1 Stochastic Demand

In this section we discuss the problem of controlling the inventory of a single item with stochastic demands. We start by studying the single period problem, also known as the Newsvendor Problem and then extend it to multi-period and infinite-horizon problems with and without setup costs.

1.1 The Newsvendor Problem

Let D denote the one period random demand, with mean $\mu = E[D]$ and variance $\sigma^2 = \text{Var}[D]$. Let c be the unit cost, $c(1+m)$ the selling price and $c(1-d)$ the salvage value. You can interpret m as the retail markup and d as the salvage discount. If Q units are ordered, then $\min(Q, D)$ units are sold and $(Q-D)_+$ units are salvaged. [1] The profit is given by $c(1+m)\min(Q, D)+c(1-d)(Q-D)_+ - cQ$. Taking expectations we find the expected profit:

$$\pi(Q) = c(1+m)E\min(Q, D) + c(1-d)E(Q-D)_+ - cQ$$

Using the fact that $\min(Q, D) = D - (D-Q)_+$ we can write the expected profit as

$$\pi(Q) = cm\mu - G(Q)$$

where

$$G(Q) = cdE(Q-D)_+ + cmE(D-Q)_+ \geq 0.$$

This allow us to view the problem of maximizing $\pi(Q)$ as the problem of minimizing the cost $G(Q)$.

For convenience let $h = cd$ and $p = cm$. It is convenient to think of h as the per unit overage cost and of p as the per unit underage cost. Sometimes the underage cost is inflated to take into account the *ill-will cost* associated with unsatisfied demand. Later, in the study of multi-period problems, we will call h the holding cost rate and p the shortage cost rate per. Let $g(x) = hx_+ + px_-$, then $G(Q)$ can be written as $G(Q) = E[g(Q-D)]$. Since g is convex and convexity is preserved by linear transformations and by the expectation operator it follows that G is also convex.

Let $G^{det}(Q) = h(\mu-Q)_+ + p(Q-\mu)_+$. This represents the cost when D is deterministic. Clearly $Q = \mu$ minimizes G^{det} and $G^{det}(\mu) = 0$, so $\pi^{det}(\mu) = cm\mu$. By Jensen's inequality $G(Q) \geq G^{det}(Q)$. As a result, $\pi(Q) \leq \pi^{det}(Q) \leq \pi^{det}(\mu) = cm\mu$.

If the distribution of D is continuous, we can find an optimal solution by taking the derivative of G and setting it to zero. Since we can interchange the derivative and the expectation operators, it follows that $G'(Q) = hE\delta(Q-D) - pE\delta(D-Q)$ where $\delta(x) = 1$ if $x \geq 0$ and zero otherwise. Consequently,

$$G'(Q) = hPr(Q-D \geq 0) - pPr(D-Q \geq 0).$$

Setting the derivative to zero reveals that

$$Pr(D \leq Q) = \beta, \tag{1}$$

where $\beta = \frac{p}{h+p}$.

If F is continuous then there is at least one Q satisfying Equation (1). We can select the smallest such solution by letting

$$Q^* = \inf\{Q \geq 0 : Pr(D \leq Q) \geq \beta\}. \tag{2}$$

[1] We will use $x_+ = x^+ = \max(x, 0)$ and $x_- = x^- = \max(-x, 0)$ to denote the positive and the negative part of a number.

If F is strictly increasing then F has an inverse and there is a unique optimal solution given by

$$Q^* = F^{-1}(\beta). \tag{3}$$

In practice, D often takes values in the set of natural numbers $\mathcal{N} = \{0, 1, \ldots\}$. In this case it is useful to work with the forward difference $\Delta G(Q) = G(Q+1) - G(Q)$, $Q \in \mathcal{N}$. By writing $E(D-Q)_+ = \sum_{j=Q}^{\infty} Pr(D > j)$, it is easy to see that

$$\Delta G(Q) = h - (h+p)Pr(D > Q)$$

is non-decreasing in Q, and that $\lim_{Q\to\infty} \Delta G(Q) = h > 0$, so an optimal solution is given by $Q = \min\{Q \in \mathcal{N} : \Delta G(Q) \geq 0\}$, or equivalently,

$$Q^* = \min\{Q \in \mathcal{N} : Pr(D \leq Q) \geq \beta\}, \tag{4}$$

The origin of the Newsvendor model appears to date back to the 1888 paper by Edgeworth [2] who used the Central Limit Theorem to determine the amount of cash to keep at a bank to satisfy random cash withdrawals from depositors. The fractile solution (1) appeared in 1951 in the classical paper by Arrow, Harris and Marchak [1].

The newsvendor solution can be interpreted as providing the smallest supply quantity that guarantees that all demand will be satisfied with probability at least $100\beta\%$. Thus, the profit maximizing solution results in a service level $100\beta\%$. In practice, managers often specify β and then find Q accordingly. The service level measure implied by the Newsvendor problem should not be confused with the fraction of demand served, or fill-rate, which is defined as $\alpha = E\min(D, Q)/ED$.

1.2 Normal Demand Distribution

An important special case arises when the distribution D is normal. The normal assumption is justified by the Central Limit Theorem when the demand comes from many different independent or weakly dependent customers. If D is normal, then we can write $D = \mu + \sigma Z$ where Z is a standard normal random variable. Let $\Phi(z) = Pr(Z \leq z)$ be the cumulative distribution function of the standard normal random variable. Although the function Φ is not available in closed form, it is available in Tables and also in electronic spreadsheets. Let z_β be such that $\Phi(z_\beta) = \beta$. In Microsoft Excel, for example, the command NORMSINV(0.75) returns 0.6745 so $z_{.75} = 0.6745$. Since $Pr(D \leq \mu + z_\beta\sigma) = \Phi(z_\beta) = \beta$, it follows that

$$Q^* = \mu + z_\beta\sigma \tag{5}$$

satisfies Equation (3), so Equation (5) gives the optimal solution for the case of normal demand. The quantity z_β is known as the safety factor and $Q^* - \mu = z_\beta\sigma$ is known as the safety stock.

It can be shown that $E(D-Q^*)_+ = \sigma E(Z - z_\beta)_+ = \sigma[\phi(z_\beta) - (1-\beta)z_\beta]$ where ϕ is the density of the standard normal random variable. As a consequence,

$$\begin{aligned} G(Q^*) &= hE(Q^* - D)_+ + pE(D - Q^*)_+ \\ &= h(Q^* - \mu) + (h+p)E(D - Q^*)_+ \\ &= hz_\beta\sigma + (h+p)\sigma E(Z - z_\beta)_+ \\ &= hz_\beta\sigma + (h+p)\sigma[\phi(z_\beta) - (1-\beta)z_\beta] \\ &= (h+p)\sigma\phi(z_\beta), \end{aligned}$$

so

$$\begin{aligned} \pi(Q^*) &= cm\mu - (h+p)\sigma\phi(z_\beta) \\ &= cm\mu - c(d+m)\sigma\phi(z_\beta). \end{aligned}$$

In addition, since $E\min(D, Q^*) = ED - E(D-Q^*)_+$, we can divide by ED and write the fill-rate as

$$\alpha = 1 - \text{cv}[\phi(z_\beta - (1-\beta)z_\beta]$$

where $\text{cv} = \sigma/\mu$ is the coefficient of variation of demand. Since $\phi(z_\beta - (1-\beta)z_\beta \geq 0$ is decreasing in β, it follows that the α is *increasing* in β and *decreasing* in cv. Numerical results show that $\alpha \geq \beta$ for all reasonable values of cv, including $\text{cv} \leq 1/3$, which is about the highest cv value for which the normal model is appropriate.

Example Normal Demand: Suppose that D is normal with mean $\mu = 100$ and standard deviation $\sigma = 20$. If $c = 5$, $h = 1$ and $p = 3$, then $\beta = 0.75$ and $Q^* = 100 + 0.6745 * 20 = 113.49$. Notice that the order is for 13.49 units (safety stock) more than the mean. Typing NORMDIST(.6574,0,1,0) in Microsoft Excel, returns $\phi(.6745) = 0.3178$ so $G(113.49) = 4 * 20 * .3178 = 25.42$, and $\pi(113.49) = 274.58$, with $\alpha = 97\%$.

Table 1 gives z_β, $\phi(z_\beta)$ and α (at cv = .2) for different values of β.

$100\beta\%$	z_β	$\phi(z_\beta)$	$100\alpha\%$
50%	0	0.3989	92.0%
75%	0.6745	0.3178	97.0%
90%	1.2816	0.1755	99.1%
95%	1.6499	0.1031	99.6%
97.5%	1.9600	0.0584	99.8%
99%	2.3263	0.0267	99.9%

Table 1: *Normal solution for several values of* β

1.3 Poisson Distribution

Another distribution that arises often in practice is the Poisson distribution. D is said to be Poisson with parameter $\lambda > 0$ if

$$Pr(D = k) = \exp(-\lambda)\frac{\lambda^k}{k!} \quad k = 0, 1, 2, \ldots$$

The Poisson distribution arises as a limit of the binomial distribution with large n and small p via the relationship $\lambda = np$. For example, the number of customers that enter a store and make a purchase can often be modeled as a Poisson distribution. It is well known that $\mu = \lambda$ and $\sigma = \sqrt{\lambda}$ so the coefficient of variation σ/μ becomes small for large λ. When λ is large, the Poisson distribution can be approximated by the Normal distribution with mean $\mu = \lambda$ and standard deviation $\sigma = \sqrt{\lambda}$.

The following recursions, starting from $Pr(D = 0) = e^{-\lambda}$ and $E[D] = \lambda$, are useful in tabulating and solving problems involving the Poisson distribution:

$$\begin{aligned} Pr(D = k) &= Pr(D = k-1)\lambda/k, \quad k = 1, 2, \ldots \\ Pr(D \leq k) &= Pr(D \leq k-1) + Pr(D = k), \quad k = 1, 2, \ldots \\ E[(D-k)_+] &= E[(D-k+1)_+] - Pr(D \geq k) \quad k = 1, 2, \ldots \end{aligned}$$

An optimal value of Q is given by the smallest integer such that $P(D \leq Q) \geq \beta$.

Example Poisson: If D is Poisson with parameter $\lambda = 25$, and $c = 5$, $h = 1$ and $p = 3$, then $\beta = 0.75$ and $Q^* = 28$ is optimal. To compute $G(Q^*)$ notice that $G(Q) = h(Q-\lambda)+(h+p)E(D-Q)_+$, so $G(28) = 6.48$. Table 2 provides some of the values associated with the Poisson distribution. At $Q = 28$, $E(D-28)_+ = 0.87$ so $\alpha = 1 - 0.87/25 = .97$.

k	$Pr(D=k)$	$Pr(D\leq k)$	$E(D-k)_+$	$G(k)$	α
22	0.07	0.32	3.80	12.21	.85
23	0.08	0.39	3.12	10.48	.88
24	0.08	0.47	2.51	9.06	.90
25	0.08	0.55	1.99	7.95	. 92
26	0.08	0.63	1.54	7.16	.94
27	0.07	0.70	1.17	6.68	.95
28	0.06	*0.76*	0.87	*6.48*	.97
29	0.05	0.82	0.63	6.54	.97
30	0.05	0.86	0.45	6.81	.98
31	0.04	0.90	0.32	7.26	.99
32	0.03	0.93	0.22	7.86	.99
33	0.02	0.95	0.14	8.57	.99
34	0.02	0.97	0.09	9.38	.99

Table 2: *Solution to Newsvendor with Poisson Demand*

1.4 Random Number of Customers

A more general demand model arises when the number of customers, say N, is itself a non-negative random variable taking integer values and each customer demands a random number of units. If the customer demands are IID, then we can model the total demand as

$$D=\sum_{k=1}^{N} X_k.$$

Notice that we recover $D=N$ if $X_k=1$ with probability one.

Using well known results of conditional expectations (see page 153 in reference [8]) it follows that:

$$E[D]=E[E[D|N]] \quad \text{and} \quad \text{Var}[D]=\text{Var}[E[D|N]]+E[\text{Var}[D|N]].$$

If $\mu_n=E[N], \sigma_n^2=\text{Var}[N]$, $\mu_x=E[X]$ and $\sigma_x^2=\text{Var}[X]$, then

$$E[D]=\mu_n\mu_x \quad \text{and} \quad \text{Var}[D]=\mu_x^2\sigma_n^2+\mu_n\sigma_x^2.$$

A little algebra reveals that the coefficient of variation of demand is given by

$$cv_d=\sqrt{cv_n^2+\frac{1}{\mu_n}cv_x^2}.$$

Since cv_d is decreasing in μ_n so, everything else being equal, it is better to have a large number of small customers than to have a small number of large customers. As an example, suppose that the average demand is $\mu_n\mu_d=100$, that $cv_x^2=0.3$ and that $cv_n^2=0.2$. Then $cv_d=0.202237$ if $\mu_x=1$ and $cv_d=0.360551$ if $\mu_x=100$. Since inventory related costs (overage and underage) are proportional to the standard deviation of demand, the cost of dealing with a small number of large customers can be significantly higher, 78% higher in this example, than the cost of dealing with a large number of small customers.

If N is Poisson with parameter λ, then D has a compound Poisson distribution and

$$E[D]=\lambda\mu_x \qquad \text{Var}[D]=\lambda(\mu_x^2+\sigma_x^2).$$

Notice that the coefficient of variation for the compound Poisson distribution

$$cv_d=cv_n\sqrt{1+cv_x^2}\geq cv_n=\frac{1}{\sqrt{\lambda}}.$$

1.5 The Lognormal Approximation

When the coefficient of variation σ/μ is large, neither the Normal nor the Poisson distributions are appropriate. The Normal is not appropriate because when σ/μ is large, it assigns a significant probability to negative demands. The Poisson is not appropriate because $\sigma = \sqrt{\mu}$ so the coefficient of variation is small for most reasonable values of λ. The Lognormal distribution provides, in many cases, an adequate distribution that allows closed form solutions when the coefficient of variation is large.

A random variable D is said to have the lognormal distribution, with parameters ν and τ, if $\ln(D)$ has the normal distribution with mean ν and standard deviation $\tau \geq 0$. The lognormal distribution is often used to model non-negative random variables such as lifetimes and total returns. It is well known that $E(X^n) = \exp(n\nu + n^2\tau^2/2)$. Thus, $\mu = \exp(\nu + \tau^2/2)$ and $\sigma^2 = \mu^2(\exp(\tau^2) - 1)$, so $\nu = \ln \mu - \ln \sqrt{1 + cv^2}$ and $\tau = \sqrt{\ln(1 + cv^2)}$.

The solution to the Newsvendor problem under the lognormal distribution is given by

$$Q^* = \exp(v + \tau z_\beta)$$

and

$$\pi(Q^*) = cm\mu - (h + p)\mu\Phi(\tau - z_\beta) + h\mu.$$

To see why this is true, notice that if D is lognormal then $Pr(D \leq Q^*) = Pr(\ln(D) \leq \ln(Q^*)) = Pr(\nu + \tau Z \leq \nu + \tau z_\beta) = Pr(Z \leq z_\beta) = \Phi(z_\beta) = \beta$. Now, using the fact that $E(D - Q^*)_+ = \mu\Phi(\tau - z_\beta) - Q^*\Phi(-z_\beta)$ and $\Phi(-z_\beta) = h/(h + p)$ we see that

$$\begin{aligned} G(Q^*) &= h(Q^* - \mu) + (h + p)E(D - Q^*)_+ \\ &= h(y^* - \mu) + (h + p)\mu\Phi(\tau - z_\beta) - (h + p)Q^*\Phi(-z_\beta) \\ &= (h + p)\mu\Phi(\tau - z_\beta) - h\mu. \end{aligned}$$

Example Log Normal: Figure 1 shows actual weekly demand data for a semiconductor product with $c = 5$, $p = 5$ and $h = 2$. The empirical distribution has a coefficient of variation equal to 2.22, a sample mean of 207, and a sample standard deviation equal to 459. Although close to three quarters of the demand observations were for fewer than 100 units, there is a chance of receiving a demand for over 1000 units. The Newsvendor solution based on the empirical *cdf* is $Q^* = 100$ resulting in an expected profit of \$63. If we assume demand is normally distributed with the moments calculated based on sample demand data, then the profit maximizing solution will be 467 units resulting in an expected *loss* of \$291 (based on the empirical distribution). To satisfy demand with probability 95%, management would have to order 1,400 units and incur a *loss* of \$1,583. If we use lognormal distribution with the sample moments, the profit maximizing solution will be 181 units giving us an expected profit of \$29.

1.6 Worst Case Distribution

Often there is not enough data to ascertain the form of the distribution or there may be no theoretical justification for demand to follow a particular distribution such as the Normal or the Poisson. In practice, one has to often work with guess-estimates of the mean and the forecast error or the standard deviation. Fortunately, there is a closed form formula that minimizes the function $G(Q)$ (maximizes $\pi(Q)$) against the worst possible distribution withe a given mean and a given standard deviation. This order quantity is due to Herbert Scarf [9] and it is given by

$$Q^S = \mu + \frac{\sigma}{2}\left(\sqrt{\frac{p}{h}} - \sqrt{\frac{h}{p}}\right). \tag{6}$$

Notice that Scarf's formula (6) suggests ordering more (resp., less) than the mean demand when $p > h$ (resp., $p < h$). Moreover, $|Q^S - \mu|$ increases linearly in σ for $h \neq p$.

The derivation of Scarf's formula and of other related results can be simplified by observing that $x_+ = 0.5(|x| + x)$ and then using the Cauchy-Schwarz inequality:

$$\begin{aligned} E(D-Q)_+ &= \frac{1}{2}E\{(|D-Q| + (D-Q)\} \\ &\leq \frac{1}{2}\{\sqrt{\sigma^2 + (\mu - Q)^2} + (\mu - Q)\}. \end{aligned}$$

From this, and some algebra, it follows that

$$G(Q^S) \leq \sqrt{ph}\sigma = c\sqrt{md}\sigma$$

with equality holding for a certain distribution of demand with mass concentrated at two points. As a result,

$$cm\mu \geq \pi(Q^S) \geq cm\mu - c\sqrt{md}\sigma,$$

so Scarf's ordering rule is particularly good when μ/σ is small relative to $\sqrt{m/d}$. Scarf's ordering rule is modified to $Q^S = 0$ when $\sigma/\mu > \sqrt{m/d} = \sqrt{p/h}$ reflecting the fact that it may be better not to be in business when demand is very uncertain.

It turns out that $E(D - Q^S)_+ \leq \frac{1}{2}\sigma\sqrt{\frac{h}{p}}$ so

$$\alpha = \frac{E\min(D, Q^S)}{ED} \geq 1 - \frac{1}{2}\frac{\sigma}{\mu}\sqrt{\frac{h}{p}},$$

so if the coefficient of variation is $1/4$ and $h = p$, we would have $\alpha \geq 7/8$ and if $p = 4h$, we would have $\alpha \geq 15/16$.

Finally, it is also possible to show that $G(\mu) \leq \frac{1}{2}(h + p)\sigma$, so ordering the mean results in an expected cost that is at most the *arithmetic* average of the overage and underage cost times the standard deviation of demand. Thus, in the worst case the improvement in bounds between ordering the mean and using Scarf's ordering rule is a reduction from the arithmetic to the geometric mean of h and p multiplied by the standard deviation of demand.

Example WCD vs. Normal: Consider the data used for the Normal Distribution: $\mu = 100$, $\sigma = 20$, If $c = 5$, $h = 1$ and $p = 3$. Then, $Q^S = 100 + 10(\sqrt{3} - 1/\sqrt{3}) = 111.55$, which is not too far from 113.49, the optimal order quantity under the Normal distribution.

Example WCD vs. Poisson: Consider the data used for the Poisson Distribution: $\lambda = 25$, and $c = 5$, $h = 1$ and $p = 3$. Then $Q = 25 + 2.5(\sqrt{3} - 1/\sqrt{3}) = 27.89$, which is not far from 28, the optimal order quantity under the Poisson distribution.

Example WCD vs. Lognormal: $c = 5.00$, $h = 2$, $p = 5$, $\mu = 207$, $\sigma = 459$. In this case $\sigma/\mu > \sqrt{p/h}$ so it would be best not to order if we expect the worst case distribution. The profit for not ordering will be zero assuming that $p = cm$ and no additional penalties accrue for shortages.

1.7 Random Demand at Salvage Value

Here we consider an extension where demand at the salvage price is a random variable V. Notice that the traditional newsvendor model implicitly assumes that $Pr(V \geq Q) = 1$ for all Q. The newsvendor model also implicitly assumes that $s < c$ or equivalently that the discount $d > 0$. Here we will allow $s > c$ but we will keep the assumption that $m + d > 0$.

Using the fact that $\min(D, Q) = D - (D - Q)_+$ and the fact that $\min(V, (Q - D)_+) = (Q - D)_+ - (Q - D - V)_+$ it follows that

$$\pi(Q) = cm\mu - H(Q)$$

where

$$H(Q) = G(Q) + sE(Q - D - V)_+.$$

Thus, the expected profit differs from that of the traditional Newsvendor Model only when $V \leq (Q-D)_+$, or equivalently, when $V + D \leq Q$ in that the revenue $s(Q-D-V)_+$ does not accrue. The problem of maximizing $\pi(Q)$ reduces to that of minimizing $H(Q)$. If the distributions of D and V are continuous, then

$$\begin{aligned} H'(Q) &= G'(Q) + sE\delta(Q-D-V) \\ &= h - (h+p)Pr(D > Q) + sPr(D+V \leq Q). \end{aligned}$$

It is clear that $H(Q)$ is non-decreasing in Q so $H(Q)$ is convex. Thus, a minimizer of H, say Q^*, can be found by finding a root of $H'(Q) = 0$. Let Q^{nv} be the solution to the traditional newsvendor problem. Then $H'(Q^{nv}) = sPr(D+V \leq Q^{nv}) \geq 0$, implying that there exists an optimal solution $Q^* \leq Q^{nv}$. Consequently, if $Pr(D+V \leq Q^{nv}) > 0$ then $Q^* < Q^{nv}$ so it is optimal to order fewer units than under the newsvendor model.

If D and V take integer values then it is convenient to work with the difference function $\Delta H(Q) = H(Q+1) - H(Q)$ for $Q \in \mathcal{N} = \{0, 1, \ldots\}$. To compute the $\Delta H(Q)$ first notice that

$$\begin{aligned} H(Q) &= h(Q - ED) + (h+b)E(D-Q)_+ + sE(Q-D-V)_+ \\ &= h(Q - ED) + (h+b)\sum_{j=Q}^{\infty} Pr(D > j) + s\sum_{j=0}^{Q-1} Pr(D+V \leq j). \end{aligned}$$

Consequently,

$$\Delta H(Q) = h - (h+b)Pr(D > Q) + sPr(D+V \leq Q).$$

Since $\Delta H(Q)$ is non-decreasing in Q, an optimal solution is given by

$$Q^* = \min\{Q \in \mathcal{N} : \Delta H(Q) > 0\}.$$

A revenue management problem arises when Q is fixed and $m + d < 0$, when we need to decide how many units to make available for sale at the *low* price $c(1+m)$ so that enough capacity is available at the *high* price $c(1-d)$ to maximize expected revenues.

1.8 The Forward Selling Newsvendor

Suppose that the Newsvendor computes and orders Q^* units in anticipation of random demand D from a single buyer. Under the traditional model, the newsvendor will wait until the demand D is realized before he collects any revenue. Suppose that the Newsvendor wants to reduce risk by forward selling (i.e., selling in advance of the realization of D) a certain number, say $y \leq Q^*$, units at a fair price $f(y)$. What should $f(y)$ be? Would the buyer agree to this price? Will they both be able to reduce their risk?

To determine $f(y)$ we need to find the value that makes the expected profit equivalent to that without forward selling. The revenue from forward sales will be $f(y)y$. Let $\tilde{D} = \min(D, Q^*)$ and notice that $\min(\tilde{D}, y) = \min(D, Q^*, y)$ on account of our assumption that $y \leq Q^*$. If $\tilde{D} > y$, then $(\tilde{D} - y)_+$ units will be sold at the regular price $c(1+m)$. If the seller provides $\min(D, y) = \min(\tilde{D}, y)$ units to the buyer in exchange of the payment $f(y)y$, then the inventory related costs are not affected and the expected profit is given by

$$\begin{aligned} \pi(Q^*; y) &= E\{c(1+m)(\tilde{D} - y)_+ + f(y)y + c(1-d)(Q^* - D)_+ - cQ^*\} \\ &= E\{c(1+m)\tilde{D} - c(1+m)\min(\tilde{D}, y) + f(y)y + c(1-d)(Q^* - D)_+ - cQ^*\} \\ &= E\{c(1+m)\tilde{D} + c(1-d)(Q^* - D)_+ - cQ + f(y)y - c(1+m)\min(\tilde{D}, y)\} \\ &= \pi(Q^*) + f(y)y - c(1+m)E\min(\tilde{D}, y). \end{aligned}$$

As a result, a fair forward price $f(y)$, is such that $f(y)y - c(1+m)E\min(D, y) = 0$, or equivalently by

$$f(y) = \frac{c(1+m)E\min(\tilde{D}, y)}{y} = c(1+m)E\min(\tilde{D}/y, 1). \tag{7}$$

You can think of $f(y)$ as an inverse supply function. To seller will agree to forward sell $\min(\tilde{D}, y)$ units at unit price $f(y) \leq c(1+m)$. Notice that the above formula assumes that if D, and hence $\tilde{D}$, is less than y, then the seller only needs to deliver $\tilde{D}$. The derivations would be different, as shown later, if the seller has to deliver y units even when the demand is for less than y.
Remark: In some cases the newsvendor does not have the luxury of deciding Q^*, and his starting inventory or capacity may be a fixed value Q. The analysis continues to hold with Q^* replaced by Q provided that $y \leq Q$.

1.8.1 Cost to Single Buyer

In this section we assume a single buyer with demand D. Under the traditional newsvendor arrangement, he would first observe D and then make his purchasing decision. Since the seller is providing up to Q^* units, the buyer will purchase $\tilde{D} = \min(D, Q^*)$ units from the buyer.

We will now determine the expected cost to the buyer of forward buying y units from the seller. While equation (7) represents the fair forward price from the seller's perspective it is not clear whether it will represent a fair price from the buyer's perspective. We will demonstrate that the forward price $f(y)$ is fair to the buyer if the seller delivers $\min(\tilde{D}, y)$ units at the forward price. To see this, we need to compare the cost $c(1+m)E\tilde{D}$ of buying only at the regular price and the cost $fy + c(1+m)E(\tilde{D}-y)_+$ of forward buying y units and then purchasing $(\min(D, Q^*) - y)_+$ units at the regular price $c(1+m)$. Notice that

$$\begin{aligned} f(y)y + c(1+m)E(\min(D,Q^*) - y)_+ &= c(1+m)E\min(\tilde{D}, y) + c(1+m)E(\tilde{D}-y)_+ \\ &= c(1+m)E[\min(\tilde{D}, y) + (\tilde{D}-y)_+] \\ &= c(1+m)E\tilde{D} \\ &= c(1+m)E\min(D, Q^*), \end{aligned}$$

so the expected cost is the same for the buyer.

1.8.2 Risk Reduction

The next question to investigate is risk. The buyer would be interested in the variance of his cost, while the seller would be interested in the variance of his profit. Since the random portion of the cost to the buyer is $p(\tilde{D}-y)_+$, the buyer is interested in how the variance of the the random variable $(\tilde{D}-y)_+$ changes for values of $y \leq Q^*$.

Recall that if X is a non-negative random variable then $E[X^k] = \int_0^\infty kx^{k-1}P(X > x)dx$ for all k for which the expectation exists. For $x > 0$, $P(\tilde{D}-y)_+ > x) = P(\tilde{D} > y + x) = \tilde{F}(y+x)$, where $\tilde{F}(x) = P(\tilde{D} > x)$. It follows that

$$V[(\tilde{D}-y)_+] = 2\int_y^\infty (z-y)\tilde{F}(z)dz - \left(\int_y^\infty \tilde{F}(z)dz\right)^2.$$

Consequently,

$$\frac{d}{dy}[(\tilde{D}-y)_+] = -2E[(\tilde{D}-y)_+] + 2E[(\tilde{D}-y)_+]\tilde{F}(y) = -2E[(\tilde{D}-y)_+]P(\tilde{D} \leq y).$$

This analysis shows that the risk is reduced as y increases. Notice that the derivative with respect to y vanishes at $y = Q^*$ and that at this point the variance is zero. As a result, the buyer will reduce his risk by increasing y over the range $y \leq Q^*$.

How does the risk for the seller changes with y? The profit to the seller can be written as

$$f(y)y + c(1+m)(\tilde{D}-y)_+ + c(1-d)(Q^* - \tilde{D})_+ - cQ^*.$$

As a result, the variance of the seller's profit depends also on the the covariance $\text{Cov}(\tilde{D}, (\tilde{D}-y)_+)$. The derivative of the covariance is given by $(E\tilde{D}-y)\tilde{F}(y)-E(\tilde{D}-y)_+]$, and the second derivative by $(\mu-y)\tilde{F}'(y)$. Because of the covariance term the value at which the risk for the seller is minimized is different than the value at which the risk for the buyer is minimized.

Example: $c = 10, c(1+m) = 20, c(1-d) = 4, \mu = 100$ and $\sigma = 25$, so $p = 10, h = 6$. If the distribution of D is uniform, we find $x^* = 117$. Suppose that the parties agree to a forward contract for $y = 100$ units. Then $f(y) = 17.83$, the expected profits remain the same, but the risk is reduced. The risk to the seller, measured by the standard deviation of his profit goes down to 118 from 175, and the risk to the buyer goes down from 395 to 144. A risk averse buyer would be happy to enter into this agreement because of the risk reduction effect.

If the seller *must* deliver $y \leq Q^*$ units to the buyer even when $D < y$, then his salvage revenue is reduced by $c(1-d)(y-D)_+ = c(1-d)(y-\min(D,y))$, so his expected profit is

$$\pi(Q^*; y) = \pi(Q^*) + (f - c(1-d))y - c(m+d)E\min(D,y).$$

A fair forward price in this case would be

$$f(y) = c(1-d) + c(m+d)E\min(D/y, 1). \tag{8}$$

Notice that the formula (8) reduces to (7) when the salvage value is zero, or equivalently $d = 1$, and that $f(y)$ increases with the salvage value at a rate less than one. To determine whether $f(y)$ is acceptable to the buyer, we would need to know the buyer's own salvage value. If the buyer's salvage value is larger than the seller's salvage value, then forward selling can be win-win.

In the case of multiple buyers, the seller can specify a menu of forward quantities $\{y_1, \ldots, y_m\}$ and a menu of forward payments $\{f(y_1)y_1, \ldots, f(y_m)y_m\}$ and let the buyers select their plan, much as wireless phone companies do.

2 Multi-period Models

In this section we consider a variety of multi-period models. Initially, we discuss models without setup costs and with zero lead times. Later we extend the analysis to the case of positive setup costs and positive lead times.

2.1 Finite Horizon Models

Let $D_1, \ldots, D_T$ be the demands for the next T periods. We assume that the D_t's are independent random variables, and that all stockouts are backordered. Let c_t denote the unit cost in period t, and let x_t denote the inventory level at the beginning of period t, where a positive x_t indicates that x_t units of inventory are carried from the previous period, and a negative x_t indicates that a backlog of $-x_t$ units is carried form the previous period. Let $y_t - x_t \geq 0$ denote the size of the order in period i resulting in a procurement cost $c_t(y_t - x_t)$ and an increase of the inventory level to y_t. Since D_t units are demanded during the period, the inventory level at the beginning of period $t+1$ is given by

$$x_{t+1} = y_t - D_t.$$

If $y_t = y$ the loss function in period t is given by

$$G_t(y) = h_t E(y - D_t)_+ + p_t E(D_t - y)_+ \tag{9}$$

where h_t is the overage or holding cost and p_t is the underage or backorder penalty cost. Notice that the interpretation of h_t and p_t for period $t \leq T$ is different from that of period $T+1$ in that in period $T+1$ we would typically salvage remaining items and either produce or reimburse customers if there are backlogs.

Before continuing with the formulation, we remark that the cost function G_t is convex and also coercive. A function is coercive if it grows to ∞ as y goes to $\pm\infty$. The results that we are about to obtain would continue to hold for general convex and coercive functions G_t, and are not limited to specific form of Equation (9).

Let $C_{T+1}(x_{T+1})$ be an arbitrary cost function on the inventory level at the end of period T (beginning of period $T+1$), let $0 < \alpha \leq 1$ be the one period discounted cost and let $C_t(x)$ denote the optimal expected discounted cost starting in period t with x_t units of inventory. Then,

$$C_t(x_t) = \min_{y \geq x_t} \{c_t(y - x_t) + G_t(y) + \alpha E C_{t+1}(y - D_t)\} \tag{10}$$

represents a recursive, Dynamic Programming equation that can be solved backwards starting with period T.

It can be shown that if $C_{T+1}(\cdot)$ is convex and coercive, then $C_t(\cdot)$ is convex and coercive for all $t = 1, \ldots, T$, and the optimal policy is to order $\max(0, y_t^* - x_t)$ units in period t where y_t^* minimizes

$$c_t y + G_t(y) + \alpha E C_{t+1}(y - D_t).$$

This class of policies is known as order-up-to policies. The idea is that we order up to y_t^* in period t if $x_t < y_t^*$ and do not to order otherwise.

As an example of a convex terminal cost functions C_{T+1}, consider the case where left over stock is salvaged at a unit price $c_T(1-d)$ and and backlogged sales are cancelled by at a unit cost $c_T(1+m)$ per unit. Then, $c_{T+1}(y) = -c_T(1-d)y_+ + c_T(1+m)y_-$ is convex as long as $m + d > 0$. Another form of C_{T+1} that is often used is $C_{T+1}(x) = -c_{T+1}x$. This is a situation where excess units are salvaged at c_{T+1} and excess demand is satisfied by producing at unit cost c_{T+1}.

Notice that the above problem needs to be solved recursively starting with period T down to period 1. This requires a computer code that can be written in less than one hour by an experienced programmer. The quality of the solution depends on the quality of the estimates of the data and the demand distributions.

2.1.1 The Myopic Policy

Here we describe a myopic policy that is frequently used in practice. The advantage of the myopic policy is that the computations reduce to that of solving one Newsvendor problem for each period $t = 1, \ldots, T$, and thus avoid the computational effort of solving the Dynamic Programming problem (10).

To develop this policy we need to write a slightly different but equivalent set of recursive equations. To this end let

$$M_t(x_t) = C_t(x_t) + c_t x_t.$$

$M_t(x_t)$ is the expected cost-to-go $C_t(x_t)$ starting with x_t units of inventory plus the value of the x_t units of inventory. With this definition, the recursion becomes

$$M_t(x_t) = \alpha c_{t+1}\mu_t + \min_{y \geq x_t} \left[m_t y + G_t(y) + \alpha E M_{t+1}(y - D_t)\right],$$

where $m_t = c_t - \alpha c_{t+1}$. The myopic policy ignores at time t, the future discounted costs

$$\alpha E M_{t+1}(y - D_t),$$

and orders $\max(0, y_t^m - x_t)$ units in period t, where y_t^m minimizes the *current* expected cost

$$m_t y + G_t(y).$$

If the demand is continuous, then y_t^m satisfies

$$P(D_t \leq y) = \frac{p_t - m_t}{h_t + p_t}.$$

How is the myopic policy related to the optimal policy? The most important known result is that

$$\min\{y_t^m, \ldots, y_T^m\} \leq y_t^* \leq y_t^m$$

which implies that $y_t^* = y_t^m$ whenever the y_t^m are non-decreasing.

Using Scarf's min-max approach, the myopic policy is to order $\max(0, y_t^S - x_t)$ where

$$y_t^S = \mu_t + \frac{\sigma_t}{2}\left(\sqrt{\frac{p_t - m_t}{h_t + m_t}} - \sqrt{\frac{h_t + m_t}{p_t - m_t}}\right).$$

Another way of deriving the myopic policy is to write down the total cost over the entire horizon and then separate the terms that depend on the decision made in period t. To this end, let

$$C_1(x_1) = \min_{y_t \geq x_t} E[\sum_{t=1}^{T} \alpha^{t-1}\{c_t(y_t - x_t) + G_t(y_t)\} + \alpha^T C_{T+1}(x_{T+1})]$$

Notice that y_t appears in the sum as $\alpha^t\{(c_t - \alpha c_{t+1})y_t + G_t(y_t)\}$. If $C_{T+1}(x) = -c_{T+1}x$, then can write

$$C_1(x_1) = -c_1 x_1 + \min_{y_t \geq x_t} E\sum_{t=1}^{T} \alpha^{t-1}\{(c_t - \alpha c_{t+1})y_t + G_t(y_t)\} + E\sum_{t=1}^{T} \alpha^t c_{t+1} E D_t.$$

We can now see that the myopic policy minimizes the cost function term-by-term, but ignores the possible interactions among the terms. However, if the myopic solution is such that $y_{t+1}^m \geq y_t^m - D_t$ with probability one, then the decisions in one-period do not preclude us from achieving the minimum cost in the next period, so the myopic policy is optimal in this case. This occurs, for example, if the y_t^m's are non-decreasing in t, so a natural question to ask is when can we guarantee that the y_t^m's are non-decreasing in t. One such case is when the ratios $(p_t - m_t)/(h_t + p_t)$ is independent of t and $Pr(D_t \leq y)$ is non-decreasing in t, or equivalently if the sequence of random variables is *stochastically* increasing.

2.2 Infinite Horizon, Stationary Models

If all the costs are stationary, i.e., $c_t = c, h_t = h$ and $p_t = p$ for all t, and the demands are independent and identically distributed (IID), then finite-horizon discounted costs (when $\alpha < 1$) converge, so the DP (10) can be written as

$$C(x) = \min_{y \geq x}\{c(y - x) + G(y) + \alpha E C(y - D)\}.$$

In terms of $M(x) = C(x) + cx$, the functional equation can be written as

$$M(x) = \alpha c\mu + \min_{y \geq x}\{c(1 - \alpha)y + G(y) + \alpha E M(y - D)\}.$$

The myopic policy orders $\max(0, y^m - x)$ units where y^m minimizes the current cost

$$c(1 - \alpha)y + G(y).$$

If the one period demand has a continuous distribution, then y^m satisfies

$$P(D \leq y) = \frac{p - c(1 - \alpha)}{h + p}.$$

Surprisingly, the myopic policy is optimal, under the mild assumption that D takes only non-negative values. This can be seen as follows. Suppose that $M(\cdot)$ is known and that y^* minimizes

$$c(1 - \alpha)y + G(y) + \alpha E M(y - D).$$

Then for $x \leq y^*$ we have

$$M(x) = \alpha c\mu + c(1-\alpha)y^* + G(y^*) + \alpha EM(y^* - D).$$

Notice that the right hand side of the last equation is independent of x, so there is a constant, say M^*, such that $M(x) = M^*$ for all $x \leq y^*$. Since $D \geq 0$, $y^* - D \leq y^*$ so $M(y^* - D) = M^*$. Therefore M^* satisfies

$$M^* = \alpha c\mu + c(1-\alpha)y^* + G(y^*) + \alpha M^*.$$

Solving for M^* yields

$$M^* = \frac{\alpha c\mu + c(1-\alpha)y^* + G(y^*)}{1-\alpha}$$

so y^* must minimize the current cost

$$c(1-\alpha)y + G(y)$$

just as y^m. Therefore $y^* = y^m$ if $c(1-\alpha)y + G(y)$ has a unique minimizer or we can select y^* as y^m if this function admits more than one minimizer.

Finally, notice that for $x \leq y^m = y^*$,

$$C(x^*) = M(x^*) - cx = c(y^* - \mu - x) + \frac{c\mu + G(y^*)}{1-\alpha}$$

and this can be interpreted as the cost of the safety stock $c(y^* - \mu)$ minus the cost of the inventory already available cx, plus the discounted purchasing and inventory related costs $(c\mu + G(y))/(1-\alpha)$.

The policy of ordering up to y^* works as follows. If x is initially greater than y^* we do nothing until x drops below y^*. Once x drops below y^* and we place the initial order $y^* - x$, all subsequent orders will be equal to the previous period demand. To see this, suppose that we order up to y^* at the beginning of period t. Then $x_{t+1} = y^* - D_t$, so $y^* - x_{t+1} = D_t$ is the amount to be ordered at the beginning of period $t+1$. This policy is also known as a base-stock policy because orders are placed in each period to restore the inventory to y^*.

Notice that as α increases to one, i.e., no discounting, the optimal policy is to order up to y^* where y^* satisfies

$$P(D \leq y) = \frac{p}{h+p}. \tag{11}$$

Also, as α increases to one, the discounted cost goes to infinity and it makes more sense to talk about the average cost per period. It can be shown, e.g., by using the vanishing discount cost method, that the policy that sets y^* as prescribed in equation (11) is indeed an optimal solution for the average cost case.

Notice also that the myopic policy is also optimal for the *finite horizon* stationary problem provided we set $c_{T+1}(x) = -cx$.

2.3 Positive Lead Times

Suppose that an order placed at the beginning of period t arrives at the beginning of period $t+L$. To work with positive, but deterministic, lead times, we need to add the inventory on order to the inventory level to summarize the state space at the beginning of each period. The resulting quantity is known as the inventory position and is equal to the inventory on hand plus on order minus backorders. When the lead time is zero, the inventory position is equal to the inventory level. Let x_t be the inventory position at the beginning of period t, after we receive the order placed L periods ago, but before we make the ordering decision for period t. Suppose that we order to bring the inventory position to $y_t \geq x_t$. This order will arrive at the beginning of period $t+L$. All orders placed prior to period t would have arrived by the beginning of period $t+L$. Moreover, orders placed after period t will not arrive until after period $t+L$. Consequently, the inventory level at the end

of period $t+L$ is given by $y_t - D[t, t+L]$ where $D[t, t+L] = \sum_{s=t}^{t+L} D_s$. The demand $D[t, t+L]$ over periods $\{t, \ldots, t+L\}$ is known as the lead time demand starting from period t. Notice that $D[t, t+L]$ contains the demand over $L+1$ periods and reduces to D_t when $L = 0$. Since the decision made at time t determines the holding and penalty costs incurred at the end of period $t+L$ it makes sense to charge these costs to period t. This is accomplished by redefining the loss function to be

$$G_t(y) = h_t E(y - D[t, t+L])_+ + p_t E(D[t, t+L] - y)_+.$$

Let $C_t(x_t)$ be the minimal expected discounted cost of managing the system starting from period t with inventory position x_t. Then,

$$C_t(x_t) = \min_{y_t \geq x_t} \{c_t(y_t - x_t) + G_t(y_t) + \alpha E C_{t+1}(y_t - D_t)\}.$$

This formulation is equivalent to (10) except that x_t is now the inventory position and G_t is defined differently. One additional difference is that the last ordering period is $T - L$ instead of T. Other than this, the problems are mathematically equivalent. The myopic policy calls for bringing the inventory position up to y_t^m in period where y_t^m satisfies

$$P(D[t, t+L] > y) = \frac{h_t + m_t}{h_t + p_t}.$$

The infinite horizon policy calls for bringing the inventory position up to y^* where y^* satisfies

$$P(D[t, t+L] > y) = \frac{h + c(1-\alpha)}{h + p}.$$

Let

- μ mean demand per period
- σ standard deviation of daily demand
- μ_d mean of the leadtime demand.
- σ_d standard deviation of the leadtime demand.

If we assume that the period demands are statistically independent, then $\mu_d = \mu(1+L)$ and $\sigma_d = \sigma\sqrt{1+L}$. Often $D[t, t+L]$ can be modeled as normally distributed with mean μ_d and standard deviation σ_d. In this case,

$$y^* = \mu_d + z\sigma_d$$

where

$$\bar{\Phi}(z) = \frac{h_t + m_t}{h_t + p_t}.$$

2.3.1 Random Lead Times

When lead times are random things become complicated because of the possibility of order crossing, i.e., a recent order arrives before an old order. There is no easy way to account for order crossings. In many practical manufacturing and distribution situations orders do not cross or they cross so rarely that it makes sense to build a model under the assumption that orders do not cross although this assumption may be inconsistent with the assumption that demands are time-independent. If we are willing to assume that orders do not cross, then the problem can be solved, at least approximately, once we find the mean and the variance over the lead time.

Let L be the lead time. To simplify the notation we will let μ_l and σ_l to denote respectively the mean and the standard deviation of $L+1$. Our objective is to write μ_d and σ_d in terms of μ, σ, μ_l and σ_l under the assumption that the period demands are statistically independent. The formula for

the mean lead time demand is again $\mu_d = \mu\mu_l$. The formula for σ_d, which is what we are interested in, is given by

$$\sigma_d = \sqrt{\mu_l\sigma^2 + \sigma_l^2\mu^2}.$$

These results are direct applications of the well known formulas:

$$E[\sum_{i=1}^{N} X_k] = E[E[\sum_{i=1}^{N} X_k|N]]$$

and

$$\text{Var}[\sum_{i=1}^{N} X_k] = \text{Var}[E[\sum_{i=1}^{N} X_k|N]] + E[\text{Var}[\sum_{i=1}^{N} X_k|N]].$$

and can be found on on page 153 in reference [8].

Numerical Example The mean daily demand for a product is $\mu = 80$ units and the standard deviation is $\sigma = 20$ units.

- Scenario 1. The leadtime is short, but unreliable: The mean leadtime is $\mu_l = 5$ days but the standard deviation is $\sigma_l = 4$ days. In this case, the standard deviation of the leadtime demand is
$$\sigma_d = \sqrt{5(20)^2 + (4)^2(80)^2} = 323.$$
- Scenario 2. The leadtime is long, but reliable: The mean leadtime is $\mu_l = 25$ days but the standard deviation is $\sigma_l = 0$ days. In this case, the standard deviation of the leadtime demand is
$$\sigma_d = \sqrt{25(20)^2 + (0)^2(80)^2} = 100.$$

Since the holding and penalty costs are proportional to the standard deviation of demand, we see that the costs are over three times higher with the shorter and more unreliable leadtime. Comparing the standard deviation of the lead time demand to the mean lead time demand shows that the insidious effect of randomness in the lead time is even worse than the direct comparison between the standard deviations would indicate.

3 Positive Ordering Costs

3.1 (Q, r) Policies

Up to now we have considered periodic review models where decisions are made at the beginning of each period. In this section we consider a continuous review model with nonzero lead times and positive setup costs. We will first restrict our attention to the class of (Q, r) policies and then discuss (s, S) policies. Under a (Q, r) policy, we monitor the inventory position and place an order of size Q whenever the inventory position falls to or below the reorder point r. Under an (s, S) policy, we monitor the inventory position and place an order to restore the inventory position to S whenever the inventory position falls to or below s. (Q, r) policies are equivalent to (s, S) policies with $s = r$ and $S = r + Q$ when all demands are for a single unit. The equivalence also holds when $Q = 1$ and integer demands, even if not all of the demands are of size one. The case $Q = 1$ reduces to a base stock policy with base stock level $S = r + Q = r + 1$. The policies behave differently when demands can be larger than one and $Q > 1$ because a demand of more than one unit may bring the inventory position strictly below $s = r$ and in this case ordering in batches of $Q > 1$ units may not restore the inventory position to $S = r + Q$.

Let $D(t)$ denote the cumulative demand up to time t. Let L denote the known and constant leadtime. Let $D(t|L) \equiv D(t) - D(t - L)$ be the number of units demanded over the time interval $(t - L, t]$. We will assume that as $t \to \infty$, $D(t|L)$ converges in distribution to a random variable that

we will denote by D. In the Poisson case $D(t)$ is Poisson with parameter λt and $D(t|L)$ is Poisson with parameter λL which is independent of t. As a result D is Poisson with parameter λL.

To keep track of the evolution of the system, let

- $I(t)$ inventory on hand at time t.
- $B(t)$ backorders at time t.
- $IN(t)$ net inventory at time t.
- $IO(t)$ inventory on order at time t.
- $IP(t)$ inventory position at time t.

The net inventory $IN(t) = I(t) - B(t)$, and it is equal to the inventory $I(t)$ when positive and equal to $-B(t)$ when negative. In other words, $I(t) = IN(t)_+$ and $B(t) = IN(t)_-$. The inventory on order $IO(t)$ at time t is equal to the number of orders placed during the interval $(t - L, t]$. The inventory position $IP(t)$ is defined as the inventory on hand plus the inventory on order minus the number of backorders. Consequently,

$$\begin{aligned} IP(t) &= I(t) + IO(t) - B(t) \\ &= IN(t) + IO(t). \end{aligned}$$

Notice that $IO(t) = 0$ when $L = 0$, and in that case the inventory position is equal to the net inventory.

Given a stationary demand processes we will show how to compute a number of performance measures under a (Q, r) policy. These measures include probability of stockouts, the average number of units on inventory, the average number of units backordered, and the average frequency of orders.

3.2 $Q = 1$

We will start by computing performance measures for the case $Q = 1$. This mode of operation is optimal when there are no setup costs or they are small relative to the cost of holding inventory, e.g., for expensive items with low demand rates. For convenience let $S = r + Q = r + 1$. Notice that in this case the (Q, r) policy is actually a base stock policy with base stock level S. Under this policy we order to keep the inventory position equal to S. As a consequence, $IP(t) = S$ except at demand epochs when the inventory position momentarily drops below S and an order is immediately placed to restore the inventory position to S. Since

$$S = IP(t) = IN(t) + IO(t)$$

we have

$$IN(t) = S - IO(t).$$

Under a base-stock policy orders are placed to keep the inventory position constant so $IO(t) = D(t|L)$. As $t \to \infty$, the random variable $D(t|L)$ converges in distribution to the stationary lead time demand D, so $IN(t)$ converges in distribution to

$$IN = S - D$$

Thus the stationary distribution of IN is determined by the stationary distribution of the lead time demand. Similarly, $I(t)$ and $B(t)$ converge to stationary random variables I and B where $I = (S - D)_+$ and $B = (D - S)_+$.

If we want to minimize the long-run expected holding and backorder costs, we need to select S to minimize $G(S)$ where

$$G(y) = hE(y - D)_+ + pE(D - y)_+.$$

This, of course, is a Newsvendor problem, so an optimal solution is given by the smallest integer, say S^*, such that

$$Pr(D \leq S) \geq \frac{p}{h+p}.$$

Let $A = Pr(B > 0)$ be the long run probability of stockouts (i.e., of having backorders). Since $Pr(B > 0) = Pr(D > S) = 1 - Pr(D \leq S)$, at $S = S^*$, $A \leq \frac{h}{h+p}$.

3.3 Q a Positive Integer

Now, suppose that Q is a positive integer. Then, under very general conditions on the demand process, it can be shown that the stationary inventory position is uniform between $r+1$, and $r+Q$. That is,

$$P(IP = j) = \frac{1}{Q} \qquad j = r+1, \ldots, r+Q.$$

Moreover, it can be shown that IP is independent of D.

When the inventory position is at $y \in \{r+1, \ldots, r+Q\}$, the holding and penalty cost rate is $G(y)$. Since the inventory position is uniformly distributed over $\{r+1, \ldots, r+Q\}$, it follows that the average holding and penalty cost is given by $\frac{1}{Q}\sum_{y=r+1}^{r+Q} G(y)$. If the average demand per unit time is λ, and a setup cost K, is incurred every time an order of size Q is placed, then the average ordering cost is given by $\frac{K\lambda}{Q}$.

The above performance measures can then be combined to form the cost function:

$$c(Q, r) = \frac{K\lambda}{Q} + \frac{1}{Q}\sum_{y=r+1}^{r+Q} G(y).$$

On the other hand, the probability of stockouts is given by $Pr(D > y)$ when the inventory position is y. Since the inventory position is uniformly distributed, it follows that

$$A = Pr(B > 0) = \frac{1}{Q}\sum_{y=r+1}^{r+Q} Pr(D > y).$$

3.4 Algorithm

We will now discuss an algorithm to find the optimal (Q, r) pair and its associated cost. The algorithm is based on three observations.

First, since $-G(y)$ is unimodal, the problem

$$c(Q) = \min_r c(Q, r)$$

is easily solved by finding the set of Q *consecutive* integers minimizing $G(\cdot)$. More precisely, we want to find the consecutive integers

$$\{y_1, \ldots y_Q\}$$

such that

$$y_1 = \text{argmin}\{G(y) : y \in \mathcal{Z}\},$$

and, given $y_1, \ldots, y_k$

$$y_{k+1} = \text{argmin}\{G(y) : y \in \mathcal{Z}, y \neq y_i, i = 1, \ldots, k\}.$$

Letting G_k denote $G(y_k)$ we can write

$$c(Q) = \frac{K\lambda + \sum_{k=1}^{Q} G_k}{Q}.$$

The second observation is that we can write $c(Q)$ as a convex combination of $c(Q-1)$ and G_Q. Indeed it is easy to verify that

$$c(Q) = \frac{Q-1}{Q}c(Q-1) + \frac{1}{Q}G_Q.$$

This implies that $c(Q) < c(Q-1)$ if and only if $C(Q-1) > G_Q$ which implies that

$$G_Q < c(Q) < c(Q-1).$$

The third observation is that $-c(Q)$ is unimodal, which implies that the optimal batch size is the largest value of Q for which

$$G_Q < c(Q-1).$$

Algorithm

1. Set $Q = 1$ and find y_1, G_1 and $c(1)$.
2. Let
$$L_Q = \min\{y_1, \ldots, y_Q\} - 1,$$
$$R_Q = \max\{y_1, \ldots, y_Q\} + 1,$$
and $G_{Q+1} = \min(G(L_Q), G(R_Q))$. If $G_{Q+1} \geq c(Q)$ then stop. Otherwise compute
$$c(Q+1) = \frac{Q}{Q+1}c(Q) + \frac{1}{Q+1}G_{Q+1}$$
and set $y_{Q+1} = L_Q$ if $G(L_Q) < G(R_Q)$ and $y_{Q+1} = R_Q$ otherwise.
3. Set $Q \leftarrow Q+1$ and return to Step 2.

This algorithm is due to Federgruen and Zheng [3]
To facilitate the use of this algorithm it is convenient to write the increment of the $G(y)$ as

$$G(y+1) - G(y) = (h+p)P(D \leq y) - p.$$

For Poisson demands we can update $P(D = y)$ and $P(D \leq y)$

$$P(D = y+1) = \frac{\lambda L}{y+1}P(D = y),$$

starting from $P(D = 0) = e^{-\lambda L}$, and

$$P(D \leq y+1) = P(D \leq y) + P(D = y+1).$$

3.5 Sensitivity, Bounds and Heuristics

Let us consider again the cost function

$$c(Q, r) = \frac{K\lambda + \sum_{y=r+1}^{r+Q} G(y)}{Q}$$

that arises when the demand rate is λ, the ordering cost is K, the holding cost is h the backorder cost is p and the lead time demand is a random variable D with mean μ and variance σ^2.

Notice that if the variance $\sigma^2 = 0$ the demand is deterministic and the resulting problem is essentially an economic order quantity where we need to balance the ordering holing and backorder costs. On the other hand, if the ordering cost $K = 0$ then the problem reduces to the newsvendor

problem where we need to decide on the stock level to minimize the holding and backorder costs. Thus, the cost function $c(Q,r)$ reduces to well known subproblems if either $\sigma^2 = 0$, or $K = 0$.

Although we have developed a fairly deep understanding of both the EOQ and the newsvendor subproblems and have an efficient algorithm to minimize the cost function $c(Q,r)$, we don't yet have a deep understanding of the cost function $c(Q,r)$. Is it more or less sensitive than the EOQ to incorrect specifications of the batch size or the cost parameters? Is it more or less sensitive than the newsvendor problem to the specification of the distribution of the lead time demand? Can we obtain effective bounds on the average cost without having to run the algorithm? How does the average cost behave as a function of the set up cost and the variance of the lead time demand? Can we find upper and lower bounds on Q? Are there effective heuristics for the batch size? We now provide answers to some of these questions. The results, except as noted, are due to Gallego [4].

3.5.1 Sensitivity

It can be shown that $c(Q) = \min_r c(Q,r)$ is less sensitive than the EOQ in the sense that

$$\frac{c(Q)}{c(Q^*)} \leq \frac{1}{2}\left(\frac{Q}{Q^*} + \frac{Q^*}{Q}\right).$$

Notice that we have an inequality for the case of random demands, where we had an equality for the EOQ cost function. This result is due to Zheng [13].

3.5.2 Bounds

We have the following closed form bounds on the cost function

$$\sqrt{c_d^2 + G_1^2} \leq c(Q^*) \leq \sqrt{c_d^2 + \overline{G}_1^2}.$$

where c_d is the average cost of the EOQ subproblem,

$$G_1 = G(y_1) = \min\{G(y) : y \in \mathcal{Z}\}$$

is the newsvendor cost, and

$$\overline{G}_1 = \sigma\sqrt{hp}$$

is Scarf's upper bound on the newsvendor cost. Recall that $c_d = \sqrt{2HK\lambda}$ where $H = \frac{hp}{h+p}$.

Closed form bounds on Q^* are given by

$$Q_d \leq Q^* \leq Q_e$$

where

$$Q_d = c_d/H$$

is the economic order quantity, and

$$Q_e = \sqrt{c_d^2 + \overline{G}_1^2}/H = \sqrt{Q_d^2 + Q_\sigma^2}$$

where

$$Q_\sigma = \frac{\overline{G}_1}{H}.$$

3.5.3 Heuristics

It can be shown that

$$\frac{c(\sqrt{2}\ Q_d)}{c(Q^*)} \leq 1.061,$$

so using a batch size that is $\sqrt{2}$ times the EOQ results in a cost increase of at most 6.1%. In practice, we get closed to this upper bound when G_1 is small relative to c_d. In practice, the $\sqrt{2}\ Q_d$ heuristic can be improved by using the batch size

$$Q_g = \min(\sqrt{2}\ Q_d, \sqrt{Q_d Q_e}).$$

3.6 General Demand Sizes

When demands are not for one unit at a time an order under an (Q, r) policy consists of the number of batches of size Q that are necessary to bring the inventory position to the interval $\{r+1, \ldots, r+Q\}$. In this case, (Q, r) policies are no longer optimal. Managerially (Q, r) policies are policies are still attractive because the more restricted order size facilitates packaging, transportation, and coordination. Let X denote the random demand size. Then, the long run average cost under an (Q, r) policy is given by

$$c(Q, r) = \frac{K\lambda E \min(Q, X) + \sum_{j=r+1}^{r+Q} G(y)}{Q} \tag{12}$$

To see how the ordering cost arises, notice that when the inventory position is $r+j$, a demand of size X triggers an order if and only if $X \geq j$. Since the inventory position is uniform $\{r+1, \ldots, r+Q\}$ the probability, and the long run average frequency, of placing an order is $\sum_{j=1}^{Q} P(X \geq j)/Q$. Since $X \geq 0$ and $E \min(Q, X) = \sum_{j=1}^{Q} P(X \geq j)$, the cost function (12) results.

4 (s, S) Policies

Under an (s, S) policy, $s < S$, the inventory manager places an order to increase the item's inventory position to the *order-up-to level* S, whenever he finds the item's inventory position to be at or below the *reorder-level* s.

Researchers have devoted a large effort to the problem of identifying single-item stochastic inventory models for which an (s, S) policy is optimal. It turns out that (s, S) policies are optimal for a large class of single-item inventory models including the one we will study in this section. Here we will take the optimality of (s, S) policies for granted and will concern ourselves with the problem of computing an optimal (s, S) policy for a model where both the demand and the relevant costs are time stationary.

We assume that orders may be placed at the beginning of each period, orders are delivered immediately, all stockouts are backordered, period demands are independent and identically distributed, and that costs are stationary over time. Later we discuss how to extend the model to positive lead times.

The objective is to minimize the long run average cost over an infinite horizon.

Notation:

D the one period demand,
$p_j = Pr(D = j),\ j = 0, 1, \ldots,$
$K > 0$ fixed cost of placing an order,
$G(y)$ one period expected cost starting with y units.

The typical form of $G(y)$ is

$$G(y) = hE(y-D)_+ + pE(y-D)^-,$$

where h is the holding cost rate and p is the stockout penalty cost rate. However, other forms of $G(\cdot)$ may also arise. In any event, all that we will require of $G(\cdot)$ is that:

(i) $-G(\cdot)$ is unimodal,
(ii) $\lim_{|y|\to\infty} G(y) > \min_x G(x) + K$.

Let $c(s,S)$ denote the long run average cost of using the policy (s,S). To obtain an expression for (s,S) we use the well known reward-renewal theorem that states that the long run average cost is equal to the expected cost per cycle divided by the expected cycle length. A cycle is interpreted as the time elapsed between the placement of two consecutive orders. We say that the system renews itself after each cycle because the item's inventory position immediately after an order is placed is equal to S.

We are now concerned with the determination of the expected cost per cycle, and the expected cycle length. For $y > s$, let $k(s,y)$ denote the total expected cost until the next order is placed when the starting inventory position is equal to y units. Our interest, of course, is in finding a formula for $k(s,S)$. Likewise, let $M(j)$ be the expected total time until an order is placed when starting with $s+j$ units. Our interest, of course, is to find a formula for $M(S-s)$. Once these formulas are obtained, we can write

$$c(s,S) = \frac{k(s,S)}{M(S-s)}.$$

It is clear that the functions $k(s,\cdot)$, and $M(\cdot)$ satisfy the discrete renewal equations

$$k(s,y) = G(y) + K\sum_{j=y-s}^{\infty} p_j + \sum_{j=0}^{y-s-1} p_j k(s,y-j), \quad y > s$$

and

$$M(j) = 1 + \sum_{i=0}^{j-1} p_t M(j-i), \quad j = 1,2,\ldots$$

Let $m(0) = 1/(1-p_0),\ M(0) = 0$, and

$$m(j) = \sum_{k=0}^{j} p_k m(j-k), \quad j = 1,2,\ldots.$$

It follows that

$$M(j) = M(j-1) + m(j-1), \quad j = 1,2,\ldots,$$

and

$$k(s,y) = K + \sum_{j=0}^{y-s-1} m(j)G(y-j) \quad y > s.$$

Consequently,

$$c(s,S) = \frac{K + \sum_{j=0}^{S-s-1} m(j)G(S-j)}{M(S-s)}.$$

Unfortunately the cost function $c(s,S)$ is not, in general, convex. For a long time this fact precluded the development of efficient algorithms. However, Zheng and Federgruen [14] have observed that

$$c(s-1,S) = \alpha_n c(s,S) + (1-\alpha_n)G(s) \tag{13}$$

where

$$\alpha_n \equiv \frac{M(n)}{M(n+1)},$$

and $n = S - s$. Based on this observation, they have derived a very effective algorithm to compute an optimal (s, S) policy. We present here some of their key results, as well as their algorithm. ¿From (1) we see that $c(s-1, S)$ is a convex combination of $c(s, S)$ and of $G(s)$, and consequently

$$\min\{c(s,S), G(s)\} \leq c(s-1,S) \leq \max\{c(s,S), G(s)\}.$$

We will use property (1) to determine necessary and sufficient conditions on s^o to be the optimal reorder-level for a fixed order-up-to level S. Then, we will obtain lower and upper bounds on an optimal reorder-level and an optimal order-up-to level.

For fixed S the reorder-level s^o is optimal if

$$c(s^o, S) \ \leq \ c(s, S) \ \ \forall s.$$

Consequently s^o must satisfy

$$c(s^o - 1, S) \ \geq \ c(s^o, S) \ \leq \ c(s^o + 1, S),$$

but then from (1)

$$G(s^o + 1) \ \leq \ c(s^o, S) \ \leq \ G(s^o). \tag{14}$$

Let $y_1^* = \min\{y : G(y) = \min_x G(x)\}$, and notice that $-\infty < y_1^* < \infty$.

We will now establishing lower and upper bounds on an optimal reorder-level s^*.

Proposition 1 *Let s_l^* denote the smallest optimal reorder-level, then*

$$s_l^* \leq \bar{s} \equiv y_1^* - 1.$$

Proof: Let s_l^* be the smallest optimal value of s that minimizes $c(s, S^*)$. Suppose for a contradiction that $s_l^* \geq y^*$, then it follows from the form of $c(s, S)$ that $c(s_l^*, S^*) \geq G(s_l^*)$ which in turn implies that $c(s_l^* - 1, S^*) \leq c(s_l^*, S^*)$ contradicting the definition of s_l^*. □.

Proposition 2 *Let s_u^* denote the largest optimal reorder-level $< y_1^*$. Then*

$$s^o \leq s_u^*$$

where s^o is the optimal order level for some arbitrary order-up-to level S.

Because s_u^* is optimal for S^* it follows that (2) must hold. In fact, we claim that $G(s_u^*+1) < c(s_u^*, S^*)$ holds. Suppose for a contradiction that $s_u^* < y^* - 1$, and that $G(s_u^* + 1) = c(s_u^*, S^*)$ holds. Then $s_u^* + 1 < y_1^*$ is also optimal, contradicting the definition of s_u^*. On the other hand, if $s_u^* = y_1^* - 1$, then, by the definition of y_1^*, $G(y_u^*) = G(s^* + 1) < c(s_u^*, S^*)$. Now, given any S, and an optimal reorder-level s^o for S, we have

$$G(s_u^* + 1) \ < \ c(s_u^*, S^*) \ \leq \ c(s^o, S) \ \leq \ G(s^o).$$

But then because $G(s)$ is unimodal, $G(s^o) \geq G(s_u^*) \geq G(s_u^* + 1)$, so $s^o \leq s_u^*$. □.

Corollary 3 *There exists an optimal solution s^* satisfying*

$$s^o \leq s^* \leq \bar{s}. \tag{15}$$

where s^o is an optimal reorder-level for an arbitrary order-up-to level S.

We now turn our attention to bounds on S^*. To this end, let $\underline{S} \equiv \max\{y : G(y) = \min_x G(x)\}$; notice that $y_1^* \leq \underline{S} < \infty$. Let $c^* = c(s^*, S^*)$ denote the optimal average cost value, and let $\bar{S} \equiv \max\{y \geq \underline{S} : G(y) \leq c^*\}$.

Proposition 4 *There exists an optimal policy (s^*, S^*) for which*

$$\underline{S} \ \leq \ S^* \ \leq \ \bar{S}. \tag{16}$$

Proof: We start by proving the lower bound. Let (s^*, S^*) be an optimal (s, S) policy that maximizes the value of S^*. Assume for a contradiction that $S^* < \underline{S}$. Note that for $j \geq 0$, $G(S^* + 1 - j) \leq G(S^* - j)$, so $c(s^* + 1, S^* + 1) \leq c(s^*, S^*)$ contradicting the maximality of S^*.

To show the upper bound, assume for a contradiction that $G(S^*) > c^*$. Notice that from the definition of $k(s, \cdot)$ and $M(\cdot)$ we can write

$$\begin{aligned} c^* &= \frac{G(S^*) + KPr(D \geq S^* - s^*) + \sum_{j=0}^{S^*-s^*-1} p_j k(s^*, S^* - j)}{1 + \sum_{j=0}^{S^*-s^*-1} p_j M(S^* - s^* - j)} \\ &> \frac{c^* + \underline{k} Pr(D < S^* - s^*)}{1 + \underline{M} Pr(D < S^* - s^* - 1}, \end{aligned}$$

where

$$\begin{aligned} \underline{k} &= \frac{\sum_{j=0}^{S^*-s^*-1} p_j k(s^*, S^* - j)}{Pr(D < S^* - s^*)} \\ \underline{M} &= \frac{\sum_{j=0}^{S^*-s^*-1} p_j M(S^* - s^* - j)}{Pr(D < S^* - s^*)}. \end{aligned}$$

Consequently,

$$c^* > \frac{\underline{k}}{\underline{M}}. \tag{17}$$

However, we can identify the right hand side of (5) as the average cost of a feasible policy! This contradicts the optimality of (s^*, S^*) so $G(S^*) \leq c^*$.

Corollary 5 *Let $c > c^*$, and $\bar{S}_c \equiv \max\{y \geq \underline{S} : G(y) \leq c\}$, then $S^* \leq \bar{S} \leq \bar{S}_c$.*

Corollary 5 can be used to identify increasingly tighter upper bounds for S^* as increasingly better (s, S) policies are found.

For any fixed order up to level S, let

$$c^*(S) = \min_{s<S} c(s, S).$$

S is said to be *improving* upon S^o, if $c^*(S) < c^*(S^o)$.

Lemma 6 *For a given order-up-to level $S^o (\geq y_1^*)$, let $s^o (< y_1^*)$ be an optimal reorder-level. Then $c^*(S) < c^*(S^o)$ if and only if $c(s^o, S) < c(s^o, S^o)$.*

Proof: Suppose $c(s^o, S) < c(s^o, S^o)$, then $c^*(S) \leq c(s^o, S) < c(s^o, S^o) = c^*(S^o)$.

Conversely, assume that $c^*(S) < c^*(S^o)$, and that $c(s^o, S) \geq c(s^o, S^o)$. To reach a contradiction it is enough to show that $c(s, S) \geq c(s^o, S^o)$ for all $s < y_1^*$. First, consider $s^o < s < y_1^*$, and notice that the optimality of s^o implies that $c(s^o, S^o) \geq G(s^o + 1)$, and since $-G(\cdot)$ is unimodal $G(S - j) \leq c(s^o, S^o)$ for $j = S - s, \ldots, S - s^o - 1$. Consequently,

$$c(s^o, S) \ = \ \frac{K + \sum_{j=0}^{S-s-1} m(j) G(S - j) + \sum_{j=S-s}^{S-s^o-1} m(j) G(S - j)}{M(S - s^o)}$$

$$\begin{aligned} &= \frac{c(s,S)M(s,S) + \sum_{j=S-s}^{S-s^o-1} m(j)G(S-j)}{M(S-s^o)} \\ &\leq \frac{c(s,S)M(s,S) + \sum_{j=S-s}^{S-s^o-1} m(j)c(s^o,S^o)}{M(S-s^o)} \\ &= \beta c(s,S) + (1-\beta)c(s^o,S^o), \end{aligned}$$

where $\beta = \frac{M(S-s)}{M(S-s^o)}$. Thus for $s^o < s < y_1^*,\ c(s^o,S)$ is dominated by a convex combination of $c(s,S)$ and $c(s^o,S^o)$. But then, $c(s^o,S) \geq c(s^o,S^o)$ implies $c(s,S) \geq c(s^o,S^o)$.

Now, for $s < s^o$, the fact that $G(S-j) \geq c(s^o,S^o)$ for $j = S-s^o,\ldots,S-s-1$, allow us to write

$$c(s,S) \geq \gamma c(s^o,S) + (1-\gamma)c(s^o,S^o),$$

where $\gamma = \frac{M(S-s^o)}{M(S-s)}$, and consequently $c(s,S) \geq c(s^o,S^o)$. □.

Thus, given (s^o,S^o), we can easily identify an improving S' by simply comparing $c(s^o,S^o)$ and $c(s^o,S')$. If S' improves on S^o, then we want to find an optimal reorder-level s' for S'. The following lemma restricts the search for s' to $s^o,\ldots,\bar{s}$.

Lemma 7 *Assume that $s^o \leq \bar{s}$ is an optimal reorder-level for S^o and that S' improves on S^o, then there exists an optimal reorder-level s' for S' with $s' \in \{s^o,\ldots,\bar{s}\}$.*

Proof: Given S' we know from Proposition 1 that there exists an optimal reorder-level $\leq \bar{s}$. Let s' be the largest optimal reorder-level ($\leq \bar{s}$) for S'. Then $G(s'+1) < c(s',S') \leq c(s^o,S') < c(s^o,S^o) \leq G(s^o)$. Since $-G(\cdot)$ is unimodal it follows that $s^o \leq s'$. □.

We are now ready to present an algorithm to find an optimal (s^*,S^*) policy.

Algorithm.

Let y^* be a minimizer of $G(\cdot)$.

Step 0. (Initial Solution)
$S^o = y^*$;
$s = y^* - 1$;
DO WHILE $c(s,S^o) > G(s)$;
$s = s - 1$;
ENDO;
$c^o = c(s,S^o),\ S = S^o + 1$;
Step 1 (Main Step)
DO UNTIL $G(S) > c^o$;
IF $c(s,S) < c^o$;
$S^o = S$;
DO WHILE $c(s,S^o) \leq G(s+1)$;
$s = s + 1$;
ENDO;
$c^o = c(s,S^o)$;
ENDIF;
$S = S + 1$;
ENDO;
END;

5 Multi-echelon Systems

We present a new Dynamic Programming formulation for the infinite-horizon multiple-stage serial production/distribution system that commonly arises in Supply Chain Management. The formulation is based on local cost accounting and the intermediate cost functions have a precise interpretation as the cost of truncated sub-systems. This formulation enables an algorithm based on simple

gradient update formulas that reduces the computational work. In addition, the formulation results in a natural heuristic that provides near-optimal policies by solving a single newsvendor problem for each stage in the system. We show, through an extensive numerical study, that the heuristic is very effective in identifying near-optimal base-stock levels. We conclude by providing a distribution-free approximate bound that accurately predicts the sensitivity of the optimal average cost to the system parameters.

The study of multi-stage serial inventory systems is central to the study of supply chain management both as a benchmark and as a building block for more complex supply networks. Unfortunately, existing policy evaluation and optimization algorithms (see Gallego and Zipkin 1999) are difficult to understand and communicate. In this paper, we provide a new dynamic programming formulation based on the idea of allocating a given echelon-inventory level for a sub-system between the local inventory level for the most upstream stage of the sub-system and the successor's echelon base stock level. This formulation yields a new algorithm that can be made efficient by updating the gradients to compute optimal base stock levels and costs for each of the sub-systems. Second, based on this formulation we develop simple spreadsheet-based heuristics that computes one newsvendor problem per stage and is more accessible to practitioners and to Production and Operations Management students. The need to develop accurate, spreadsheet based heuristics that are easy to understand has been correctly identified by Shang and Song (2003), who develop a spreadsheet based heuristic based on solving two newsvendor problems per stage. We evaluate our heuristic and compare it to that of Shang and Song by testing it on the set of test problems in Gallego and Zipkin (1999) and Shang and Song (2003), and in additional experiments designed to test the performance when different stages have different lead times. Our numerical results indicate that our heuristic is near optimal, with an average error that is lower than the Shang and Song heuristic. Finally, we provide an approximate distribution-free bound that accurately reflects the sensitivity of the optimal average cost to changes in system parameters.

Consider a series system that consists of J stages as illustrated in the figure. Stage $j < J$ procures from sage $j+1$ and stage J replenishes from an outside supplier with ample stock. Customer demand occurs only at stage 1 and follows a (compound) Poisson process, $\{D(t), t \geq 0\}$. It takes $L_j, j = 1, \ldots, J$, units of time for a shipment to arrive to Stage j once it is shipped from its predecessor.

Unsatisfied demand is backordered at each stage but only Stage 1 incurs a linear backorder penalty cost b, per unit per unit time. We assume without loss of generality that each stage adds value as the item moves through the supply chain, so echelon holding costs h_j^e are positive. The local holding cost for stage j is $h_j = h^e[j, J] \equiv \sum_{i=j}^{J} h_j^e$, where sums over empty sets will be defined as zero. The system is operated under continuous review, so management orders every time a demand occurs. As pointed out by Zipkin (2000), this is justified for expensive and/or slow moving items.

It is known that an echelon base stock policy $\mathbf{s} = (s_J, \ldots, s_1)$ is optimal for this series system (Zipkin (2000), Federgruen and Zipkin (1984) and the original work by Clark and Scarf (1960)). Under this policy, the manager continuously monitors the echelon inventory-order position at each stage and places an order from stage $j+1$ to bring it up to s_j whenever it is below this level.

We now provide a recursive formulation based on local holding cost accounting to calculate optimal echelon base stock levels. We first construct the recursive optimization and verify that the solution of this new formulation essentially produces the same result as the traditional, echelon cost accounting algorithm explained in Gallego and Zipkin (2001).

Let D_j be the leadtime demand for Stage j in equilibrium, $j = 1, \ldots, J$. When $J = 1$, the serial system reduces to a single stage model. The cost for this problem is given by

$$c_1(s) = h_1 E(s - D_1)^+ + bE(D_1 - s)^+. \tag{18}$$

Notice that $\Delta c_1(s) \equiv c_1(s+1) - c_1(s) = (h_1 + b)Pr(D_1 \leq s) - b$ is non-decreasing in s so $c_1(s)$ is convex. The optimal base stock level is simply given by $s_1 \equiv \min\{s \in \mathcal{Z}_+ : \Delta c_1(s) > 0\}$.

When $J > 1$, the problem is more complex. Consider the sub-system consisting of Stages $\{1, \ldots, j\}$. Assume that stage j replenishes its inventory from an external supplier with ample

supply. We refer to this new serial system as *sub-system*-$\{1,2,\ldots,j\}$. We define $c_j(s)$ to be the expected cost of *optimally* managing this sub-system when the *echelon* base stock level for Stage j is s. In other words, this sub-system is equivalent to the original J stages series system with $D_i \equiv 0, h_i \equiv 0$ for all $i > j$.

Consider now the sub-system consisting of stages $\{1,\ldots,j+1\}$. Our goal is to compute $c_{j+1}(\cdot)$ from the knowledge of $c_j(\cdot)$. To link these two sub-systems, we decompose the echelon base stock level s over stages $\{1,\ldots,j+1\}$ into the local base stock level x for Stage $j+1$ and the echelon base stock level $s-x$ for Stage j. If the local base stock level at Stage $j+1$ is x, then the net inventory at Stage $j+1$ will be $(x-D_{j+1})^+$. The net inventory of this stage accrues at cost rate h_{j+1}. The sub-system $\{1,\ldots,j\}$ has echelon base stock level $s-x$ but Stage $j+1$ now has limited inventory. Since Stage $j+1$ faces a shortage when $D_{j+1}-x>0$, the effective echelon inventory for sub-system $\{1,\ldots,j\}$ is limited to $s-x-(D_{j+1}-x)^+ = s-\max(D_{j+1},x) = \min(s-x, s-D_{j+1})$. Thus, a finite base stock level at Stage $j+1$ imposes an externality to the sub-system $\{1,\ldots,j\}$ whose average cost is now $Ec_j(\min(s-x,s-D_{j+1}))$. Finally, to find $c_{j+1}(s)$ we need to take into account the holding cost, $h_{j+1}ED_j$, of the units in-transit from Stage $j+1$ to Stage j. When we allocate x units of local base stock level to stage $j+1$, the cost of managing a series system with $j+1$ stages is, therefore, given by

$$c_{j+1}(x;s) = h_{j+1}E(x-D_{j+1})^+ + h_{j+1}ED_j + Ec_j(\min(s-x,s-D_{j+1})). \tag{19}$$

Let $c_{j+1}(s)$ denote the cost of an *optimal* allocation of s units. To find its value, we minimize $c_{j+1}(x;s)$ over integer values of $x \in \{0,\ldots,s\}$. Consequently,

$$c_{j+1}(s) = \min_{x\in\{0,\ldots,s\}}\{c_{j+1}(x;s)\}, \text{ for } j=1,\ldots,J, \tag{20}$$

where we define $h_{J+1}\equiv 0$ and $D_{J+1}\equiv 0$. The solution to this problem prescribes how to allocate s units of echelon for the subsystem $\{1,\ldots,j+1\}$. In particular, it tells us how much local base stock to hold at stage $j+1$ and how much echelon base stock to allocate for stage j.

We use the recursion in Equation (20), we obtain the optimal echelon base stock levels via

$$s_j = \min\{s\in\mathcal{Z}_+ : \Delta c_j(s) > h_{j+1}\}, \text{ for } j=1,\ldots,J. \tag{21}$$

We now prove that these echelon base stock levels are indeed optimal.

Proposition 1 *1. $c_j(s)$ is convex in s for $j=1,\ldots,J$,*

2. $x_j(s) = (s-s_{j-1})^+$ minimizes $c_j(x;s)$ for $j=2,\ldots,J$,

3. The policy is optimal with echelon base stock levels $(s_1,\ldots,s_J)$ and $c_J(s_J)$ is the expected cost for the entire system.

Proof

We base the proof on an induction argument on the number of stages in the series system. To do so, we first show that the statements are true for a two-stage series system; that is, for $J=2$. Part 1 for j=1 is trivially true; that is, $c_1(s)$ in Equation (18) is convex. Let $\Delta c_j(x;s) \equiv c_j(x+1;s) - c_j(x;s)$. For $J=2$ we have

$$\Delta c_2(x;s) = [h_2 - \Delta c_1(s-x-1)]Pr(D_2 \le x).$$

Notice that $\Delta c_2(x;s) = 0$ for all $x<0$ on account of $D_2 \ge 0$. The convexity of c_1 implies that $\Delta c_1(s-x-1)$ is decreasing in x. As a consequence, $\Delta c_2(x;s)$ has at most one sign change from $-$ to $+$ over the range $x\in\{0,\ldots,s\}$. From Equation (21), s_1 is the smallest integer y such that $\Delta c_1(y) > h_2$.[2] This implies that $\Delta c_2(x;s)$ changes sign from $-$ to $+$ for the first time when

[2]Or equivalently, the smallest integer y such that $Pr(D_1\le y) > \frac{b+h_2}{b+h_1}$. In other words, s_1 is the largest minimizer of the newsvendor problem with holding cost h_1-h_2, backorder cost $b+h_2$ and demand D_1. In particular, s_1 is independent of the distribution of D_2.

$x = (s - s_1)^+$, hence this is a minimizer of $c_2(\cdot; s)$. Note that this result shows that allocating s_1 units of echelon base stock level to stage 1 is optimal when $s \geq s_1$. We have

$$\begin{aligned} c_2(s) &= c_2((s-s_1)^+; s) = E[h_2((s-s_1)^+ - D_2)^+ + h_2 D_1 + c_1(\min(s_1, s - D_2))] \\ &= E[h_2(s - s_1 - D_2)^+ + h_2 D_1 + c_1(\min(s_1, s - D_2)], \end{aligned}$$

where the last equation follows since $(x^+ - a)^+ = (x - a)^+$ when $a \geq 0$. Since $c_1(\min(y, s - x))$ is convex in s for all x and y and convex combinations of a convex function and convexity is preserved by sums and expectations, it follows that $c_2(s)$ is convex. So far, we proved parts 1 and 2 and the optimality of allocating s_1 units to stage one. Note that $h_3 \equiv 0$ and $D_3 \equiv 0$ for a two stage serial system. Hence the minimizer of $c_2(s)$ is given by Equation (21). With this final observation, we have shown that s_1, s_2 are the optimal echelon base stock levels, concluding the proof of part 3 for $J = 2$.

Assume now that all three statements are true for some $n < J$. In that case c_j is convex for all $j \leq n$ and an optimal echelon base stock policy is given by $(s_1, \ldots, s_n)$. Now consider adding one more stage to the this sub-system with local holding cost h_{n+1}. Stage n will replenish from stage $n+1$ that has limited supply. Then for stage $n+1$, we need to allocate local base stock level. To find the optimal allocation, we look at the difference $c_{n+1}(x+1; s) - c_{n+1}(x, s)$, which is non-zero only when $D_{n+1} \leq x$. In this case, the difference is given by $h_{n+1} - c_n(s-x-1) + c_n(s-x) = h_{n+1} - \Delta c_n(s-x-1)$. Consequently,

$$\Delta c_{n+1}(x; s) = [h_{n+1} - \Delta c_n(s - x - 1)] Pr(D_{n+1} \leq x).$$

Now, since c_n is convex it follows that $\Delta c_{n+1}(x; s)$ has at most one sign change and this must be from $-$ to $+$. Since the sign change occurs at $(s - s_n)^+$, it follows that $x_{n+1}(s) = (s - s_n)^+$ minimizes $c_{n+1}(x; s)$ so $c_{n+1}(s) = c_{n+1}((s - s_n)^+; s)$. This result implies that allocating s_n units of echelon base stock level to stage n is optimal. This proves part 2 for $n+1$ and part 3. Therefore, we have

$$c_{n+1}(s) = h_{n+1} E(s - s_n - D_{n+1})^+ + h_{n+1} E D_{n+1} + E c_n(\min(s_n, s - D_{n+1})),$$

which is convex in s, proving part 1 for $n+1$. For an $n+1$ stage series system, by definition $h_{n+2} \equiv 0$ and $D_{n+2} \equiv 0$, it follows that the minimizer of $c_{n+1}(s)$ is given by Equation (21) and $c_{n+1}(s_{n+1})$ is the optimal expected cost of managing a series system with $n+1$ stages. This concludes the induction argument for $n+1$ and hence the proof. □

The optimal echelon base stock levels can also be found through solving the *traditional* recursive optimization for $j = 1, 2, \ldots, J$. This formulation is based on echelon cost accounting.

$$C_j(y) = E\{h_j^e(y - D_j) + C_{j-1}(\min[y - D_j, s_{j-1}])\} \quad (22)$$

$$s_j = \max\{y : C_j(y) \leq C_j(x) \text{ for all } x \neq y\}, \quad (23)$$

where $C_0(y) = (b + h_1)[y]^-$, see Gallego and Zipkin (1999). The optimal system wide average cost is given by $C_J(s_J^*)$. We now verify that the new algorithm produces the same echelon base stock levels as the traditional algorithm.

Proposition 2 *1.* $C_j(s) = c_j(s) - h_{j+1} E(s - D_j)$, *and 2.* $C_J(s_J) = c_J(s_J)$.

Proof

The proof is based on an induction argument. For the case $j = 1$, we have $C_1(s) = E[h_1^e(s - D_1) + (b + h_1)(D_1 - s)^+] = E[h_1^e(s - D_1) - h_1(s_1 - D_1) + h_1(s - D_1) + (b + h_1)(D_1 - s)^+] = c_1(s) + E[h_1^e(s - D_1) - h_1(s_1 - D_1)] = c_1(s) - h_2 E(s - D_1)$.

Suppose the result holds for j, then $C_{j+1}(s) = E[h_{j+1}^e(s - D_{j+1}) + C_j(\min(s_j, s - D_{j+1})) = E[h_{j+1}^e(s - D_{j+1}) + c_j(\min(s_j, s - D_{j+1})) - h_{j+1} E(\min(s_j, s - D_{j+1}) - D_j)] + \pm h_{j+1} E(s - s_j - D_{j+1})^+ = c_{j+1}(s) - h_{j+2} E(s - D_{j+1})$. The last equality can be verified easily.

Part 2 follows directly from the fact that $h_{J+1} \equiv 0$. □

From part 1 it follows that $\Delta C_j(s) = \Delta c_j(s) - h_{j+1}$. Consequently, the largest minimizer of $C_j(s)$ will be the smallest integer s; that is s_j, such that $\Delta c_j(s) > h_{j+1}$. Since this is consistent with the definition of s_j given by Equation (21) it follows that the two algorithms result in the same policy. The second part shows that there is also an agreement in the cost over the entire system.

5.1 A New Algorithm with Gradient Updates

To obtain optimal echelon base stock level for stage $j+1$ using Equation (21), we need to be compute $\Delta c_{j+1}(x)$. This computation requires us to first calculate $\Delta c_j(x)$, which in turn requires us to calculate $\Delta c_{j-1}(x)$. This recursive computation for $\Delta c_{j+1}(x)$ can be improved significantly if we use what we already know about Δc_j. The next proposition establishes the link among these functions.

Proposition 3 *For $j = 1, \ldots, J$, we have*

$$\Delta c_{j+1}(s) = h_{j+1} Pr(D_{j+1} \leq (s-s_j)^+) + \sum_{k=0}^{\min(s, s_j-1)} \Delta c_j(k) Pr(D_{j+1} = s-k) - b Pr(D_{j+1} > s). \quad (24)$$

Proof
We show first that Equation (24) holds when $s < s_j$. Note that for this case from Equation (20), $c_{j+1}(s) = h_{j+1} E D_j + E c_j(s - D_{j+1})$. Therefore,

$$\begin{aligned} \Delta c_{j+1}(s) &= E\Delta c_j(s - D_{j+1}) = \sum_{k=0}^{\infty} \Delta c_j(s-k) Pr(D_{j+1} = k) \\ &= \sum_{k=0}^{s} \Delta c_j(s-k) Pr(D_{j+1} = k) - b Pr(D_{j+1} > s) \\ &= \sum_{k=0}^{s} \Delta c_j(k) Pr(D_{j+1} = s-k) - b Pr(D_{j+1} > s), \end{aligned}$$

where the last two equation is a consequence of $\Delta c_j(s) = -b$ for $s < 0$. The last equation is equivalent to (24) for $s < s_j$. Next we show the result for $s \geq s_j$. For this case, we subtract $c_{j+1}(s) = E[h_{j+1}(s - s_j - D_{j+1}) + h_{j+1} E D_j + c_j(\min(s_j, s - D_{j+1}))]$ from $c_{j+1}(s+1) = E[h_{j+1}(s + 1 - s_j - D_{j+1}) + h_{j+1} E D_j + c_j(\min(s_j, s + 1 - D_{j+1}))]$. After some algebra we arrive at

$$\Delta c_{j+1}(s) = h_{j+1} Pr(D_{j+1} \leq s - s_j) + \sum_{k=s-s_j+1}^{\infty} \Delta c_j(s_j - k) Pr(D_{j+1} = k).$$

By noticing that $\Delta c_j(s) = -b$ for $s < 0$, we can rewrite the gradient as

$$\Delta c_{j+1}(s) = h_{j+1} Pr(D_{j+1} \leq s - s_j) + \sum_{k=0}^{s_j-1} \Delta c_j(k) Pr(D_{j+1} = s-k) - b Pr(D_{j+1} > s).$$

This is equivalent to Equation (24) for $s \geq s_j$, concluding the proof. □

Next we describe an algorithm to obtain best echelon base stock levels and the resulting cost.

$s_1 \leftarrow \min\{y \in \mathcal{Z}_+ : \Delta c_1(y) > h_2\}$,
FOR $j = 2$ to J DO
  $\Delta c_{j+1}(s) \leftarrow h_{j+1} Pr(D_{j+1} \leq (s-s_j)^+) + \sum_{k=0}^{\min(s, s_j-1)} \Delta c_j(k) Pr(D_{j+1} = s-k) - b Pr(D_{j+1} > s)$.
  $s_j \leftarrow \min\{y \in \mathcal{Z}_+ : \Delta c_j(y) > h_{j+1}\}$.
END
PRINT $(s_1, \ldots, s_J)$ and $c_J(s_J) = c_J(0) + \sum_{y=0}^{s_J-1} \Delta c_J(y)$.

5.2 Newsvendor Bounds and Heuristics

The new Dynamic Programming formulation in Equations (19) and (20) is intuitive and enables us to design a fast algorithm based on gradient updates. Yet, both the new and the tractional formulation

are difficult to explain to non-mathematically oriented students and practitioners. We now provide a heuristic that can be implemented in a spreadsheet by solving one newsvendor problem per stage.

Consider the subsystem $\{1,\ldots,j+1\}$, for some j such that $1 \le j < J$. Assume that $h_{j+1} < h_j$ and that all the stages $\{1,\ldots,j\}$ have the same holding cost $h_j = h_{j-1} = \ldots = h_1 \equiv H$. Since it is equally expensive to hold stock at stages $1,\ldots,j$, it is clearly optimal to hold stock only at stages $j+1$ and 1. In other words, allocating zero local base stock levels to stages $2,\ldots,j-1$ is optimal. With this allocation scheme Equation (19) simplifies to

$$c_{j+1}(x;s|H) = h_{j+1}(x-D_{j+1})^+ + h_{j+1}ED_j + H\sum_{k=1}^{j-1} ED_k + Ec_1(\min(s-x-D[2,j], s-D[2,j+1])). \quad (25)$$

Hence, the first difference is given by

$$\Delta c_{j+1}(x;s|H) = [b + h_{j+1} - (H+b)Pr(D[1,j] \le s-x-1)]Pr(D_{j+1} \le x). \quad (26)$$

Note that $\Delta c_{j+1}(x;s|H)$ crosses from $-$ to $+$ at $x = (s - s_j^{NV}(H))^+$, where $s_j^{NV}(H)$ is the solution of a newsvendor problem with holding cost $H - h_{j+1}$, backorder cost $b + h_{j+1}$ and demand $D[1,j]$; that is,

$$\begin{aligned} G_j(s|H) &\equiv E\{(H-h_{j+1})(y-D[1,j])^+ + (b+h_{j+1})(D[1,j]-y)^+\}, \\ s_j^{NV}(H) &\equiv \min\{s \in \mathcal{Z}_+ : \Delta G_j(s|H) > 0\}. \end{aligned} \quad (27)$$

Consider now the general case where $h_j < h_{j-1} < \ldots < h_1$. Assume first that we *increase* the holding costs of stages $j = 2,\ldots,j$ to h_1. For this new series system, the cost of allocating x units to $j+1$ and s units to j is given by $c_{j+1}(x;s|h_1)$ as defined in (25). Assume now that instead of increasing, we *decrease* the holding cost of stages $1,\ldots,j-1$ to h_j. The corresponding cost for this system is given by $c_{j+1}(x;s|h_j)$.

Proposition 4 *The following are true for all $s, x > 0$.*

1. $c_{j+1}(x;s|h_1) \ge c_{j+1}(x;s) \ge c_{j+1}(x;s|h_j)$
2. $\Delta c_{j+1}(x;s|h_1) \le \Delta c_{j+1}(x;s) \le \Delta c_{j+1}(x;s|h_j)$,
3. $s_j^{NV}(h_j) \le s_j \le s_j^{NV}(h_1)$.

Proof

Part 1 is trivially true because we force a larger holding cost to obtain $c_{j+1}(x;s|h_1)$, hence the upper bound on the original system cost $c_{j+1}(x;s)$. We also force a smaller holding cost to obtain $c_{j+1}(x;s|h_j)$, hence the lower bound. To prove Part 2 observe that from Equation (26) we have $\Delta c_{j+1}(x;s|h_1) \le \Delta c_{j+1}(x;s|h_j)$, proving Part 2 To prove Part 3, observe that the function $\Delta c_{j+1}(x;s|h_1)$ changes sign from $-$ to $+$ at $x = (s - s_j^{NV}(h_1))^+$ for the first time and that $\Delta c_{j+1}(x;s|h_j)$ changes sign from $-$ to $+$ at $x = (s - s_j^{NV}(h_j))^+$ for the first time. Together with part 2 these two observations imply Part 3. □

This proposition suggests that instead of solving the recursive algorithm, we can approximate optimal echelon base stock levels simply by s_j^{NV} for $j = 1,\ldots,J$, which are based on newsvendor solutions. We note that the bounds in Part 3 are the same newsvendor bounds as in Shang and Song (2003). They propose to solve the two newsvendor problems given in Equations with h_1 and h_j for each stage $\{1,\ldots,J\}$ to obtain $s_j^{NV}(h_j)$ and $s_j^{NV}(h_1)$. Next they either truncate or round the average of the solution to these two newsvendor problems.

We now propose an approach that consists of solving a *single* newsvendor problem based on approximate holding cost rate $h_j^{GO} \in (h_j, h_1)$. The idea is based on the approximate time an item spends at each stage of the subsystem. To obtain this approximation, we set

$$h_j^{GO} \equiv \sum_{k=1}^{j} L_k h_k / L[1,j].$$

We solve the newsvendor problem in Equation (27) with $H = h_j^{GO}$ to approximate the optimal echelon base stock levels for each stage $j = 1, \ldots, J$.

Proposition 5 *For any given j and s we have:*

1. $G_j(s|h_j) \leq G_j(s|h_j^{GO}) \leq G_j(s|h_1)$,
2. $s_j^{NV}(h_j) \leq s_j^{NV}(h_j^{GO}) \leq s_j^{NV}(h_1)$,
3. $G_j(s|h_j^{GO}) \leq \sqrt{(b+h_{j+1})(h_j^{GO}+h_{j+1})}\sqrt{\lambda L[1,j]E[X^2]}$, *where X is the random demand size of the compound Poisson process.*

Proof
Notice that we have $h_j \leq h_j^{GO} \leq h_1$. Part 1 follows immediately from this inequality. Since the newsvendor cost functions are convex we also have $\nabla G_j(s|h_j) \geq \nabla G_j(s|h_j^{GO}) \geq \nabla G_j(s|h_1)$ where $\nabla f(x) = f(x+1) - f(x)$. This implies Part 2. Finally Part 3 is the distribution-free bound in Gallego and Moon (1993) and Scarf (1953).

The last two propositions imply that if the bounds in Part 3 of Proposition 4 are tight then $s_j^{NV}(h_j^{GO})$ would be very close to the optimal base stock level, s_j. In the following section we illustrate how accurate this approximation is. If our approximation is close-to-optimal, the cost of managing the series system can also be bounded by a distribution-free bound, that is

$$c_J(s_J) \leq \sqrt{bh_J^{GO}}\sqrt{\lambda L[1,J]E[X^2]} + \sum_{i=1}^{J} h_{i+1}ED_i, \tag{28}$$

where the last term is to account for pipeline inventory. This simple form enables sensitivity analysis. In particular, (1) the system cost is proportional to $\sqrt{b}$, (2) downstream leadtimes have a larger impact on system performance than upstream leadtimes, (3) upstream echelon holding cost rates have a larger impact on the system performance than downstream echelon holding cost rates, (4) the system cost is proportional to $\sqrt{\lambda}$ and proportional to $\sqrt{E[X^2]}$.

This type of parametric analysis enables a near characterization of system performance. Some system design issues may require investments in new processing plans or quicker but more expensive shipment methods. Marketing strategies could influence the demand as well as altering the cost of backlogging a customer. The closed form expression (28) facilitates gauging the benefit of any action on the inventory management costs, at least as a first cut. Our analysis suggests, for example, that management should focus on reducing the lead time at the upstream stages while reducing the holding cost at the down stream stages. If process re-sequencing is an option, the lowest value added processes with the longest processing times should be carried out sooner than later.

5.3 Numerical Study

Here we report the performance of our heuristic and of the distribution-free bound. We compare the exact solution based on equations (22) and (23) and report the percentage error $\epsilon^i\% = \frac{c_J(s_J^i) - c_J(s_J)}{c_J(s_J)}$ for $i = \{SS, GO\}$. Shang and Song (2003) use $s_j^{SS} \equiv \frac{s_j^{NV}(h_j) + s_j^{NV}(h_1)}{2}$ and truncate this average when $b \leq 39$ and round it otherwise. We use $s_j^{GO} \equiv s_j^{NV}(h_j^{GO})$. By considering a larger set of experiments, we complement the numerical study in Shang and Song (2002). In particular, our numerical study includes unequal leadtimes.

To manage the series system, we use an echelon base stock policy with echelon base stock levels s_j^{GO} for all j. The approximate cost is given by $G_J(s^{GO}) + \sum_{i=1}^{J} h_{i+1}ED_i$. Shang and Song (2003) approximate the optimal cost by $G_J(s_J^{NV}(h_J) + \sum_{i=1}^{J} h_{i+1}ED_i$ instead of the average since the lower bounds become looser as the number of stages in the system increases. We study two sets of experiments: constant leadtime set and the randomized parameters set.

The first set of experiments is similar to that of Gallego and Zipkin (1999) and Shang and Song (2002). The holding cost and the lead times are normalized so $h_1 = 1$ and $L[1, J] = 1$. We consider $J \in \{2, 4, 8, 16, 32, 64\}$; $\lambda \in \{16, 64\}$; and $b \in \{9, 39\}$ (corresponding to fill rates of $90\%, 97.5\%$). Within this group we consider *linear* holding-cost form ($h_j^e = 1/J$); *affine* holding cost form ($h^e[1, j] = \alpha + (1 - \alpha)j/J$ with $\alpha = 0.25$ and 0.75); *kink* holding cost form ($h_j^e = (1 - \alpha)/J$ for $j \geq J/2+1$ and $h_j^e = (1+\alpha)/J$ for $j < J/2+1$ with $\alpha = 0.25$ and 0.75) and *jump* holding cost form ($h_j^e = \alpha + (1 - \alpha)/J$ for $j = N/2$ and $h_j^e = (1 - \alpha)/J$ for $j \neq J/2$ with $\alpha = 0.25$ and 0.75). Notice that Shang and Song (2002) consider only the case for $\lambda = 64$ and $b = 39$.

Out of 108 problem instances, in 24 cases the s^{GO} and in 20 cases the s^{SS} heuristic resulted in the same solution as the recursive optimization. The s^{GO} (resp., s^{SS}) heuristic outperforms in 48 (resp., 44) cases and they tie in 17 cases. The average error for s^{GO} (resp., s^{SS}) heuristic is 0.195% (resp., 0.385%), while the maximum error is 3.68% and 1.24% for the GO and the SS heuristics respectively. The quality of the heuristics seems to deteriorate as the number of stages in the system exceeds 32. The SS heuristic seems to perform better for the jump holding cost case, while the GO heuristic tends to dominate in the other cases.

The second set of experiments allow for unequal leadtimes. It is here that we expect the GO heuristic to perform better. To cover a wider range of problem instances we generate the leadtimes and holding costs from uniform distributions. In particular, we use the following set of parameters:

$$h_j^e \in \{\text{Unif}(0, 1), \text{Unif}(0, 5), \text{Unif}(1, 10)\},$$

$$L_j \in \{\text{Unif}(1, 2), \text{Unif}(1, 10), \text{Unif}(1, 40)\},$$

$$J \in \{2, 4, 8, 16, 32\} \quad b \in \{1, 9, 39, 49\} \quad \lambda \in \{1, 3, 6\}.$$

We consider 25 combinations, taken at random, from the above parameters. For each subgroup we generate 40 problem instances and calculate the worst case as well as the average performances.

Out of 1000 problem instances, in 188 cases the s^{GO} and in 133 cases the s^{SS} heuristic resulted in the exact solution. In 849 cases the error term for s^{GO} heuristic is smaller or equal to that of s^{SS} heuristic. The average error for the s^{GO} (resp., s^{SS}) heuristic is 0.23% (resp., 0.83%). We observe that as the variance of the leadtimes across stages increases the average error term for s^{GO} decreases (the average error for $L_j \sim \text{Unif}(1, 10)$ is 0.14% whereas it is 0.39% for $L_j \sim \text{Unif}(1, 2)$). Similarly the s^{GO} heuristic performs even better as the variance of echelon costs across stages in a series system increases.

In light of our numerical observations we suggest the s^{GO} heuristic for a series system with up to eight stages. Caution should be used for system with a large number of stages and for systems with jump holding costs.

We have also performed a numerical study comparing the actual cost to the distribution-free bound by performing simple linear regressions of the bound to the actual cost by fixing all but one of the parameters. The coefficients of determination R^2 for the different regressions are all close to one. This observation suggests that the bound can safely be used to investigate the impact of process and design changes on the cost of managing a series system. Notice that the bound only requires knowing h_J^{GO}, b, $L[1, J]$, λ and $E[X^2]$.

The simple newsvendor heuristic and the bound enable a manager to quantify with ease, for example, the impact of re-sequencing a process. Consider, for example, a four stage series system where $h_1 = L[1, J] = 1$, $b = 1$ and $\lambda = 16$. We now compare two systems with different configurations of leadtimes. The first system has leadtimes $(0.1, 0.1, 0.1, 0.7)$ and the second has leadtimes $(0.7, 0.1, 0.1, 0.1)$. The costs based on the distribution free bound (resp., recursive optimization) are 13.29 (resp., 12.77) for the first system and 5 (resp.,4.93) for the second system. The distribution free bound predicts a cost reduction of 62.4% while the actual cost reduction based on recursive optimization is 61.39%. This indicates that the distribution free bound enables a quick, yet accurate, what if analysis. In this case, we observe that postponing the shortest and the most expensive processes to a later stage can significantly reduce inventory related costs.

We mention in passing that we also explored using the holding cost $\sum_{k=1}^{j}(L_k^\alpha / \sum_{i=1}^{j} L_i^\alpha)h_k$ for different $\alpha \in [0, 1]$. We were unable to identify an α that results in lower error terms than $\alpha = 1$. In

addition, for some problem instances we have also calculated the *implied* holding costs h_j^{im}. These holding cost when used in the newsvendor problem of Equation (27) yield the optimal echelon base stock levels s_j obtained through the exact algorithm. In other words we set $h_j^{\min} \equiv \min\{h \in \mathcal{R}_+ : s_j^{NV}(h) = s_j\}$ and $h_j^{\max} \equiv \min\{h \in \mathcal{R}_+ : s_j^{NV}(h) = s_j - 1\}$. Note that using an implied holding cost $h_j^{im} \in [h_j^{\min}, h_j^{\max})$ in Equation (27) yields the optimal echelon base stock level. The range for possible implied holding cost is typically large and frequently contains h_j^{GO}.

We end by noticing that our heuristic can also be applied to assembly systems by applying Rosling (1989)'s ideas. For distribution systems, the heuristic can be applied after using the decomposition principles in Gallego, Özer and Zipkin (1999).

References

[1] Arrow, K., T. Harris, J. Marshack. 1951. Optimal Inventory Policy. *Econometrica*. **19**, 250-272.

[2] Edgeworth, F. 1888. The Mathematical Theory of Banking. *J. Royal Statistical Society*. **51**, 113-127.

[3] Federgruen, A. and Y.S. Zheng (1992) "An Efficient Algorithm for Computing an Optimal (Q, r) Policy in Continuous Review Stochastic Inventory Systems," *Operations Research*, 40, 808-813.

[4] Gallego, G. (1998) "New Bounds and Heuristics for (Q, r) Policies. *Management Science*, 44, 219-233.

[5] Gallego G. and I. Moon. 1993. The Distribution Free Newsboy Problem: Review and Extensions. *Journal of Operational Research Society.* 44, 825-834.

[6] Gallego G, Ö. Özer, and P. Zipkin. 1999. Bounds, Heuristics and Approximations for Distribution Systems. Working Paper.

[7] Gallego G. and P. Zipkin. 1999. Stock Positioning and Performance Estimation in Serial Production-Transportation Systems. *Manufacturing & Service Operations Management* 1, 77-87.

[8] Hadley, G., and T. Whitin˙ *Analysis of Inventory Systems*, Prentice Hall. 1963.

[9] Scarf, H. (1958) "A Min-Max Solution of an Inventory Problem," *Ch. 12 in Studies in The Mathematical Theory of Inventory and Production, Stanford Univ. Press.*

[10] Rosling, K. 1989. Optimal Inventory Policies for Assembly Systems Under Random Demands. *Operations Research* 37, 565-579.

[11] Shang H. K. and J. Song. 2003. Newsvendor Bounds and Heuristics for Optimal Policies in Serial Supply Chains. *Management Science* 49, 618-638.

[12] Veinott, A. F. (1965) Lecture Notes Stanford University.

[13] Zheng, Y.S. (1992) "On Properties of Stochastic Inventory Systems," *Management Science*, 38, 87-103.

[14] Zheng, Y.S. and A. Federgruen (1991) "Finding Optimal (s, S) Policies is About as Simple as Evaluating a Single Policy," *Operations Research*, OR 39, 654-665.