007) consider the uniform distribution case only.

We briefly describe our methodology. To evaluate $C(\theta, \beta)$ for any given θ and β , we simulate the inventory system $\mathcal{S}(\theta, \beta)$ for $T^\circ + T$ periods, and use the last T periods to compute the time-average cost for this sample path. We replicate the experiment N times, and obtain the mean of the average costs, which we denote by $\hat{C}(\theta, \beta)$. We let $T^\circ = 2000$, $T = 5000$, and $N = 2000$.

To find the optimal $\theta^*(\beta)$ for any given β , we simulate $\mathcal{S}(0, \beta)$ for $T^\circ + T$ periods, replicating the experiment N times. From each replication, we collect the ending inventory levels IN_t of the last T periods, i.e., $t = T^\circ + 1, \dots, T^\circ + T$; thus, there are a total of $T \cdot N$ samples of ending inventories. We find θ such that $-\theta$ is the $h/(b+h)$ -fractile of the sample distribution (see Lemma 2). We denote it by $\hat{\theta}(\beta)$, and use it as an estimate for $\theta^*(\beta)$. Once we obtain $\hat{\theta}(\beta)$, we estimate $\hat{C}(\hat{\theta}(\beta), \beta)$.

To find the optimal choice of the parameters, we use MATLAB's minimization function to find $\hat{\beta}$ minimizing $\hat{C}(\hat{\theta}(\beta), \beta)$. Let $\hat{\theta} = \hat{\theta}(\hat{\beta})$. In our computational results, $\hat{C}(\hat{\theta}(\beta), \beta)$ appears to be a convex function of β . As discussed in Section 3.2, the convexity or the unimodality of $C(\theta(\cdot), \cdot)$ is an open research question. However, our experiments indicate that $C(\theta, \beta)$ is not jointly convex.

We have used common number variables to reduce the variance of the simulated results. For each combination of ρ_D , ρ_Y and $b/(b+h)$, we have used the same realized $N \cdot (T + T^\circ)$ samples of D and Y to compute the results for all heuristics. We found that the half-width of the 95% confidence interval for each $C(\theta(\beta), \beta)$ was less than

approximately 0.20% of the corresponding sample average $\hat{C}(\hat{\theta}(\beta), \beta)$.

We compare the performance of LIR policies with other heuristics in (Bollapragada and Morton 1999).

- NLH2. The order quantity is a piece-wise linear function with two non-zero slopes. More precisely,

$$Q(I) = \frac{E[D]}{E[Y]} + \frac{s \cdot E[D] - I}{E[Y] \sqrt{1 - s^2 \cdot \rho_Y^2 / (\rho_D^2 + \rho_Y^2)}}$$

if $I \leq E[D]$,

$$Q(I) = Q(E[D]) - (I - E[D]) / E[Y]$$

if $E[D] < I \leq E[D] + E[Y]Q(E[D])$, and $Q(I) = 0$ otherwise, where s is the $b/(b+h)$ -fractile of $D/E[D] - Y/E[Y]$. Note that $Q(I)$ is piece-wise linear, with non-zero slopes being β_A and β_C . As discussed in Section 3.2, the above function may not be well-defined.

- Newsvendor Heuristic (NH). (Bollapragada and Morton 1999) give the expression for the case of the normal distribution only. Let v be the $b/(b+h)$ -fractile of standard normal distribution. If $I < E[D] \cdot (1 + v \cdot \rho_D)$, then

$$Q(I) = \frac{E[D]}{E[Y]} \cdot \frac{(1 - \frac{I}{E[D]}) + v \sqrt{(1 - \frac{I}{E[D]})^2 \rho_Y^2 + (1 - v^2 \rho_Y^2) \rho_D^2}}{1 - v^2 \rho_Y^2};$$

otherwise, $Q(I) = 0$. As in NLH2, the $Q(\cdot)$ function may not be well-defined.

In the case of uniform distribution, we numerically find $Q(I)$ such that equation (12) of (Bollapragada and Morton 1999) is satisfied, i.e., $P[D \leq I + Y \cdot Q(I)] = b/(b+h)$.

The computational results are reported in the full version of this paper. For each combination of ρ_D , ρ_Y and the $b/(b+h)$ -ratio, we report compare the cost of the best LIR policy $C(\hat{\theta}, \hat{\beta})$ to the optimal costs reported in the appendix of (Bollapragada and Morton 1999); we caution that these ‘‘optimal costs’’ correspond to discretized problems and thus they differ from the exact optimal costs. Overall, the numerical experiments indicate that the class of LIR policies perform fairly well.

We also examine the performance of other LIR policies and non-LIR policies; the cost associated with these policies are presented in comparison to the cost of the best LIR policy. A positive number represents a cost worse than the best

LIR policy, and a negative number represents a cost better than the best LIR policy. As expected, the best LIR policy outperforms any other LIR policy; the reported costs of MULT and NLH1 exceed the cost associated with the best LIR policy on average by 79.2% and 5.0%, respectively. Compared to the non-LIR methods, the best LIR policy outperforms NH in all of the cases, and reports worse results than NLH2 only in three cases; on average, it is better than NH and NLH2 by 6.5% and 1.4%, respectively. (We note that NH actually performs quite well when ρ_Y is small, but its performance deteriorates as ρ_Y increases.) We stress that both NH and NLH2 are not always well-defined, and is thus questionable whether these heuristics can be broadly applicable.

We also observe that while all of β_Z , β_A , β_B and β_C are reasonably good choices of β (having costs exceeding the best LIR policy by 4.2%, 2.8%, 4.7% and 1.3%, respectively), β_{AB} displays the strongest performance (0.6%).

In summary, we recommend, in all cases, to always use the LIR heuristic with the optimal value of the parameter β . However, whenever this is not feasible, a very good alternative is to always use the LIR policy with the choice of β_{AB} . Our results indicate that this is consistently a good approach, independent of service levels and variability.

ACKNOWLEDGEMENT

The authors are thankful to Jonathan J. Tay, a Masters student at Columbia University, who conducted most of the reported computational results. We also thank Paul Zipkin, Ganesh Janakirman, two referees, and the associate editor for helpful comments on an earlier version of this paper. The first author was supported partially by NSF grant DMS-0732169.

REFERENCES

- Bollapragada, S., and T. E. Morton. 1999. Myopic heuristics for the random yield problem. *Operations Research* 47 (5): 713–722.
- Ehrhardt, R., and L. Taube. 1987. An inventory model with random replenishment quantities. *International Journal of Production Research* 25:1975–1803.
- Glasserman, P., and S. Tayur. 1995. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science* 41 (2): 263–281.
- Henig, M., and Y. Gerchak. 1990. The structure of periodic review policies in the presence of random yield. *Operations Research* 38 (4): 634–643.
- Henig, M., and N. Levin. 1992. Joint production planning and product delivery commitment with random yield. *Operations Research* 40 (2): 404–408.

- Inderfurth, K., and S. Transchel. 2007. Note on “myopic heuristics for the random yield problem”. *Operations Research* 55 (6): 1183–1186.
- Janakiraman, G., and R. O. Roundy. 2004. Lost-sales problems with stochastic lead times: Convexity results for base-stock policies. *Operations Research* 52 (5): 795–803.
- Li, Q., H. Xu, and S. Zheng. 2006. Periodic review inventory systems with random yield: Bounds and heuristics. *Working Paper, HKUST*.
- Yano, C. A., and H. L. Lee. 1995. Lot sizing with random yields: A review. *Operations Research* 43 (2): 311–334.
- Zipkin, P. H. 2000. *Foundations of inventory management*. The McGraw-Hill Companies, Inc.

AUTHOR BIOGRAPHIES

WOONGHEE TIM HUH is an assistant professor in the Department of Industrial Engineering and Operations Research at Columbia University. He received a Ph.D. degree in Operations Research at Cornell Research, where his research was supported by Semiconductor Research Corporation. His email address is <th2113@columbia.edu>.

MAHESH NAGARAJAN is an assistant professor in Operations and Logistics Division, Sauder School of Business, University of British Columbia. He received a Ph.D. degree from Marshall Business School, University of Southern California. His email address is <mahesh.nagarajan@sauder.ubc.ca>.