Practical Tools for OR, ML and DS

Lecture 9: Inventory #3 (Q, R)-Policies

February 18th, 2020

EOQ model: demand occurs at a fixed known rate

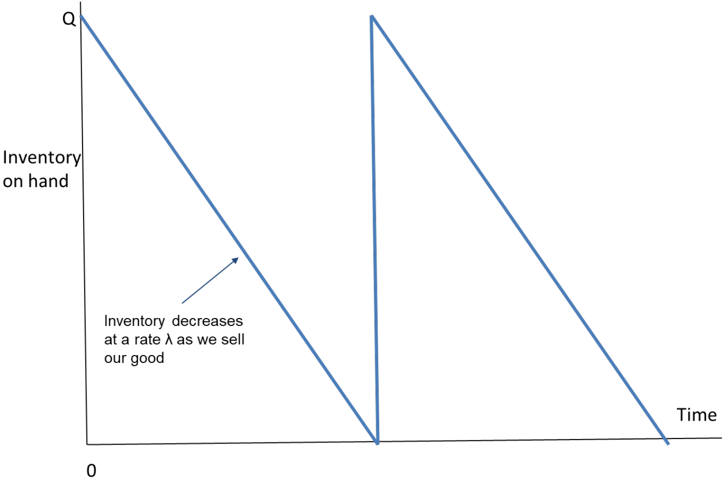

Today

Random demand, multiple periods (non-perishable goods)

Here's what demand looks like over time when the demand rate is random

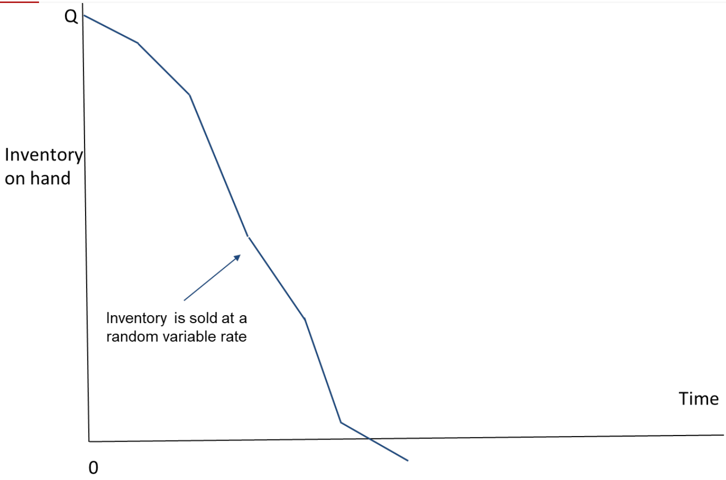

Recall goal: keeping costs low

We usually consider 3 kinds of costs when managing inventory:

- Holding costs
- Order costs
- *Penalty costs* (these were 0 in EOQ, because there were no stock-outs — we could determine exactly when to order so that inventory was always nonnegative)

Model of random demand

$$\lambda_t = \text{mean of demand in time } t$$
$$\sigma_t^2 = \text{variance of demand in time } t$$

- Demand is independent in each time period
- These two parameters are a reasonable way to characterize demand
- Most of our analysis will assume λ_t and σ_t^2 don't change with t. When we do this, we'll write them as λ and σ^2.

Notation

$$D_t = \text{demand in time } t$$
$$C(t) = \text{total demand by time } t$$
$$C(t) = D_1 + D_2 + \ldots + D_t$$

Example

- Mean demand is $\lambda = 30$ units per week
- Variance of the weekly demand is $\sigma^2 = 15$
- Suppose the lead time is $\tau = 5$ weeks.

What is $\mathbb{E}[C(\tau)]$?

1. $15 \times 5 = 75$
2. $30 \times 5 = 150$
3. $15 \times 5^2 = 375$
4. $30 \times 5^2 = 750$
5. $30/2 = 15$

Example

- Mean demand is $\lambda = 30$ units per week
- Variance of the weekly demand is $\sigma^2 = 15$
- Suppose the lead time is $\tau = 5$ weeks.

What is $\mathbb{V}\text{ar}[C(\tau)]$?

1. $15 \times 5 = 75$
2. $30 \times 5 = 150$
3. $15 \times 5^2 = 375$
4. $30 \times 5^2 = 750$
5. $30/2 = 15$

(Q, R) policy

- The policies we will consider are (Q, R) *policies*.
- In a (Q, R) policy, we set two parameters:
  - the replenishment inventory level R
  - the size of each replenishment order Q.
- The EOQ model also had these parameters

  In EOQ we could calculate one from the other because the demand was deterministic, and we could just aim for having 0 inventory when the new order arrived
- In the (Q, R) model will set Q and R to deal with *random* demand, also including penalties for having not enough inventory

Here's how a (Q, R) policy works

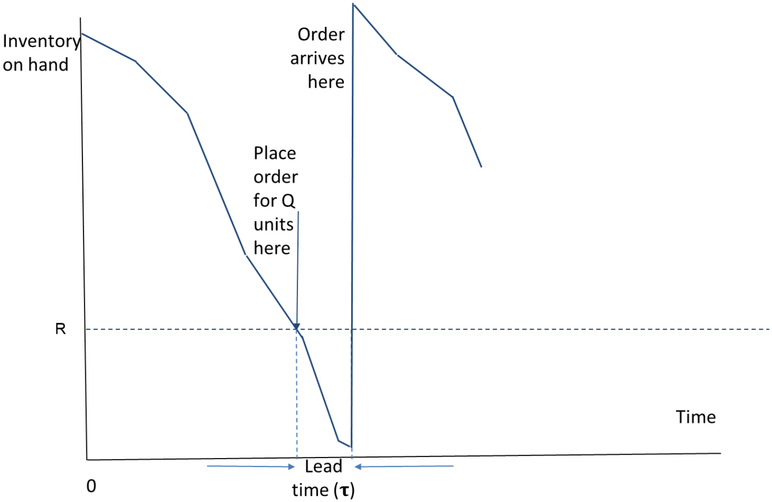

Here's how a (Q, R) policy works

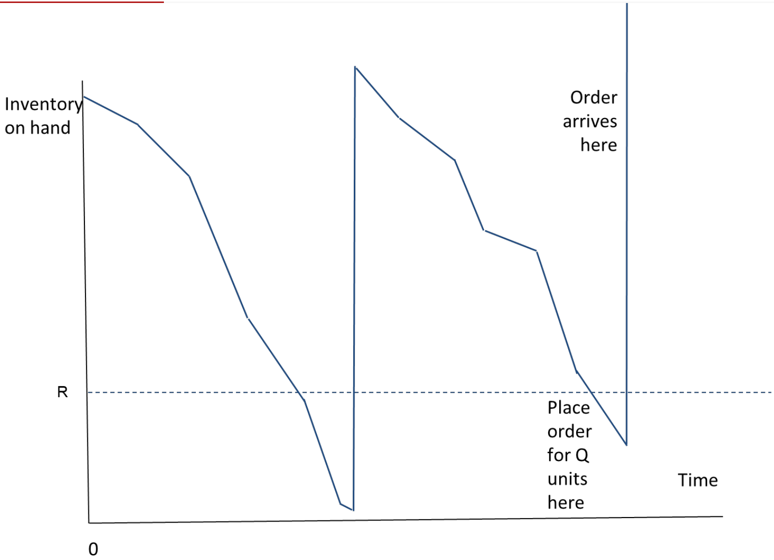

Continuous review can be a lot of work in the real world

- (Q, R) policies assume the inventory level is reviewed continuously
- To implement we either must have:
    - A computer system with point-of-sale scanners that updates inventory levels after every sale
    - A person assigned to monitor inventory every hour of every day
    - As soon as we run out, we have to order more
- This can be annoying, and a lot of work

We allow stockouts, but pay a penalty

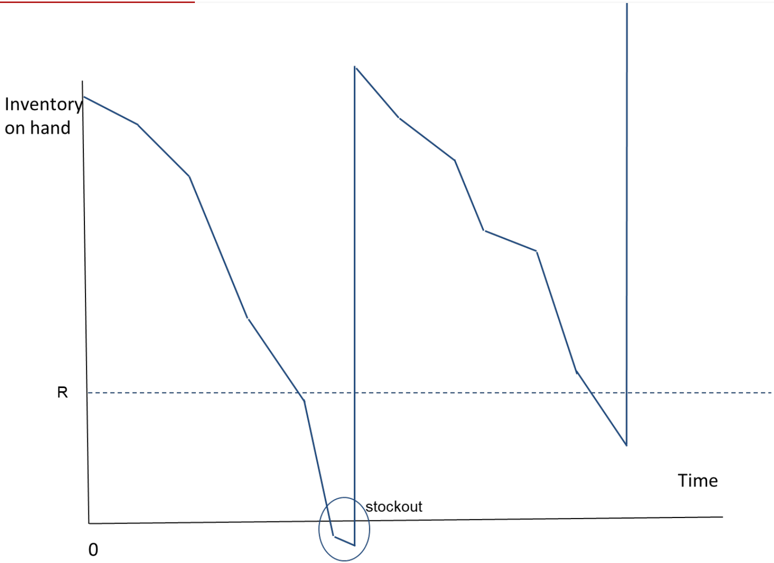

Intuitively, here's what we are doing when we set Q and R

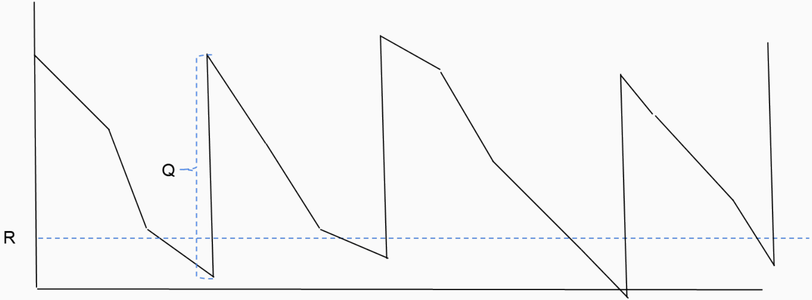

Bigger R means more inventory, and fewer stockouts

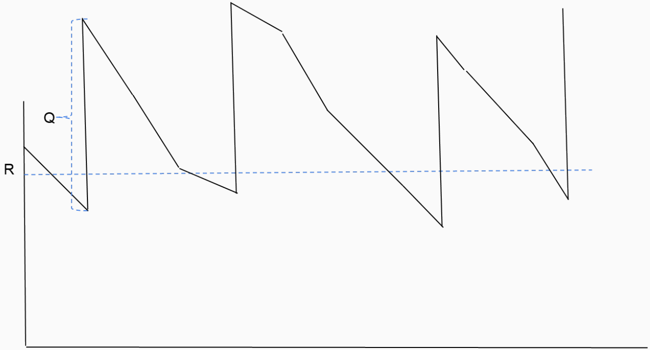

Bigger R means more inventory, and fewer stockouts

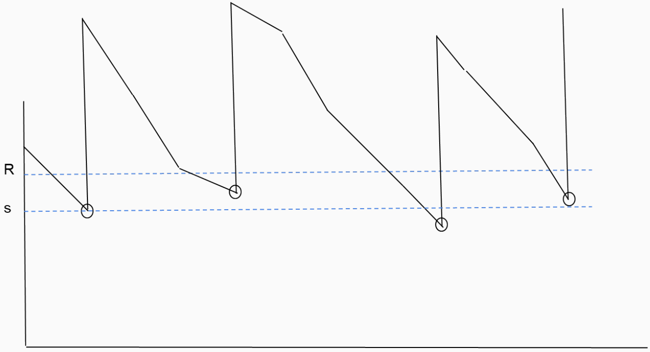

The average value of these minimum inventory values is called the "safety stock" and is indicated by "s"

More safety stock reduces stockouts.

Bigger Q means more inventory and fewer orders

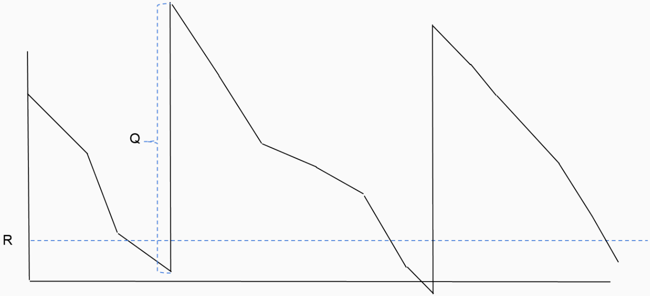

Bigger Q means more cycle stock and fewer orders

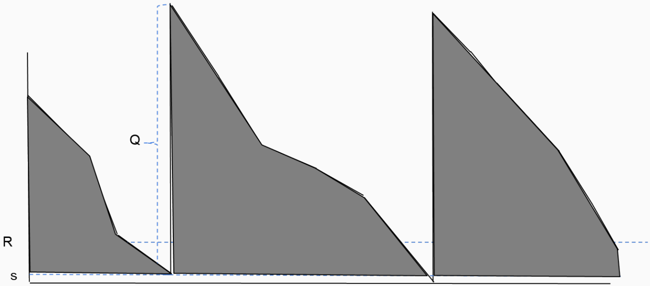

The inventory in excess of s is called the “cycle stock”

This stock is depleted at the end of the cycle, and replenished at the beginning of the cycle

Intuitively, here's what we are doing when we set Q and R

- Q controls the tradeoff between order frequency and average inventory levels.
  If Q is large, there are fewer orders and larger inventory levels.
- R controls the tradeoff between inventory levels and likelihood of stockouts.
  If R is large, there is a low probability of stocking out, but the average inventory level will be higher.

Intuitively, here's what we are doing when we set Q and R

- Q affects cycle stock, the inventory held to avoid excessive replenishment costs.
- R affects safety stock, the inventory held to avoid stockouts.
- The EOQ model kept cycle stock but no safety stock.

Optimizing the total cost

We will derive an expression for the expected cost per unit time in terms of our decision variables Q and R, and then find optimal values of Q and R to minimize this cost.

We'll model our cost like this

- Ordering cost: An order for Q units costs $K + cQ$ (same as EOQ)
- Holding cost: Holding cost are incurred at a rate of h per unit inventory per unit time (same as EOQ)
- Penalty: If we run out of inventory, we place the demand we cannot satisfy on backorder and pay a cost of p for each backordered unit (this is new)

Note that when items are backordered, our inventory position is negative, and there is no holding cost

We model our demand like this

$$
\begin{aligned}
D_t &= \text{ demand in time } t\text{, independent across } t \\
\mathbb{E}[D_t] &= \lambda \quad \text{(same for every } t\text{)} \\
\mathbb{V}\text{ar}[D_t] &= \sigma^2 \quad \text{(same for every } t\text{)} \\
C(t) &= D_1 + D_2 + \ldots + D_t
\end{aligned}
$$

Order Cost

- To compute the order cost, we want to know the number of orders per unit time
- This will be more complex than in EOQ because the time between orders varies

Order Cost

- $I(t) =$ inventory on hand at time t
- $O(t) =$ number of orders before time t
- $I(t) = I(0) + Q \times O(t) - C(t)$
- $O(t) = \dfrac{C(t) + I(t) - I(0)}{Q}$
- $\mathbb{E}(O(t)) = \dfrac{\lambda t + \mathbb{E}[I(t)] - I(0)}{Q}$
- $\mathbb{E}(O(t)/t) = \lambda/Q + \dfrac{\mathbb{E}[I(t)] - I(0)}{tQ} \to \lambda/Q$ (as $t \to \infty$)
- This is the expected number of orders per unit time in the long run — we will call it the cycle frequency
  We define $T = Q/\lambda$ and call it the (expected) cycle period or cycle length

Expected Order Cost per Unit Time

Each cycle requires one order, and has a cost of $K + cQ$.

Therefore, the expected cost of ordering per unit time is

$$(K + cQ) \times (\lambda/Q) = K\lambda/Q + c\lambda$$

Expected Holding Cost per Unit Time

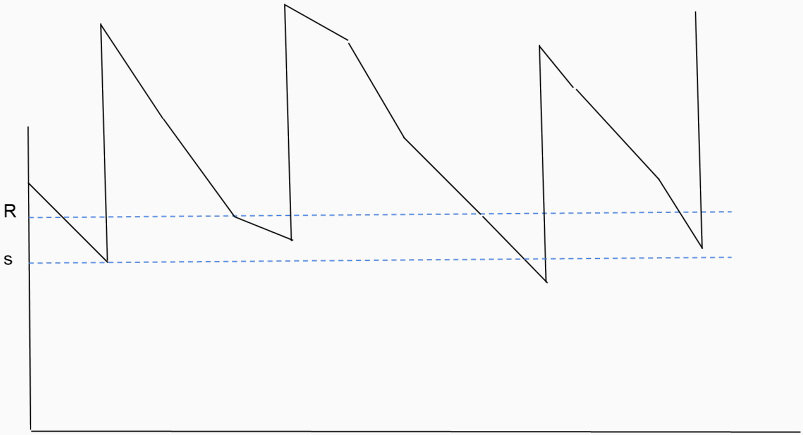

We'd like to compute the average inventory under this curve

Expected Holding Cost per Unit Time

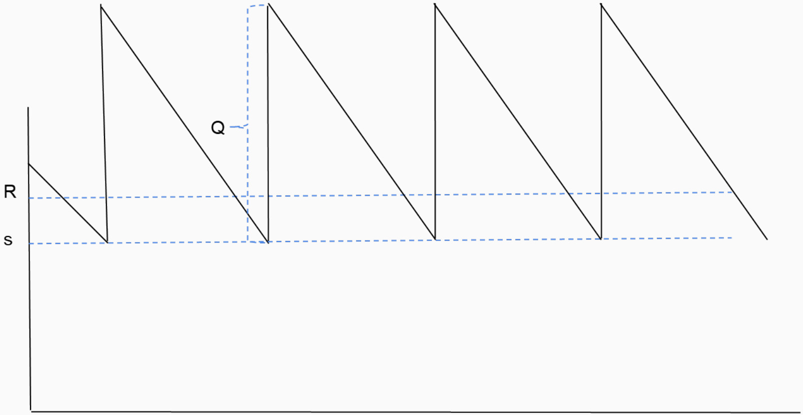

This is what it would look like if demand were deterministic

Expected Holding Cost per Unit Time

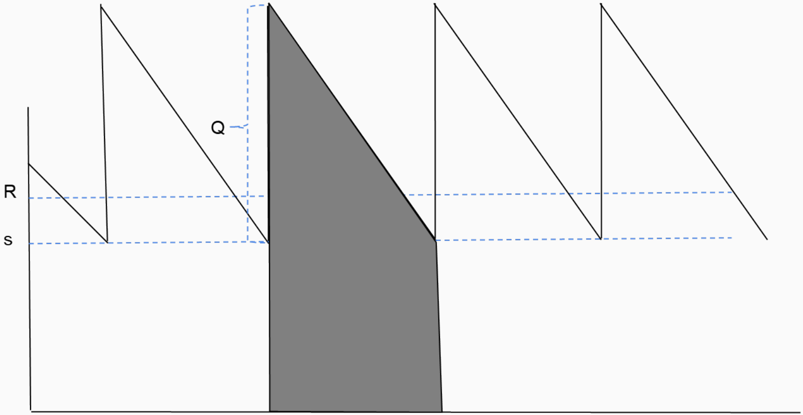

The average height of the curve during each cycle would be Q/2 + s

Expected Holding Cost per Unit Time

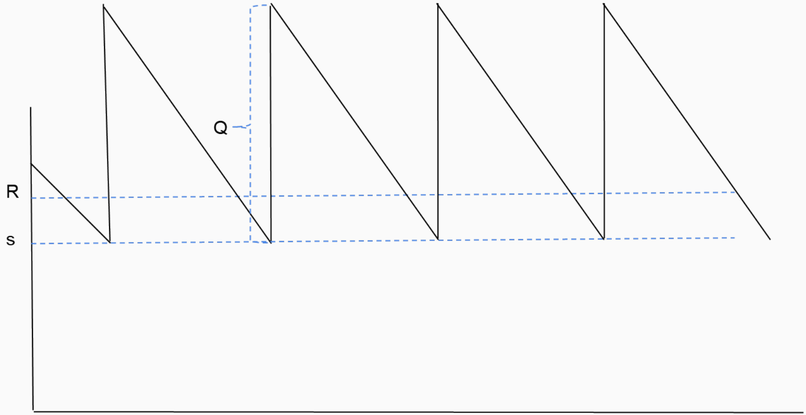

The long-run average inventory level would be Q/2 + s

Expected Holding Cost per Unit Time

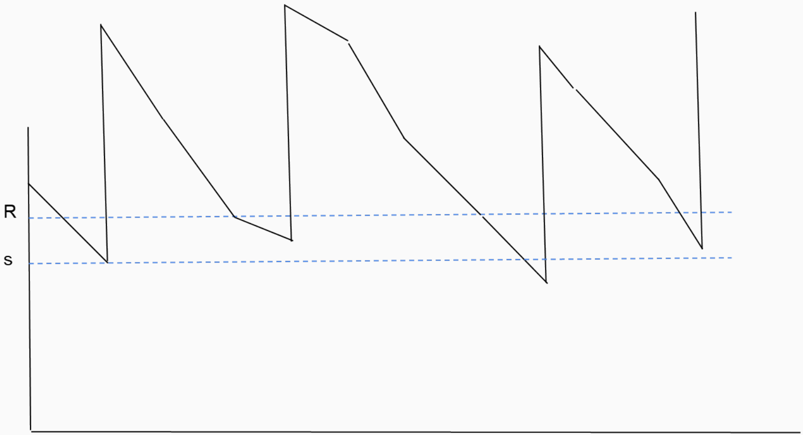

Amazingly enough, the long-run average inventory level is also $Q/2 + s$ when demand is stochastic

Knowing why is outside of the scope of this course.

What is the safety stock?

What is the safety stock s?

1. $R - \lambda T$
2. $Q - \lambda T$
3. $R - \lambda \tau$
4. $Q - \lambda \tau$
5. $Q/2$

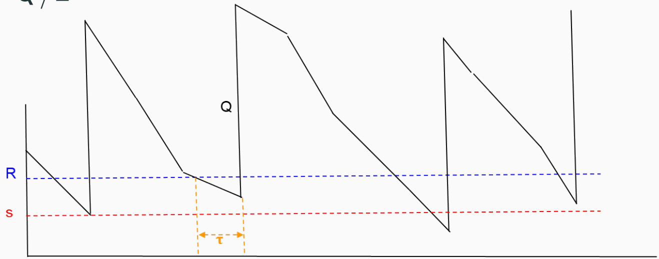

Expected Holding Cost per Unit Time

The (long-run) average inventory level is:

$$s + \frac{Q}{2} = R - \lambda\tau + \frac{Q}{2}.$$

Expected Holding Cost per Unit Time

- Next, we find the expected holding cost per unit time by multiplying the long-run average inventory by the holding cost per unit and per unit time.
- The problem is that we have not restricted the inventory to positive values, and our next step assumes positive values of inventory at all times!
- However, stockouts and the resulting backorders are usually rare, so this approximation is not likely to result in major errors.

Expected Holding Cost per Unit Time

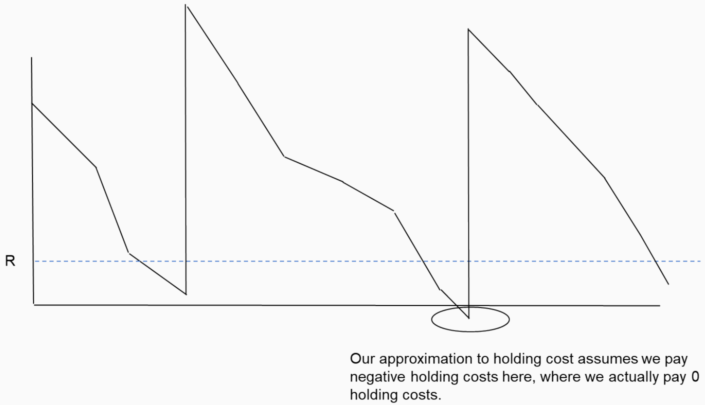

Our approximation to holding cost assumes we pay negative holding costs here, where we actually pay zero holding costs.

Expected Holding Cost per Unit Time

$$h\left(R - \lambda\tau + \frac{Q}{2}\right)^{+} \approx h\left(R - \lambda\tau + \frac{Q}{2}\right)$$

Our approximation to holding cost will be good enough because we will allow only a small number of stockouts.

Stockout Cost

- While we allow backorders, they are undesirable and we should limit them.
- To do this, we include a penalty cost in the model for each stockout.
- Each unit of demand that cannot be met from stock will incur a penalty cost of p.

How many stockout incidences we expect to see (per unit time)?

- We are exposed to stockout risk during the lead time.
- Before placing the replenishment order, we have at least R units in stock, and there is no risk of running out.
- Once the replenishment order is placed, we know that we will have to wait a known and fixed amount of time before any more stock comes in. We get a stockout if the demand over the lead time exceeds the reorder quantity R.

Stockout Cost

- Let N be the number of items short in a particular cycle.
- We emphasize that N is a function of R by writing it as $N(R)$.
- Let t be the time when inventory hits R in this cycle.
- $S = D_{t+1} + D_{t+2} + \ldots + D_{t+\tau}$ is the demand during the lead time in this cycle
- $N(R) = \max\{S - R, 0\} = (S - R)^+$ (for this particular cycle)

Stockout Cost

- Assuming demand is independent and identically distributed (over time), we have
  $S = C(\tau)$ (the total amount of demand in τ time units)
- Now we can talk about the expected number of stock outs in any cycle (because demand is time independent). We denote this by $n(R)$:
  $n(R) = \mathbb{E}N(R) = \mathbb{E}[(S - R)^+] = \mathbb{E}[(C(\tau) - R)^+]$
- So the expected number of units short per unit time is
  $n(R)/\mathbb{E}(\text{length of a cycle}) =$

Stockout Cost

- Assuming demand is independent and identically distributed (over time), we have
  $S = C(\tau)$ (the total amount of demand in τ time units)
- Now we can talk about the expected number of stock outs in any cycle (because demand is time independent). We denote this by $n(R)$:
  $n(R) = \mathbb{E}N(R) = \mathbb{E}[(S - R)^+] = \mathbb{E}[(C(\tau) - R)^+]$
- So the expected number of units short per unit time is
  $n(R)/\mathbb{E}(\text{length of a cycle}) = n(R)/T =$

Stockout Cost

- Assuming demand is independent and identically distributed (over time), we have
  $S = C(\tau)$ (the total amount of demand in τ time units)
- Now we can talk about the expected number of stock outs in any cycle (because demand is time independent). We denote this by $n(R)$:
  $n(R) = \mathbb{E}N(R) = \mathbb{E}[(S - R)^+] = \mathbb{E}[(C(\tau) - R)^+]$
- So the expected number of units short per unit time is
  $n(R)/\mathbb{E}(\text{length of a cycle}) = n(R)/T =$
  $n(R)/(Q/\lambda) = \lambda n(R)/Q$

Stockout Cost

- Assuming demand is independent and identically distributed (over time), we have $S = C(\tau)$ (the total amount of demand in τ time units)
- Now we can talk about the expected number of stock outs in any cycle (because demand is time independent). We denote this by $n(R)$: $n(R) = \mathbb{E}N(R) = \mathbb{E}[(S - R)^+] = \mathbb{E}[(C(\tau) - R)^+]$
- So the expected number of units short per unit time is $n(R)/\mathbb{E}(\text{length of a cycle}) = n(R)/T = n(R)/(Q/\lambda) = \lambda n(R)/Q$
- Stockout *cost*: penalty for a stockout is p per unit.
- Thus, the expected total stockout cost per unit time is

$$\frac{p\lambda n(R)}{Q}$$

Total Cost

The total cost per unit time is found by summing the previous expressions:

$$G(Q,R) = \underbrace{\frac{K\lambda}{Q} + c\lambda}_{\text{expected order cost per unit time}} + \underbrace{h\left(R - \lambda\tau + \frac{Q}{2}\right)}_{\text{expected holding cost per unit time } (\approx)} + \underbrace{\frac{p\lambda n(R)}{Q}}_{\text{expected stockout cost per unit time}}$$

Total Cost

The total cost per unit time is found by summing the previous expressions:

$$G(Q,R) = \underbrace{\frac{K\lambda}{Q} + c\lambda}_{\text{expected order cost per unit time}} + \underbrace{h\left(R - \lambda\tau + \frac{Q}{2}\right)}_{\text{expected holding cost per unit time } (\approx)} + \underbrace{\frac{p\lambda n(R)}{Q}}_{\text{expected stockout cost per unit time}}$$

This is what we want to minimize

Total Cost

The total cost per unit time is found by summing the previous expressions:

$$G(Q,R) = \underbrace{\frac{K\lambda}{Q} + c\lambda}_{\text{expected order cost per unit time}} + \underbrace{h\left(R - \lambda\tau + \frac{Q}{2}\right)}_{\text{expected holding cost per unit time } (\approx)} + \underbrace{\frac{p\lambda n(R)}{Q}}_{\text{expected stockout cost per unit time}}$$

This is what we want to minimize

The decision variables are Q and R

Optimizing $G(Q, R)$

Find critical points: $\dfrac{\partial G(Q, R)}{\partial Q} = 0$ and $\dfrac{\partial G(Q, R)}{\partial R} = 0$.

This will be a minimizer, because the cost is convex (no proof).

Optimizing $G(Q, R)$

$$G(Q,R) = \frac{K\lambda}{Q} + c\lambda + h\Big(R - \lambda\tau + \frac{Q}{2}\Big) + \frac{p\lambda n(R)}{Q}$$

$$\frac{\partial G(Q,R)}{\partial R} =$$

Optimizing $G(Q, R)$

$$G(Q, R) = \frac{K\lambda}{Q} + c\lambda + h\Big(R - \lambda\tau + \frac{Q}{2}\Big) + \frac{p\lambda n(R)}{Q}$$

$$\frac{\partial G(Q, R)}{\partial R} = h + \frac{p\lambda}{Q}\frac{dn(R)}{dR}$$

What is $\dfrac{dn(R)}{dR}$?

Recall: $n(R) = \mathbb{E}[(C(\tau) - R)^+]$

What is $\frac{dn(R)}{dR}$?

Recall: $n(R) = \mathbb{E}[(C(\tau) - R)^+]$

$$
\begin{aligned}
\frac{dn(R)}{dR} &= \frac{d}{dR}\mathbb{E}[(C(\tau) - R)^+] \\
&= \mathbb{E}[\frac{d}{dR}(C(\tau) - R)^+] \text{ (under certain assumptions)} \\
&= \mathbb{E}[-\mathbb{1}_{\{C(\tau)>R\}}] \\
&= -\mathbb{P}(C(\tau) > R) \\
&= -(1 - F(R)) \\
&= F(R) - 1
\end{aligned}
$$

where $F(\cdot)$ is the cdf of $C(\tau)$, i.e., $F(x) = \mathbb{P}(C(\tau) \leq x)$.

Optimizing $G(Q, R)$

$$
\begin{aligned}
G(Q,R) &= \frac{K\lambda}{Q} + c\lambda + h\Big(R - \lambda\tau + \frac{Q}{2}\Big) \\
&\quad + \frac{p\lambda n(R)}{Q} \\
\frac{\partial G(Q,R)}{\partial R} &= h + \frac{p\lambda}{Q}\frac{dn(R)}{dR} \\
&= h + \frac{p\lambda}{Q}(F(R) - 1).
\end{aligned}
$$

Optimizing $G(Q, R)$

$$\frac{\partial G(Q,R)}{\partial R} = h + \frac{p\lambda}{Q}(F(R) - 1)$$

So $\dfrac{\partial G(Q,R)}{\partial R} = 0$ means R is so that

$$F(R) = -h\frac{Q}{p\lambda} + 1.$$

Optimizing $G(Q, R)$

When we know the inverse of $F(x) = \mathbb{P}(C(\tau) \leq x)$, then we have an explicit expression for R such that $\dfrac{\partial G(Q, R)}{\partial R} = 0$:

$$R = F^{-1}(1 - h\frac{Q}{p\lambda})$$

Optimizing $G(Q, R)$

When we know the inverse of $F(x) = \mathbb{P}(C(\tau) \leq x)$, then we have an explicit expression for R such that $\dfrac{\partial G(Q,R)}{\partial R} = 0$:

$$R = F^{-1}(1 - h\frac{Q}{p\lambda})$$

provided we know Q.

Optimizing $G(Q, R)$

When we do not know the inverse of $F(x) = \mathbb{P}(C(\tau) \leq x)$, then we can still use bisection search to find R such that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$

Optimizing $G(Q, R)$

When we do not know the inverse of $F(x) = \mathbb{P}(C(\tau) \leq x)$, then we can still use bisection search to find R such that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$

provided we know Q.

Optimizing $G(Q, R)$

Find critical points: $\dfrac{\partial G(Q, R)}{\partial Q} = 0$ and $\dfrac{\partial G(Q, R)}{\partial R} = 0$.

Optimizing $G(Q, R)$

$$G(Q,R) = \frac{K\lambda}{Q} + c\lambda + h\Big(R - \lambda\tau + \frac{Q}{2}\Big) + \frac{p\lambda n(R)}{Q}$$

$$\frac{\partial G(Q,R)}{\partial Q} =$$

Optimizing $G(Q, R)$

$$G(Q,R) = \frac{K\lambda}{Q} + c\lambda + h\Big(R - \lambda\tau + \frac{Q}{2}\Big) + \frac{p\lambda n(R)}{Q}$$

$$\frac{\partial G(Q,R)}{\partial Q} = -\frac{K\lambda}{Q^2} + \frac{h}{2} - \frac{p\lambda n(R)}{Q^2}$$

Optimizing $G(Q, R)$

$$\frac{\partial G(Q,R)}{\partial Q} = -\frac{K\lambda}{Q^2} + \frac{h}{2} - \frac{p\lambda n(R)}{Q^2} = 0$$

Optimizing $G(Q, R)$

$$\frac{\partial G(Q,R)}{\partial Q} = -\frac{K\lambda}{Q^2} + \frac{h}{2} - \frac{p\lambda n(R)}{Q^2} = 0$$

Solving for Q:

$$Q^2 = \frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)$$

So

$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Optimizing $G(Q,R)$

$$\frac{\partial G(Q,R)}{\partial Q} = -\frac{K\lambda}{Q^2} + \frac{h}{2} - \frac{p\lambda n(R)}{Q^2} = 0$$

Solving for Q:

$$Q^2 = \frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)$$

So

$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

provided we know R (or more precisely $n(R)$).

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$

$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

are true simultaneously.

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$
$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

are true simultaneously.

Solve iteratively!

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$
$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Start with $Q_0 = \sqrt{\frac{2K\lambda}{h}}$.

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$

$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Start with $Q_0 = \sqrt{\frac{2K\lambda}{h}}$.

Find R_0 so that

$$F(R_0) = 1 - h\frac{Q_0}{p\lambda}$$

(using bisection search, or F^{-1} if you know it.)

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$

$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Calculate $Q_1 = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R_0)\Big)}$.

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$
$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Calculate $Q_1 = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R_0)\Big)}$.

Find R_1 so that

$$F(R_1) = 1 - h\frac{Q_1}{p\lambda}$$

(using bisection search, or F^{-1} if you know it.)

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$
$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Calculate $Q_2 = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R_1)\Big)}$.

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$

$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Calculate $Q_2 = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R_1)\Big)}$.

Find R_2 so that

$$F(R_2) = 1 - h\frac{Q_2}{p\lambda}$$

(using bisection search, or F^{-1} if you know it.)

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$
$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Calculate $Q_3 = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R_2)\Big)}$.

Optimizing $G(Q, R)$

Want Q and R so that

$$F(R) = 1 - h\frac{Q}{p\lambda}$$
$$Q = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R)\Big)}$$

Calculate $Q_3 = \sqrt{\frac{2}{h}\Big(K\lambda + p\lambda n(R_2)\Big)}$.

Find R_3 so that

$$F(R_3) = 1 - h\frac{Q_3}{p\lambda}$$

(using bisection search, or F^{-1} if you know it.)

Optimizing $G(Q, R)$

Etcetera! Keep going around in circles!

Optimizing $G(Q, R)$

Stop when $Q_{n+1} \approx Q_n$ and $R_{n+1} \approx R_n$ (up to some specified accuracy)

Done!

We developed a method to find optimal Q and R in this model, with assumptions

- Holding costs: rate of h
- Order costs: fixed component K and variable component c
- Penalty costs: p per unit stock out

Done?

We developed a method to find optimal Q and R in this model, with assumptions

- Holding costs: rate of h
- Order costs: fixed component K and variable component c
- Penalty costs: p per unit stock out

How to specify p?