1 The Newsvendor Problem

In this chapter we discuss the problem of controlling the inventory of a single item with stochastic demands over a single period. This problem is also known as the Newsvendor Problem because the prototype is the problem faced by a newsvendor trying to decide how many newspapers to stock on a newsstand before observing demand. The newsvendor faced both overage and underage costs if he orders too much or if he orders too little. The Newsvendor Problems is therefore the problem of deciding the size of a single order that must be placed before observing demand when there are overage and underage costs. The problem is particularly important for items with significant demand uncertainty and large overage and underage costs.

Let D denote the one period random demand, with mean $\mu = E[D]$ and variance $\sigma^2 = V[D]$. Let c be the unit cost, $p > c$ the selling price and $s < c$ the salvage value. If Q units are ordered, then $\min(Q, D)$ units are sold and $(Q-D)^+ = \max(Q-D, 0)$ units are salvaged. The profit is given by $p\min(Q, D) + s(Q-D)^+ - cQ$. The expected profit is well defined and given by:

$$\pi(Q) = pE\min(Q, D) + sE(Q-D)^+ - cQ.$$

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

$$\pi(Q) = (p-c)\mu - G(Q) \qquad (1)$$

where

$$G(Q) = (c-s)E(Q-D)^+ + (p-c)E(D-Q)^+ \geq 0.$$

Let $h = c - s$ and $b = p - c$. It is convenient to think of h as the per unit overage cost and of b as the per unit underage cost. Sometimes the underage cost is inflated to take into account the *ill-will cost* associated with unsatisfied demand.

Equation (1) allow us to view the problem of maximizing $\pi(Q)$ as that of minimizing the expected overage and underage cost $G(Q)$.

Let $G^{det}(Q) = h(\mu - Q)^+ + b(Q-\mu)^+$. This represents the cost when D is deterministic, i.e., $Pr(D = \mu) = 1$. Clearly $Q = \mu$ minimizes $G^{det}(Q)$ and $G^{det}(\mu) = 0$, so $\pi^{det}(\mu) = (p-c)\mu$. Thus, the Newsvendor Problem is only interesting when demand is random. Notice that the problem also becomes trivial when $s = c$ for in this case we can order an infinite amount, satisfy all demand, and then return all unsold items.

Let $g(x) = hx^+ + bx^-$, 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. By Jensen's inequality $G(Q) \geq G^{det}(Q)$. As a result, $\pi(Q) \leq \pi^{det}(Q) \leq \pi^{det}(\mu) = (p-c)\mu$. Thus, the expected profit is lower than it would be in the case of deterministic demand.

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 $G'(Q) = hE\delta(Q-D) - bE\delta(D-Q)$ where $\delta(x) = 1$ if $x > 0$ and zero otherwise. Since $E\delta(Q-D) = Pr(Q-D > 0)$ and $E\delta(D-Q) = Pr(D-Q > 0)$, it follows that

$$G'(Q) = hPr(Q-D > 0) - bPr(D-Q > 0).$$

Setting the derivative to zero reveals that

$$F(Q) \equiv Pr(D \leq Q) = \frac{b}{b+h} = \frac{p-c}{p-s} \equiv \beta. \qquad (2)$$

If F is continuous then there is at least one Q satisfying Equation (2). We can select the smallest such solution by letting

$$Q^* = \inf\{Q \geq 0 : F(Q) \geq \beta\}. \tag{3}$$

It is clear that Q^*, selected this way, is increasing in β and therefore it is increasing in b and decreasing in h.

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

$$Q^* = F^{-1}(\beta). \tag{4}$$

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

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

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,

$$Q^* = \min\{Q \in \mathcal{N} : F(Q) \geq \beta\}, \tag{5}$$

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 with high probability. The fractile solution (2) 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. This service level should not be confused with the fraction of demand served from stock, or fill-rate, which is defined as $\alpha = E\min(D, Q)/ED$.

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 $z_\beta = \Phi^{-1}(\beta)$. 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

$$Q^* = \mu + z_\beta \sigma \tag{6}$$

satisfies Equation (4), so Equation (6) 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,

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

so

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

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. Notice, for example, that $\alpha = 97\%$ when $\beta = 75\%$ and $\text{cv} = 0.2$, while $\alpha = 99.1\%$ when $\beta = 90\%$ and $\text{cv} = 0.2$.
Example Normal Demand: Suppose that D is normal with mean $\mu = 100$ and standard deviation $\sigma = 20$. If $c = 5$, $h = 1$ and $b = 3$, 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 $\alpha = 97\%$.

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:

$$Pr(D = k) = Pr(D = k-1)\lambda/k, \quad k = 1, 2, \ldots$$

$$\begin{aligned} Pr(D \le k) &= Pr(D \le k-1) + Pr(D = k), \quad k = 1, 2, \ldots \\ E[(D-k)^+] &= E[(D-k+1)^+] - Pr(D \ge k) \quad k = 1, 2, \ldots \end{aligned}$$

An optimal value of Q is given by the smallest integer such that $P(D \le Q) \ge \beta$.
Example Poisson: If D is Poisson with parameter $\lambda = 25$, and $c = 5$, $h = 1$ and $b = 3$, then $\beta = 0.75$ and $Q^* = 28$ is optimal. To compute $G(Q^*)$ notice that $G(Q) = h(Q-\lambda)+(h+b)E(D-Q)^+$, 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 $\alpha = 1 - 0.87/25 = .97$.

[width=4.0in]norm$_{d}ist.bmp$

Figure 1: Series

4 The Lognormal Distribution

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 $\tau \ge 0$. The lognormal distribution is often used to model non-negative random variables such as lifetimes of electronic devices and the total returns of risky securities 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

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

and

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

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

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

Example Lognormal: Figure 1 shows actual weekly demand data for a semiconductor product with $c = 5$, $b = 5$ 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.

5 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 with a given mean and a given standard deviation. This order quantity is due to Herbert Scarf [9] and it is given by

$$Q^S = \mu + \frac{\sigma}{2}\left(\sqrt{\frac{b}{h}} - \sqrt{\frac{h}{b}}\right). \tag{7}$$

Notice that Scarf's formula (7) suggests ordering more (resp., less) than the mean demand when $b > h$ (resp., $b < h$). Moreover, $|Q^S - \mu|$ increases linearly in σ for $h \neq b$.

Scarf's formula and of other related results, see Gallego and Moon [5], follow from $x^+ = 0.5(|x| + x)$ and the Cauchy-Schwartz 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

$$G(Q^S) \leq \sqrt{bh}\sigma = \sqrt{(p-c)(c-s)}\sigma$$

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

$$(p-c)\mu - \sqrt{(p-c)(c-s)}\sigma \leq \pi(Q^*) \leq (p-c)\mu,$$

and

$$1 - \sqrt{\frac{c-s}{p-c}}\frac{\sigma}{\mu} \;\leq\; \frac{\pi(Q^*)}{(p-c)\mu} \;\leq\; 1. \tag{8}$$

This last expression allow us to see how far from optimal Scarf's solution is in the worst case. Notice that this depends an the distribution only through the coefficient of variation.

Scarf's ordering rule is modified to $Q^S = 0$ when the left hand side of (8) is negative. reflecting the fact that it may be better not to be in business when demand is very uncertain.

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

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

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

Finally, it is also possible to show that $G(\mu) \leq \frac{1}{2}(h+b)\sigma$, 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 $b = 3$. 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 $b = 3$. 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: $c = 5.00$, $h = 2$, $b = 5$, $\mu = 207$, $\sigma = 459$. In this case $\sigma/\mu > \sqrt{b/h}$ 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 $b = p - c$ and no additional penalties accrue for shortages.

6 Compound Demand

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

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

The case $Pr(N = 1) = 1$ and X normal reduces to $D = X$ normal and the case N Poison λ and $Pr(X = 1) = 1$ reduces to the Poisson case.

In general, it is difficult solve the Newsvendor Problem in closed form when $D = \sum_{k=1}^{N} X_k$. One alternative is to find Q by simulation. Another alternative is to compute $\mu_d = E[D]$ and $\sigma_d^2 = \text{Var}[D]$ and then approximate D by a known distribution with these moments, e.g., the Normal or Lognormal. The normal approximation is recommended only when $cv_d \leq 0.33$. The lognormal tends to work better for large values of cv_d. One could also optimize Q against the worst case distribution given the mean and the variance of D.

Using well known results on 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

$$\mu_d = \mu_n \mu_x \quad \text{and} \quad \sigma_d^2 = \mu_x^2 \sigma_n^2 + \mu_n \sigma_x^2.$$

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

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

Since cv_d is decreasing in μ_n, 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 roughly proportional to the standard deviation of demand, the cost of dealing with a small number of large customers can be significantly higher, about 80% 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

$$\mu_d = \lambda \mu_x \qquad \sigma_d^2 = \lambda(\mu_x^2 + \sigma_x^2).$$

Notice that the coefficient of variation for the compound Poisson distribution

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

7 Forecast Updates and Advance Demand Information

Our analysis, see (8), indicates that high risk items, i.e., those with large σ and large overage costs $c - s$ are those for which the Newsvendor Problem is more relevant. However, solving the Newsvendor Problem is only the first step in the analysis. The next step is to find how to change the *givens* to our advantage.

Typically, high risk items also have high margins and in many instances it is a good idea to sacrifice the margin to reduce the risk. This can be done in several ways. One is to use higher cost, but more responsive, suppliers that allow us to place orders later after we leave more information about demand. Another way is to give price incentives to elicit advance demand information from customers in the form of forward purchases.

7.1 Supplier Selection

Let

$$D = \mu + \epsilon_1 + \epsilon_2$$

where μ is a known constant and ϵ_1 and ϵ_2 are normal, mean zero, random variables with standard deviations σ_1 and σ_2 and correlation ρ where it is possible to observe ϵ_1 before observing ϵ_2 by postponing the purchase and agreeing to pay a higher unit cost. To be more precise, suppose that we can purchase at unit cost c before observing ϵ_1 or at unit cost $c(1+\delta) \geq c$ after observing ϵ_1. Obviously if $\delta = 0$, we would wait and order after observing ϵ_1. How large must δ be before we are better off ordering without observing ϵ_1?

If we order before observing ϵ_1, then the expected profit is

$$(p-c)\mu - (p-s)\sqrt{\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2}\,\phi(z_\beta).$$

Suppose we wait until observe ϵ_1. Then ϵ_2, conditioned on ϵ_1, is normal with mean $\rho\sigma_2\epsilon_1/\sigma_1$ and variance $\sigma_2^2(1-\rho^2)$. Therefore, conditioned on ϵ_1 the expected profit is given by

$$(p-c-c\delta)(\mu + \rho\sigma_2\epsilon_1/\sigma_1) - (p-s)\sigma_2\sqrt{1-\rho^2}\,\phi(z_{\beta'})$$

where $\beta' = (p - c(1+\delta))/(p-s)$. Taking expectation with respect to ϵ_1, the expected profit of waiting to order until ϵ_1 is observed is

$$(p-c-c\delta)\mu - (p-s)\sigma_2\sqrt{1-\rho^2}\,\phi(z_{\beta'}).$$

The maximum δ is therefore, the largest root of

$$c\delta\mu + (p-s)\sigma_2\sqrt{1-\rho^2}\,\phi(z_{\beta'}) = (p-s)\sqrt{\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2}\,\phi(z_\beta).$$

For example: If $c = 100$, $\mu = 100$, $\sigma_1 = 30$, $\sigma_2 = 20$, $p = 200$, $s = 0$, $\rho = .4$, then we would be willing to pay a $\delta = 19.48\%$ cost premium. The premium jumps to $\delta = 58.53\%$ if $\rho = .8$.

7.2 Incentives to Induce Forward Purchases

Consider a seller selling to a single buyer with demand D and let

$$Q^* = \inf\{Q : Pr(D \leq Q) \geq \frac{p-c}{p-s}\}$$

where c is the unit cost, p is the unit sale price and s is the salvage value to the seller. Let $\pi(Q^*)$ be the expected profit to the seller. The buyer will purchase $\min(D, Q^*)$ from the seller at an expected cost $pE\min(D, Q^*)$.

Suppose that the seller wants to induce the buyer to buy $\Delta \leq Q^*$ units *before* observing D. The seller wants to find the discount price, say $p(1-\delta)$, that would induce the buyer to agree. Assume that the the buyer orders Δ units at $p(1-\delta)$ before observing D. Then, after observing D, the buyer purchases $(\min(D, Q^*) - \Delta)^+$ additional units at price p if $D \geq \Delta$ or salvages $(\Delta - D)^+$ units at the salvage value s_b if $D < \Delta$. The buyer's expected cost is equal to

$$p(1-\delta)\Delta + p(E\min(D, Q^*) - \Delta)^+ - s_b E(\Delta - D)^+.$$

Given Δ, let $\delta(\Delta)$ be the solution to the equation

$$p(1-\delta)\Delta + p(E\min(D, Q^*) - \Delta)^+ - s_b E(\Delta - D)^+ = pE\min(D, Q^*).$$

Then

$$\delta(\Delta) = \frac{p - s_b}{p}\frac{E(\Delta - D)^+}{\Delta}.$$

At discount $\delta(\Delta)$, the buyer is indifferent between forward purchasing Δ units and waiting until he observes demand to purchase up to Q^* units at the full price. The buyer may prefer the forward purchase agreement if this reduces the variance of his cost. Notice that the variance of $p(\min(D, Q^*) - \Delta)^+$ is decreasing in Δ while the variance of $s_b(\Delta - D)^+$ is increasing in Δ.

The seller now faces the random demand $D' = (D - \Delta)^+$ instead of D. We will show that $q^* = Q^* - \Delta$ is the order quantity, in excess of Δ, that maximizes the seller's total expected profit:

$$\begin{aligned}\pi(q) &= (p(1-\delta) - c)\Delta + pE\min(D', q) + sE(q - D')^+ - cq \\ &= (p(1-\delta) - c)\Delta + (p - c)ED' - H(q)\end{aligned}$$

where $H(q) = (c - s)E(q - D')^+ + (p - c)E(D' - q)^+$.

Notice that the smallest optimal solution is given by

$$q^* = \inf\{q : Pr(D' \leq q) \geq \frac{p - c}{p - s}\}.$$

But then, $Pr(D \leq Q^*) = Pr(D' \leq Q^* - \Delta) = Pr(D' \leq q^*)$, so $q^* = Q^* - \Delta$.

As a result,

$$\begin{aligned}\pi(q^*) &= (p(1-\delta) - c)\Delta + (p - c)ED' - H(q^*) \\ &= p(1-\delta)\Delta + pED' - cE\max(D, \Delta) - H(q^*) \\ &= p(1-\delta)\Delta + p(ED' \pm E(D - Q^*)^+) - cE\max(D, \Delta) - H(q^*) \\ &= pE\min(D, Q^*) - s_bE(\Delta - D)^+ + pE(D - Q^*)^+ - cE(D + (\Delta - D)^+) - H(q^*) \\ &= (p - c)ED - (c - s_b)E(\Delta - D)^+ - H(q^*) \\ &= (p - c)ED - G(Q^*) + [G(Q^*) - H(q^*)] - (c - s_b)E(\Delta - D)^+ \\ &= (p - c)ED - G(Q^*) + (c - s)E(\Delta - D)^+ - (c - s_b)E(\Delta - D)^+ \\ &= (p - c)ED - G(Q^*) + (s_b - s)E(\Delta - D)^+.\end{aligned}$$

This shows that the scheme improves the expected profits for the seller when $s_b \geq s$. If $s_b \geq s$, the seller will select the largest Δ that results in a reduction of risk for the buyer.

The scheme seems to break down when $s_b < s$, but the seller can propose a different agreement where he buys back $(\Delta - D)^+$ units from the buyer at his own salvage value s. This is equivalent to delivering $\min(D, \Delta)$ instead of Δ and paying back $s(\Delta - D)^+$. The cash flow can be simplified further by embedding the buy back cash flow into the discount. This results in a steeper discount

$$\delta(\Delta) = \frac{E(\Delta - D)^+}{\Delta}$$

without buy backs where the seller delivers $\min(D, \Delta)$. Notice that in this case the buyer's expected cost remains constant while his risk is *decreasing* in Δ. The seller's expected profit now stays constant, but his risk first decreases and then increases with Δ. The goal of the seller is to select Δ to minimize the variance of his profit resulting in a win-win solution for the buyer and the seller.

7.2.1 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 $p(\min(D, Q^*) - \Delta)^+$, the buyer is interested in how the variance of the the random variable $(\min(D, Q^*) - \Delta)^+$ changes for values of $\Delta \leq Q^*$.

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,

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

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

How does the risk for the seller changes with Δ? We analyze this question for the case where $s_b = s$. The profit to the seller can be written as

$$p(1-\delta(\Delta))\Delta + p(\min(D,Q^*) - \Delta)^+ + s(Q^* - \min(D,Q^*))^+ - cQ^*.$$

As a result, the variance of the seller's profit depends also on the the covariance $\mathrm{Cov}(\min(D,Q^*), (\min(D,Q^*) - \Delta)^+)$. The derivative of the covariance is given by $(E\min(D,Q^*) - \Delta)\tilde{F}(\Delta) - E(\min(D,Q) - \Delta)^+]$, and the second derivative by $(\mu - \Delta)\tilde{F}'(\Delta)$. 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: Suppose that $c = 10, p = 20, s = 0$ and that D is uniform $[0, 100]$. Then $\mu = 50$, $\sigma = 28.87$ and $Q^* = 50$. Suppose that the parties agree to a forward contract for $\Delta = 20$ units, at a discount $\delta = 10\%$. The expected profits to the seller remains equal, but the risk is reduced. The risk to the seller, measured by the standard deviation of his profit goes down from 144.99 to 123.79. The risk to the buyer goes down from 320 to 251. A risk averse buyer would be happy to enter into this agreement because of the risk reduction effect. For the seller, the probability that profit falls is more than one standard deviation below its mean, drops from 20.3% to 14.9%.

8 Random Demand at Salvage Value

Consider now an extension where demand at the salvage price is a random variable V. Notice that the traditional Newsvendor model implicitly assumes that $Pr(V \ge Q) = 1$ for all Q. The Newsvendor model also implicitly assumes that $s < c$. Here we will allow $s \ge c$ but we will keep the assumption that $p > s$.

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

$$\pi(Q) = (p-c)\mu - H(Q)$$

where

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

Thus, the expected profit differs from that of the traditional Newsvendor Model only when $V \le (Q-D)^+$, or equivalently, when $V + D \le Q$ 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

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

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 $H'(Q^{nv}) = sPr(D+V \le Q^{nv}) \ge 0$, implying that there exists an optimal solution $Q^* \le Q^{nv}$. Consequently, if $Pr(D+V \le Q^{nv}) > 0$ then $Q^* < Q^{nv}$ so it is optimal to order fewer units than under the traditional 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

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

Consequently,

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

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

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

8.1 Revenue Management

A revenue management problem arises when Q is fixed and $p < s$. In this case, we need to decide how many units to make available for sale at the *low* fare p so that *enough* capacity is protected for sale at the *high* fare s to maximize expected revenues. Let q be the smallest integer such that $Pr(V \ge q) > p/s$. Then, it is optimal to protect q units for sale at the high fare s and to make $(Q-q)^+$ units available for sale at low fare p.

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.