

Key Concepts: Week 7 Lesson 1: Probabilistic Inventory Models

Learning Objectives

- Understanding of safety stock and its role in protecting for excess demand over lead time
- Ability to develop base stock and order-point, order-quantity continuous review policies
- Ability to determine proper safety factor, k , given the desired CSL or IFR or the appropriate cost penalty for CSOE or CIS

Lesson Summary:

In this lesson, we continue to develop inventory replenishment models when we have uncertain or stochastic demand. We built off of both the EOQ and the single period models to introduce three general inventory policies: the Base Stock Policy, the (s,Q) continuous review policy and the (R,S) periodic review policy (the R,S model will be explained in the next lesson). These are the most commonly used inventory policies in practice. They are imbedded within a company's ERP and inventory management systems.

To put them in context, here is the summary of the five inventory models covered so far:

- **Economic Order Quantity** —Deterministic Demand with infinite horizon
 - Order Q^* every T^* periods
 - Order Q^* when $IP = \mu_{DL}$
- **Single Period / Newsvendor** — Probabilistic Demand with finite (single period) horizon
 - Order Q^* at start of period where $P[x \leq Q]=CR$
- **Base Stock Policy** — Probabilistic Demand with infinite horizon
 - Essentially a one for one replenishment
 - Order what was demanded when it was demanded in the quantity it was demanded
- **Continuous Review Policy (s,Q)** — Probabilistic Demand with infinite horizon
 - This is event based – we order when, and if, inventory passes a certain threshold
 - Order Q^* when $IP \leq s$
- **Periodic Review Policy (R,S)** — Probabilistic Demand with infinite horizon
 - This is a time based policy in that we order on a set cycle
 - Order up to S units every R time periods

All of the models make trade-offs: EOQ between fixed and variable costs, Newsvendor between excess and shortage inventory, and the latter three between cost and level of service. The concept of level of service, LOS, is often murky and specific definitions and preferences vary between firms. However, for

our purposes, we can break them into two categories: targets and costs. We can establish a target value for some performance metric and then design the minimum cost inventory policy to achieve the level of service. The two metrics that we covered were Cycle Service Level (CSL) and Item Fill Rate (IFR). The second approach is to place a dollar amount on a specific type of stock out occurring and then minimize the total cost function. The two cost metrics we covered were Cost of Stockout Event (CSOE) and Cost of Item Short (CIS). They are related to each other.

Regardless of the metrics used, the end result is a safety factor, k , and a safety stock. The safety stock is simply $k\sigma_{DL}$. The term σ_{DL} is defined as the standard deviation of demand over lead time, but it is more technically the root mean square error (RMSE) of the forecast over the lead time. Most companies do not track their forecast error to the granular level that you require for setting inventory levels, so defaulting to the standard deviation of demand is not too bad of an estimate. It is essentially assuming that the forecast is the mean. Not too bad of an assumption.

Notation:

B_1 :	Cost associated with a stockout event (\$/event)
c :	Purchase cost (\$/unit)
c_t :	Ordering Costs (\$/order)
c_e :	Excess holding Costs (\$/unit/time); Equal to ch
c_s :	Shortage costs (\$/unit)
D :	Average Demand (units/time)
D_S :	Demand over short time period (e.g. week)
D_L :	Demand over long time period (e.g. month)
h :	Carrying or holding cost (\$/inventory \$/time)
L :	Replenishment Lead Time (time)
Q :	Replenishment Order Quantity (units/order)
T :	Order Cycle Time (time/order)
μ_{DL} :	Expected Demand over Lead Time (units/time)
σ_{DL} :	Standard Deviation of Demand over Lead Time (units/time)
k :	Safety Factor
s :	Reorder Point (units)
S :	Order up to Point (units)
R :	Review Period (time)
N :	Orders per Time or $1/T$ (order/time)
IP :	Inventory Position = Inventory on Hand + Inventory on Order – Backorders
IOH :	Inventory on Hand (units)
IOO :	Inventory on Order (units)
IFR :	Item Fill Rate (%)
CSL :	Cycle Service Level (%)
$CSOE$:	Cost of Stock Out Event (\$ / event)
CIS :	Cost per Item Short

- E[US]: Expected Units Short (units)
- G(k): Unit Normal Loss Function

Formulas:

Base Stock Policy

The Base Stock policy is a one-for-one policy. If I sell 4 items, I order 4 items to replenish the inventory. The policy determines what the stocking level, or the base stock, is for each item. The base stock, S^* , is the sum of the expected demand over the lead-time plus the RMSE of the forecast error over lead time multiplied by some safety factor k . The LOS for this policy is simply the Critical Ratio. Note that the excess inventory cost, c_e , in this case (and all models here) assumes you can use it later and is the product of the cost and the holding rate, ch .

- Optimal Base Stock, S^* : $S^* = \mu_{DL} + k_{LOS}\sigma_{DL}$
- Level of Service (LOS): $LOS = P[\mu_{DL} \leq S^*] = CR = \frac{c_s}{c_s + c_e}$

Continuous Review Policies (s, Q)

This is also known as the Order-Point, Order-Quantity policy and is essentially a two-bin system. The policy is “Order Q^* units when Inventory Position is less than the re-order point s ”. The re-order point is the sum of the expected demand over the lead-time plus the RMSE of the forecast error over lead-time multiplied by some safety factor k .

- Reorder Point: $s = \mu_{DL} + k\sigma_{DL}$
- Order Quantity (Q): Q is typically found through the EOQ formula

Level of Service Metrics

We present four methods for determining the appropriate safety factor, k , for use in any of the inventory models. They are Cycle Service Level, Cost per Stock Out Event, Item Fill Rate, and Cost per Item Short.

Cycle Service Level (CSL)

The CSL is the probability that there will not be a stock out within a replenishment cycle. This is frequently used as a performance metric where the inventory policy is designed to minimize cost to achieve an expected CSL of, say, 95%. Thus, it is one minus the probability of a stock out occurring. If I know the target CSL and the distribution (we will use Normal most of the time) then we can find the s that satisfies it using tables or a spreadsheet where $s = \text{NORMINVDIST}(\text{CSL}, \text{Mean}, \text{StandardDeviation})$ and $k = \text{NORMSINV}(\text{CSL})$.

$$\text{CSL} = 1 - \text{P}[\text{Stockout}] = 1 - \text{P}[X > s] = \text{P}[X \leq s]$$

Cost Per Stockout Event (CSOE) or B_1 Cost

The CSOE is related to the CSL, but instead of designing to a target CSL value, a penalty is charged when a stock out occurs within a replenishment cycle. The inventory policy is designed to minimize the total costs – so this balances cost of holding inventory explicitly with the cost of stocking out. Minimizing the total costs for k , we find that as long as $\frac{B_1 D}{c_e \sigma_{DL} Q \sqrt{2\pi}} > 1$, then we should set:

$$k = \sqrt{2 \ln \left(\frac{B_1 D}{c_e \sigma_{DL} Q \sqrt{2\pi}} \right)}$$

If $\frac{B_1 D}{c_e \sigma_{DL} Q \sqrt{2\pi}} < 1$, we should set k as low as management allows.

Item Fill Rate (IFR)

The IFR is the fraction of demand that is met with the inventory on hand out of cycle stock. This is frequently used as a performance metric where the inventory policy is designed to minimize cost to achieve an expected IFR of, say, 90%. If I know the target IFR and the distribution (we will use Normal most of the time) then we can find the appropriate k value by using the Unit Normal Loss Function, $G(k)$.

$$\text{IFR} = 1 - \frac{E[US]}{Q} = 1 - \frac{\sigma_{DL} G[k]}{Q}$$
$$G(k) = \frac{Q}{\sigma_{DL}} (1 - \text{IFR})$$

$G(k)$ is the Unit Normal Loss Function, which can be calculated in Spreadsheets as

$$G(k) = \text{NORMDIST}(k, 0, 1, 0) - k * (1 - \text{NORMSDIST}(k))$$

Cost per Item Short (CIS)

The CIS is related to the IFR, but instead of designing to a target IFR value, a penalty is charged when for each item short within a replenishment cycle. The inventory policy is designed to minimize the total costs – so this balances cost of holding inventory explicitly with the cost of stocking out. Minimizing the total costs for k , we find that as long as $\frac{Q c_e}{D c_s} \leq 1$, then we should find k such that:

$$P[\text{StockOut}] = P[x \geq k] = \frac{Qc_e}{Dc_s}$$

Otherwise, we should set k as low as management allows. In a spreadsheet, this becomes

$$k = \text{NORMSINV}\left(1 - \frac{Qc_e}{Dc_s}\right)$$

A Tip on Converting Times

You will typically need to convert annual forecasts to weekly demand or vice versa or something in between. This is generally very easy – but some students get confused at times:

Converting long to short (n is number of short periods within long):

$$\begin{aligned} E[D_S] &= E[D_L]/n \\ \text{VAR}[D_S] &= \text{VAR}[D_L]/n \\ \sigma_S &= \sigma_L/\sqrt{n} \end{aligned}$$

Converting from short to long:

$$\begin{aligned} E[D_L] &= nE[D_S] \\ \text{VAR}[D_L] &= n\text{VAR}[D_S] \\ \sigma_L &= \sqrt{n}\sigma_S \end{aligned}$$

Additional References:

Base stock and continuous inventory models are covered in Nahmias Chpt 5 and Silver, Pyke & Peterson Chpt 7, and Ballou Chpt 9.