Chapter 16. Service Levels and Lead Times in Supply Chains: The Order-up-to Inventory Model

Medtronic's InSync pacemaker supply chain and objectives

- Supply chain:
  - One distribution center (DC) in Mounds View, MN.
  - About 500 sales territories throughout the country.
    - Consider Susan Magnotto's territory in Madison, Wisconsin.

- Objective:
  - Because the gross margins are high, develop a system to minimize inventory investment while maintaining a very high service target, a 99.9% in-stock probability.

Timing in the order up-to model

- Time is divided into periods of equal length, e.g., one hour, one month.
- During a period the following sequence of events occurs:
  - A replenishment order can be submitted.
  - Inventory is received.
  - Random demand occurs.
- Lead times:
  - An order is received after a fixed number of periods, called the lead time.
  - Let l represent the length of the lead time.

An example with $l = 1$

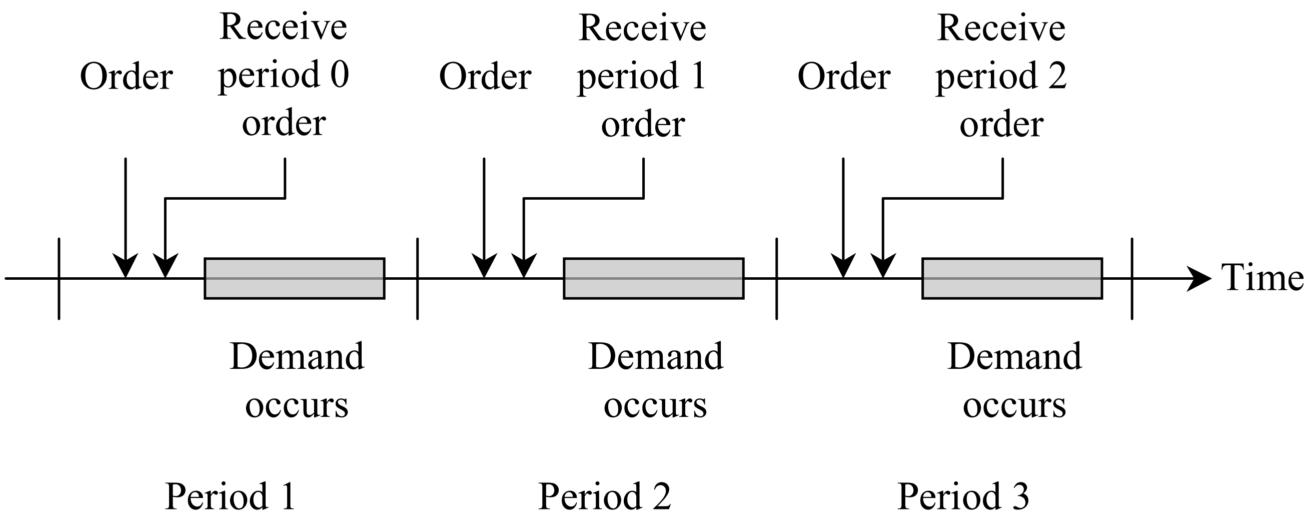

Order up-to model vs. newsvendor model

- Both models have uncertain future demand, but there are differences…

	Newsvendor	Order up-to
Inventory obsolescence	After one period	Never
Number of replenishments	One (maybe two or three with some reactive capacity)	Unlimited
Demand occurs during replenishment	No	Yes

- Newsvendor applies to short life cycle products with uncertain demand and the order up-to applies to long life cycle products with uncertain, but stable demand.

16.2 The Order Up-To Model:

Model design and implementation

Order up-to model definitions

- ***On-order inventory / pipeline inventory*** = number of units that have been ordered but have not been received.
- ***On-hand inventory*** = number of units physically in inventory ready to serve demand.
- ***Backorder*** = total amount of demand that has not been satisfied:
  - All backordered demand is eventually filled, i.e., there are no lost sales.
- ***Inventory level*** = *On-hand inventory* - *Backorder*.
- ***Inventory position*** = *On-order inventory* + *Inventory level*.
- ***Order up-to level, S***
  - maximum inventory position we allow.
  - sometimes called the ***base stock level***.

Order up-to model implementation

- *Each period's order quantity = S – Inventory position*
  - Suppose *S* = 4.
    - *If a period begins with an inventory position = -3, how many units will be ordered?*
    - *If demand were 10 in period 1, then how many units are ordered at the start of period 2?*

- A period's order quantity = the previous period's demand:
  - The order up-to model is a *pull system* because inventory is ordered in response to demand.
  - The order up-to model is sometimes referred to as a ***1-for-1 ordering policy*****.**

16.4 The Order Up-To Model:

Choosing demand distributions

InSync demand and inventory at the DC

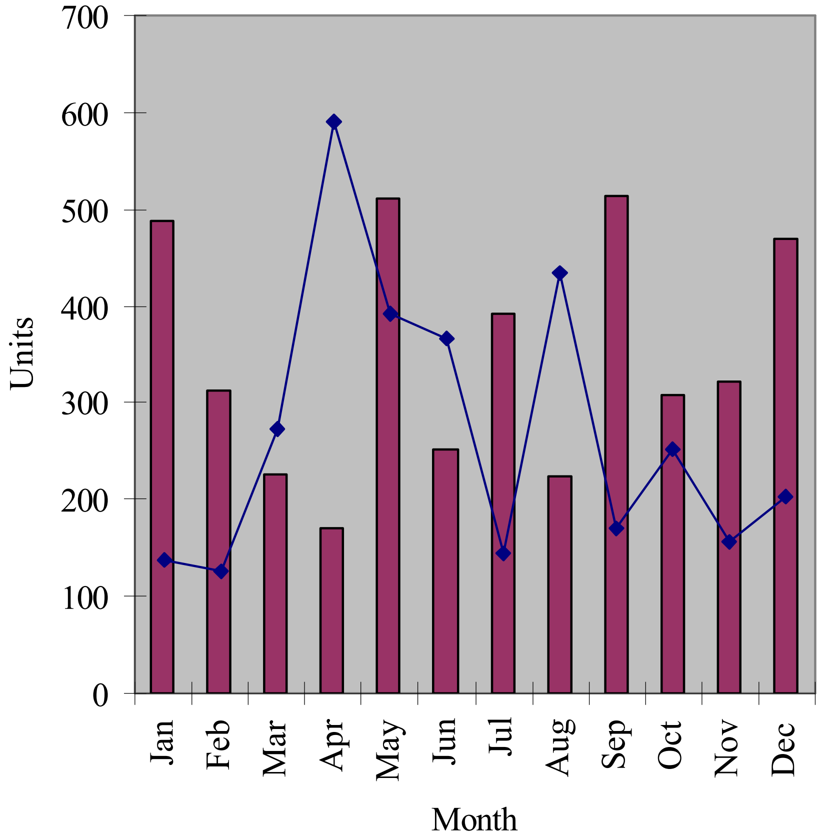

Monthly implants (columns) and end of month inventory (line)

Average monthly demand = 349 units

Standard deviation of monthly demand = 122.4

InSync demand and inventory in Susan's territory

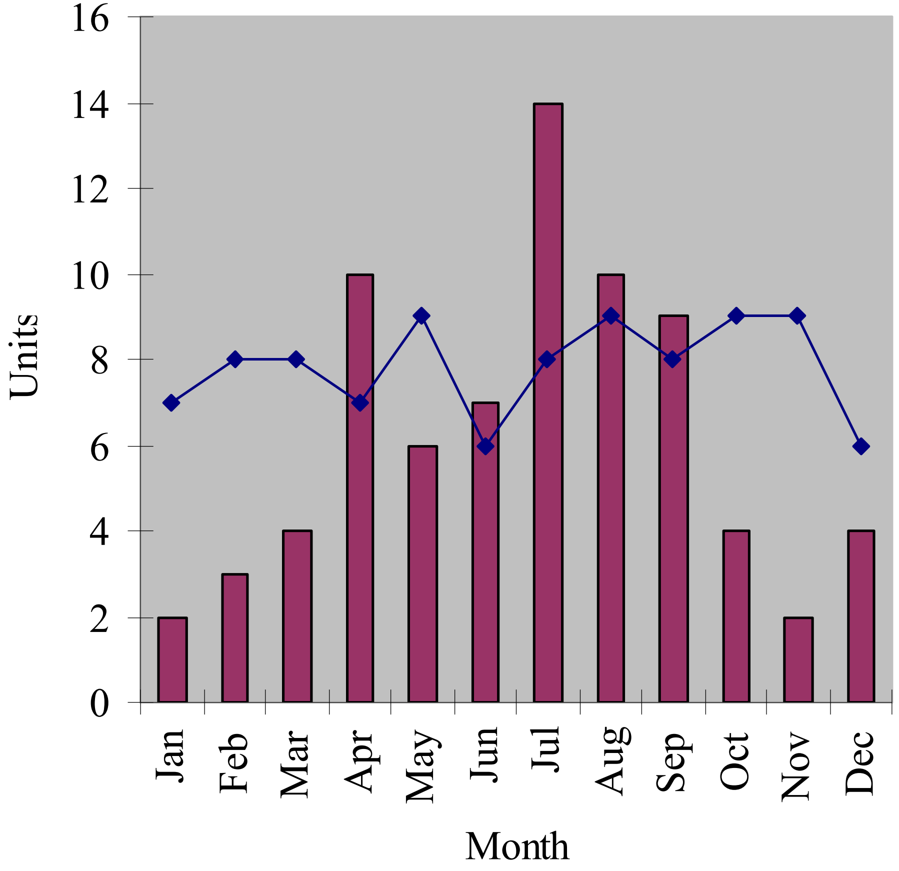

Monthly implants (columns) and end of month inventory (line)

Total annual demand = 75 units

Average daily demand = 0.29 units (75/260), assuming 5 days per week.

The Poisson demand distribution works better for slow moving items.

Poisson distribution

- The outcome of a random variable with a Poisson distribution is *discrete* (0,1,2,3,…) and *positive* (i.e., no negative outcomes).

- The Poisson is characterized by a *single parameter*, its mean:
  - Recall, the Normal distribution is defined by two parameters, the mean and standard deviation.

- The *standard deviation* of a Poisson distribution is equal to the square root of its mean.

- The Poisson is related to the exponential distribution:
  - If the inter-arrival times of customers are exponentially distributed, then the number of customers that arrive in a fixed interval of time is Poisson distributed.

- The Poisson is ideal for describing the demand of *slow moving products*, e.g., products that have average sales of 20 or fewer units over a particular period of time.

The shape of the Poisson

- The Poisson's density function takes on different shapes, depending on its mean
  - With a large mean (more than 10) the Poisson takes on a "bell shape"
  - With a small mean (less than 5) the Poisson takes on different shapes that slowly change in the direction of a bell curve.

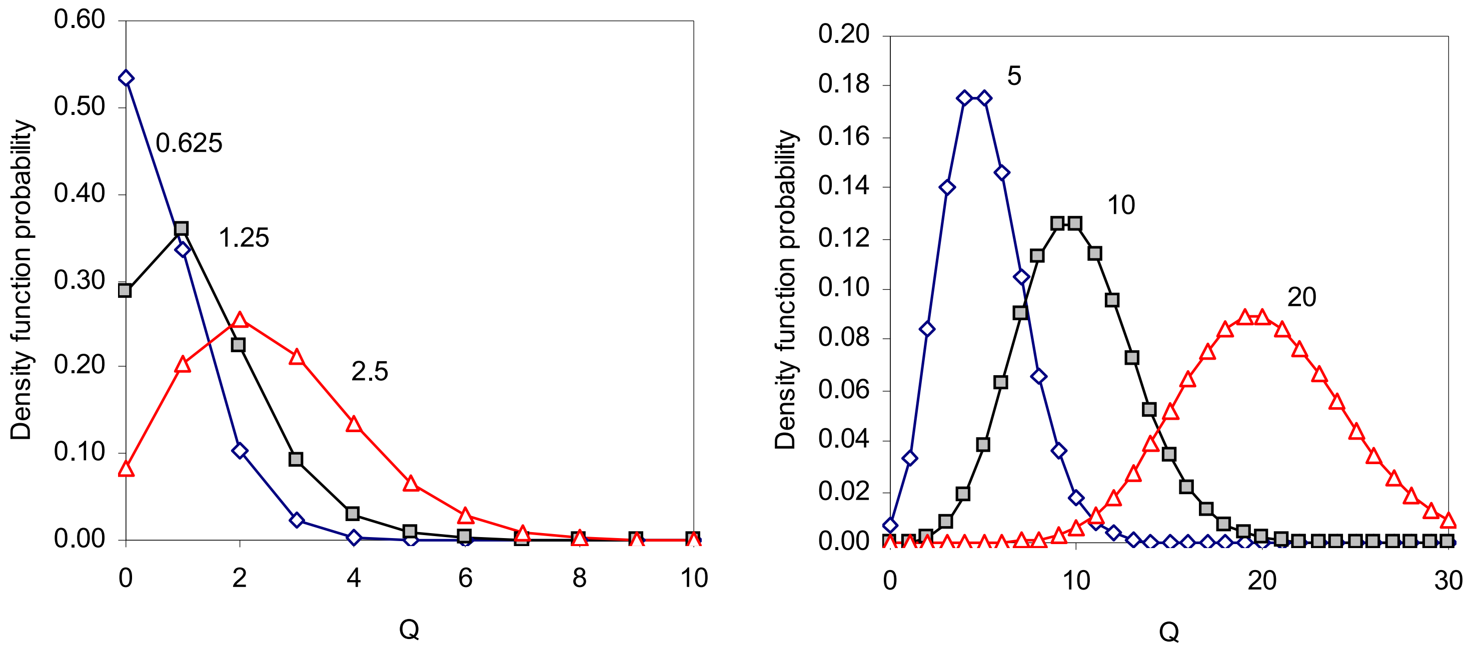

Density function of 6 Poisson distributions with means 0.625, 1.25, 2.5, 5, 10 and 20

Poisson's Distribution and Loss Functions

- There is no equivalent of the "z-statistic" for the Poisson like there is for the Normal distribution.
  - Hence, you need a distribution and loss function table for each Poisson distribution.

- Consider the density function, $f(Q)$, distribution function, $F(Q)$, and the loss function, $L(Q)$ of the Poisson with mean 1.25 in the table to the right:
  - *What is the probability that demand is exactly 2?*
  - *What is the probability that demand is no more than 3?*
  - *If you have 2 units to sell, then what is the expected lost sales?*

Q	$f(Q)$	$F(Q)$	$L(Q)$
0	0.286505	0.286505	1.25000
1	0.358131	0.644636	0.53650
2	0.223832	0.868468	0.18114
3	0.093263	0.961731	0.04961
4	0.029145	0.990876	0.01134
5	0.007286	0.998162	0.00221
6	0.001518	0.999680	0.00038
7	0.000271	0.999951	0.00006
8	0.000042	0.999993	0.00001
9	0.000006	0.999999	0.00000
10	0.000001	1.000000	0.00000

16.5 The Order Up-To Model:

Performance measures

What determines the inventory level?

- Short answer:
  - *Inventory level* at the end of a period = *S* minus demand over l +1 periods.

- Explanation via an example with *S* = 6, l = 3, and 2 units on-hand at the start of period 1

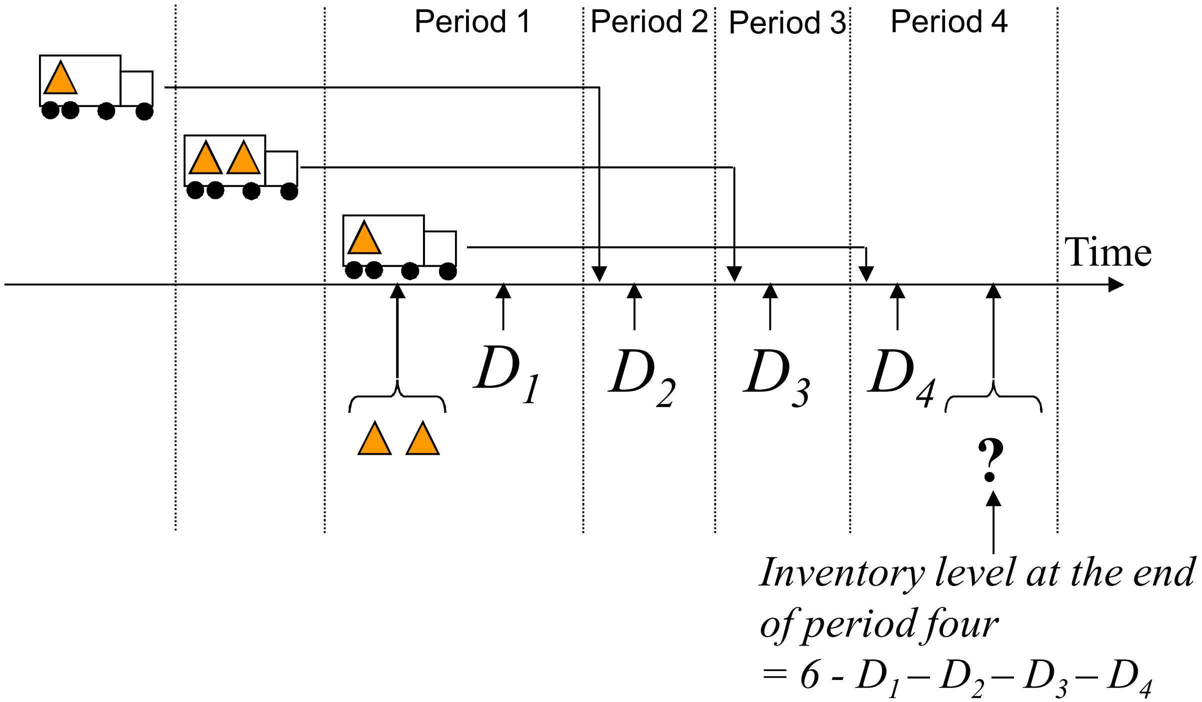

Keep in mind:

At the start of a period the Inventory level + On-order equals S.

All inventory on-order at the start of period 1 arrives before the end of period 4

Nothing ordered in periods 2-4 arrives by the end of period 4

All demand is satisfied so there are no lost sales.

Stockout and in-stock probabilities and on-order inventory

- The **stockout probability** is the probability at least one unit is backordered in a period:

$$\begin{aligned} \textit{Stockout probability} &= \textit{Prob}\{\textit{Demand over } (l+1) \textit{ periods} > S\} \\ &= 1 - \textit{Prob}\{\textit{Demand over } (l+1) \textit{ periods} \leq S\} \end{aligned}$$

- The **in-stock probability** is the probability all demand is filled in a period:

$$\begin{aligned} \textit{In-stock probability} &= 1 - \textit{Stockout probability} \\ &= \textit{Prob}\{\textit{Demand over } (l+1) \textit{ periods} \leq S\} \end{aligned}$$

- ***Expected on-order inventory*** = *Expected demand over one period* x *lead time*
  - This comes from Little's Law. Note that it equals the expected demand over l periods, not l +1 periods.

Demand over l +1 periods at the DC

- The period length is 1 week and the DC's lead time is 3 weeks

- Monthly demand is normally distributed with…
  - Mean = 349
  - Standard deviation = 122.4

- Weekly demand is normally distributed with …
  - Mean = 349 / 4.33 = 80.6
  - Standard deviation = 122.4 / sqrt(4.33) = 58.8

- l +1 demand is normally distributed with …
  - Mean = 80.6 x (3 + 1) = 322.4
  - Standard deviation = 58.8 x sqrt(3 + 1) = 117.6

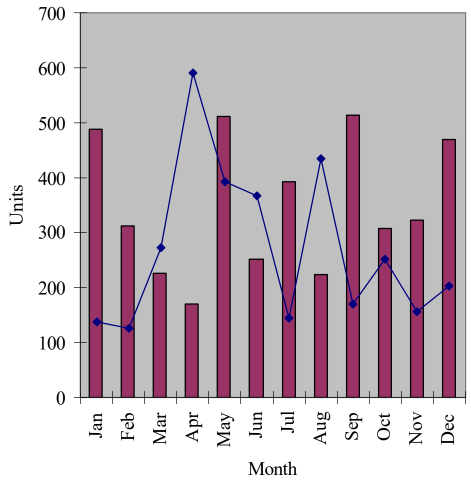

Demand over l +1 periods in Susan's Territory

- The period length is one day, and Susan's replenishment lead time is one day, l =1

- Daily demand is Poisson with …
  - Mean = 0.29

- l +1 demand is Poisson with …
  - Mean = (1+1) x 0.29 = 0.58

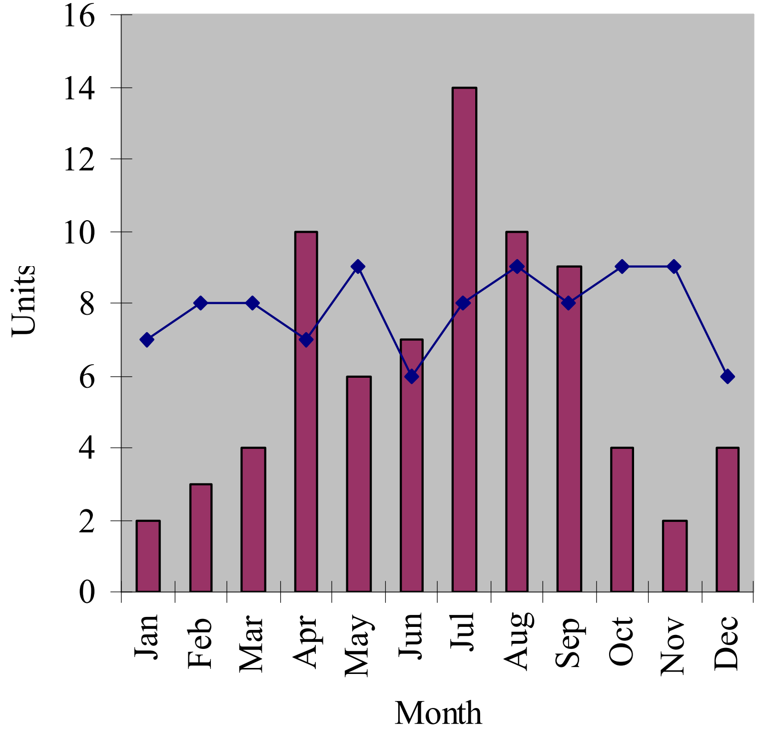

DC's *Expected backorder* assuming *S* = 625

- *Expected backorder* is analogous to the *Expected lost sales* in the Newsvendor model:
  - Suppose *S* = 625 at the DC
  - Normalize the order up-to level:

$$z = \frac{S - \mu}{\sigma} = \frac{625 - 322.4}{117.6} = 2.57$$

  - Lookup *L*(*z*) in the *Standard Normal Loss Function Table*: *L*(2.57)=0.0016
  - Convert expected lost sales, *L*(*z*), for the standard normal into the expected backorder with the actual normal distribution that represents demand over l+1 periods:

$$Expected\ backorder = \sigma \times L(z) = 117.6 \times 0.0016 = 0.19$$

  - Therefore, if *S* = 625, then on average there are 0.19 backorders at the end of any period at the DC.

Other DC performance measures with *S* = 625

- *Expected on-hand inventory* is analogous to *Expected left-over inventory* in the Newsvendor model:

$$z = \frac{S-\mu}{\sigma} = \frac{625-322.4}{117.6} = 2.57$$

- Lookup *I*(z) in the *Standard Normal Inventory Function Table*: *I*(2.57)=2.5716

$$\text{Expected on-hand inventory} = \sigma\, I(z) = 117.6 \text{ X } 2.5716 = 302.4$$

- So on average there are 302.4 units on-hand at the end of a period.

- Due to Little's Law, expected on-order inventory does not depend on demand variability:

$$\begin{aligned} \textit{Expected on-order inventory} &= \textit{Expected demand in one period} \times \textit{Lead time} \\ &= 80.6 \times 3 = 241.8. \end{aligned}$$

- So there are 241.8 units on-order at any given time.

Performance measures in Susan's territory

	Mean demand = 0.29			Mean demand = 0.58	
S	$F(S)$	$L(S)$	S	$F(S)$	$L(S)$
0	0.74826	0.29000	0	0.55990	0.58000
1	0.96526	0.03826	1	0.88464	0.13990
2	0.99672	0.00352	2	0.97881	0.02454
3	0.99977	0.00025	3	0.99702	0.00335
4	0.99999	0.00001	4	0.99966	0.00037
5	1.00000	0.00000	5	0.99997	0.00004

$F(S) = Prob\ \{Demand\ is\ less\ than\ or\ equal\ to\ S\}$

$L(S)$ = loss function = *expected backorder* = expected amount demand exceeds S

- Suppose Susan operates with S = 3:
  - *What is the expected backorder?*
  - *What is the in-stock probability?*
  - *What is expected on-order inventory?*
  - *What is expected on-hand inventory? 2.42335*

(Appendix B. Poisson Inventory Function Table)

16.6 The Order Up-To Model:

Choosing an order up-to level, *S*, to meet a service target

Choose *S* to hit a target in-stock with normally distributed demand

- Suppose the target in-stock probability at the DC is 99.9%:
  - From the Standard Normal Distribution Function Table (p461), $\Phi(3.08)=0.9990$
  - So we choose $z = 3.08$
  - To convert z into an order up-to level:

$$S = \mu + z \times \sigma = 322.4 + 3.08 \times 117.6$$

$$= 685$$

  - Note that μ and σ are the parameters of the normal distribution that describes demand over $l + 1$ periods.

Choose S to hit a target in-stock with Poisson demand

- *What S should Susan choose to yield an in-stock probability of at least 99%?*
- *What if the target is 99.96%?*
- *What if the target is 99.99%?*

	Mean demand = 0.29			Mean demand = 0.58	
S	$F(S)$	$L(S)$	S	$F(S)$	$L(S)$
0	0.74826	0.29000	0	0.55990	0.58000
1	0.96526	0.03826	1	0.88464	0.13990
2	0.99672	0.00352	2	0.97881	0.02454
3	0.99977	0.00025	3	0.99702	0.00335
4	0.99999	0.00001	4	0.99966	0.00037
5	1.00000	0.00000	5	0.99997	0.00004

$F(S) = Prob\ \{Demand\ is\ less\ than\ or\ equal\ to\ S\}$

$L(S)$ = loss function = *expected backorder* = expected amount demand exceeds S

16.7 The Order Up-To Model:

Choosing an appropriate service level

What does a 99.9% in-stock target mean?

- Let *x* be the in-stock probability target for a period.
- Let *N* be the number of periods in an interval of interest (e.g., one week, 365 days, etc).
- The expected number of periods in which an out-of-stock occurs = $N \times (1- x)$
- The expected number of periods before the first out-of-stock occurs = $1 / (1- x)$

	Expected time to 1st day out of stock*		Expected number of days out of stock per year*		
In-stock	Days	Years	per year	per 3 years	per 5 years
99.00%	100	0.27	3.7	11.0	18.3
99.50%	200	0.55	1.8	5.5	9.1
99.90%	1000	2.74	0.4	1.1	1.8
99.95%	2000	5.48	0.2	0.5	0.9
99.99%	10000	27.40	0.0	0.1	0.2

* one day = one period

Justifying a service level via cost minimization

- Let h equal the holding cost per unit per period
  - e.g. if p is the retail price, the gross margin is 75%, the annual holding cost rate is 35% and there are 260 days per year,

    then **$h = p$ x (1-0.75) x 0.35 / 260** = 0.000337 x p
- Let b equal the penalty per unit backordered
  - e.g., let the penalty equal the 75% gross margin, then **b = 0.75 x p**
- "Too much-too little" challenge:
  - If S is too high, then there are holding costs, $C_o = h$
  - If S is too low, then there are backorders, $C_u = b$
- Cost minimizing order up-to level satisfies

$$Prob\left\{Demand\ over\ (l+1)\ periods \le S\right\} = \frac{C_u}{C_o + C_u} = \frac{b}{h+b} = \frac{(0.75 \times p)}{(0.000337 \times p) + (0.75 \times p)}$$

$$= 0.9996$$

- Optimal in-stock probability is 99.96% because

$$In\text{-}stock\ probability = Prob\left\{Demand\ over\ (l+1)\ periods \le S\right\} = Critical\ ratio$$

The optimal in-stock probability is usually quite high

- Suppose the annual holding cost rate is 35%, the backorder penalty cost equals the gross margin and inventory is reviewed daily.

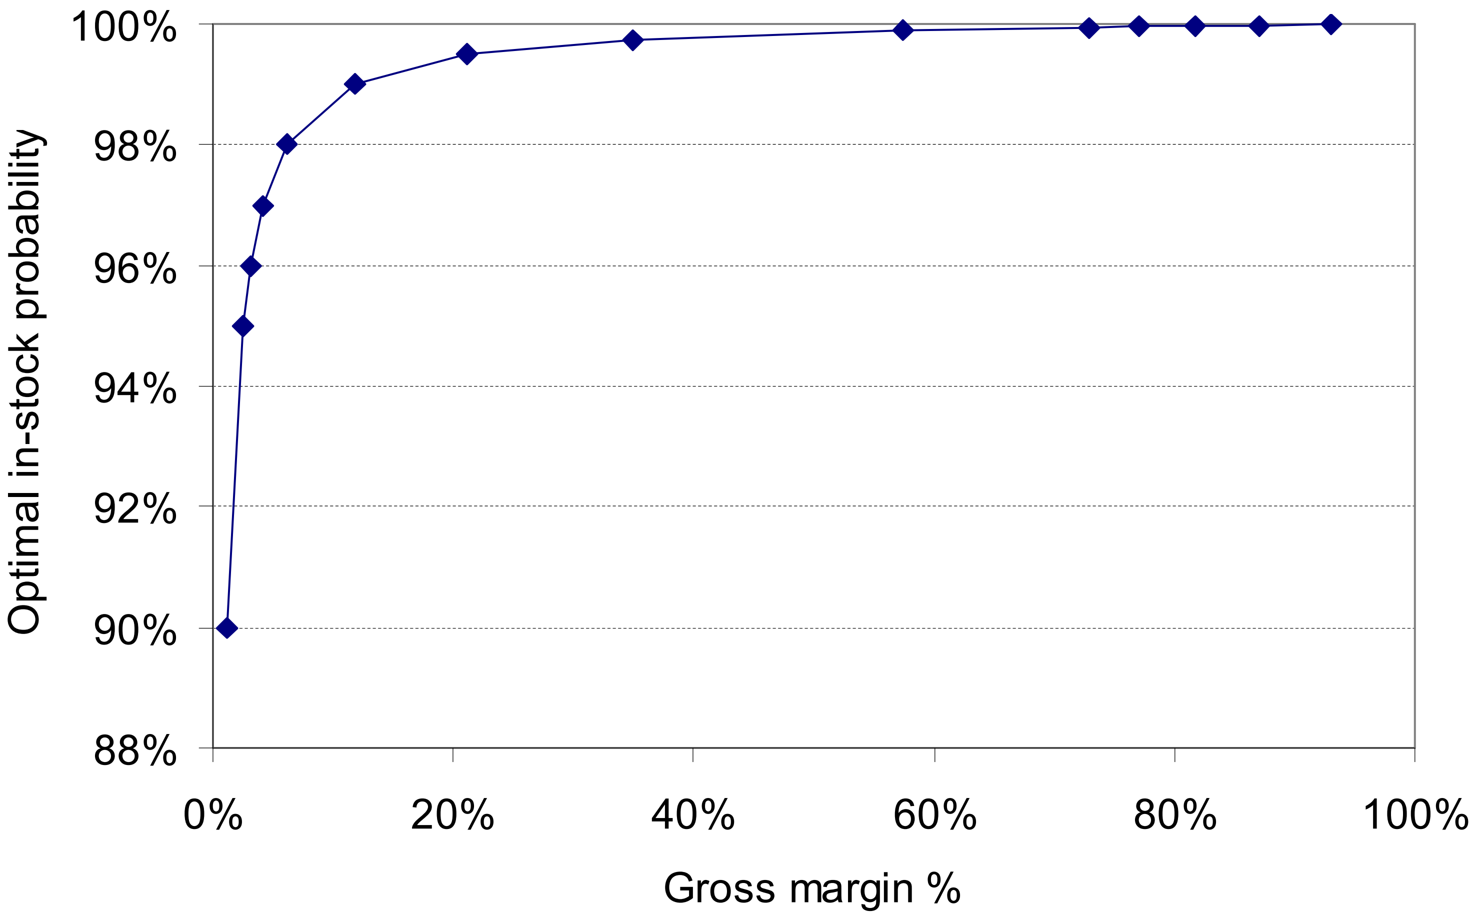

18.9 The Order Up-To Model:

Managerial insights

What determines the average inventory level?

- Using the performance metric equations discussed, we can derive an equation for the days of supply of on-hand inventory that has three components:

$$Days\ of\ supply = \underbrace{\sqrt{l+1}}_{\text{Lead time}} \times \underbrace{\left(\frac{\sigma}{\mu}\right)}_{\text{Variability of demand}} \times \underbrace{\left(\phi(z) + z\Phi(z)\right)}_{\text{In-stock/service level}}$$

  - l = lead time
  - σ = standard deviation of one period demand
  - μ = mean of one period demand
  - z = z-statistic
  - $\phi(z)$, $\Phi(z)$ = normal density and distribution functions

Shorten lead times and you will reduce inventory

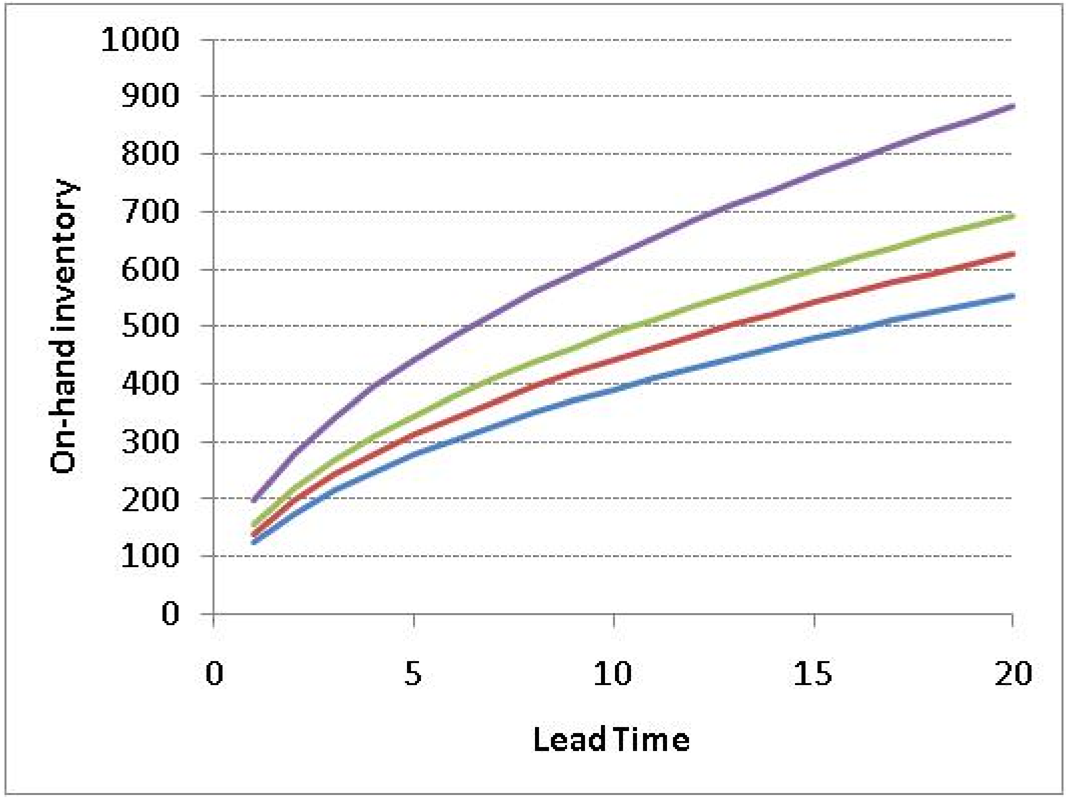

- **Reducing the lead time** reduces expected inventory, especially as the target in-stock increases

The impact of lead time on expected inventory for four in-stock targets, 99.95%, 99.5%, 99.0% and 98%, top curve to bottom curve respectively. Demand in one period is normally distributed with mean 100 and standard deviation 60.

Do not forget about pipeline inventory

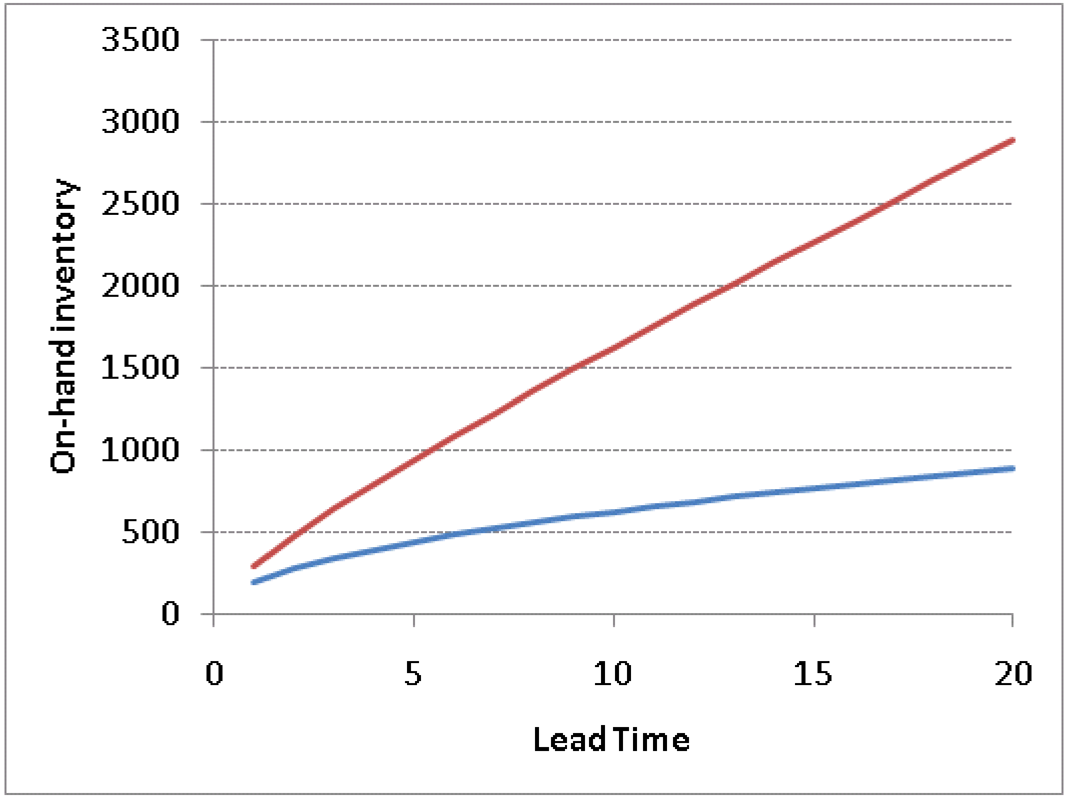

Expected on-hand inventory (blue/lower) and total inventory (red/upper), which is expected on-hand inventory plus pipeline inventory, with a 99.9% in-stock requirement and demand in one period is Normally distributed with mean 100 and standard deviation 60

- Reducing the lead time reduces expected inventory and pipeline inventory
- **The impact on pipeline inventory** can be even more dramatic that the impact on on-hand inventory

Increase service and you need to add inventory

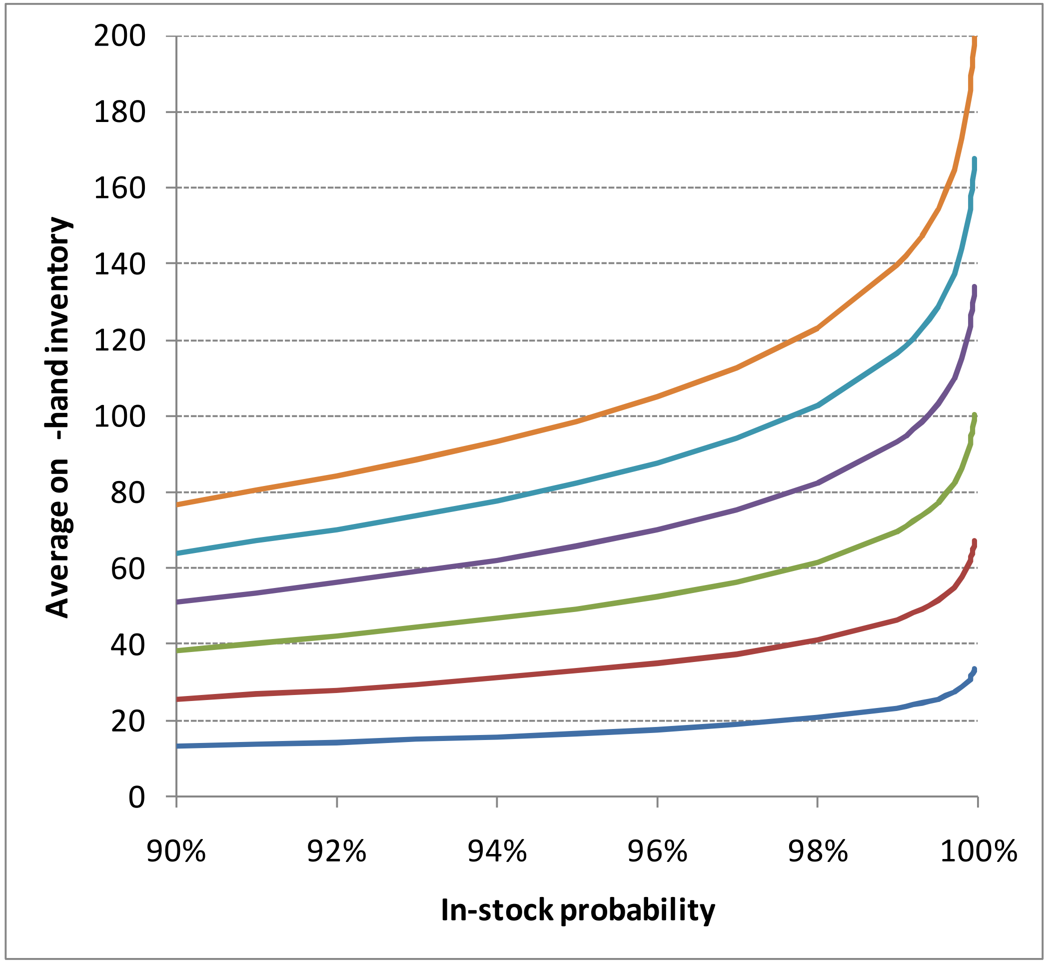

- **The higher the targeted service**, the more inventory needed.
- Inventory increases at an increasing rate as the target service approaches 100%.

The tradeoff between inventory and in-stock probability with Normally distributed demand and a mean of 100 over (l+1) periods. The curves differ in the standard deviation of demand over (l+1) periods: 60,50,40,30,20,10 from top to bottom.

More variability means more inventory

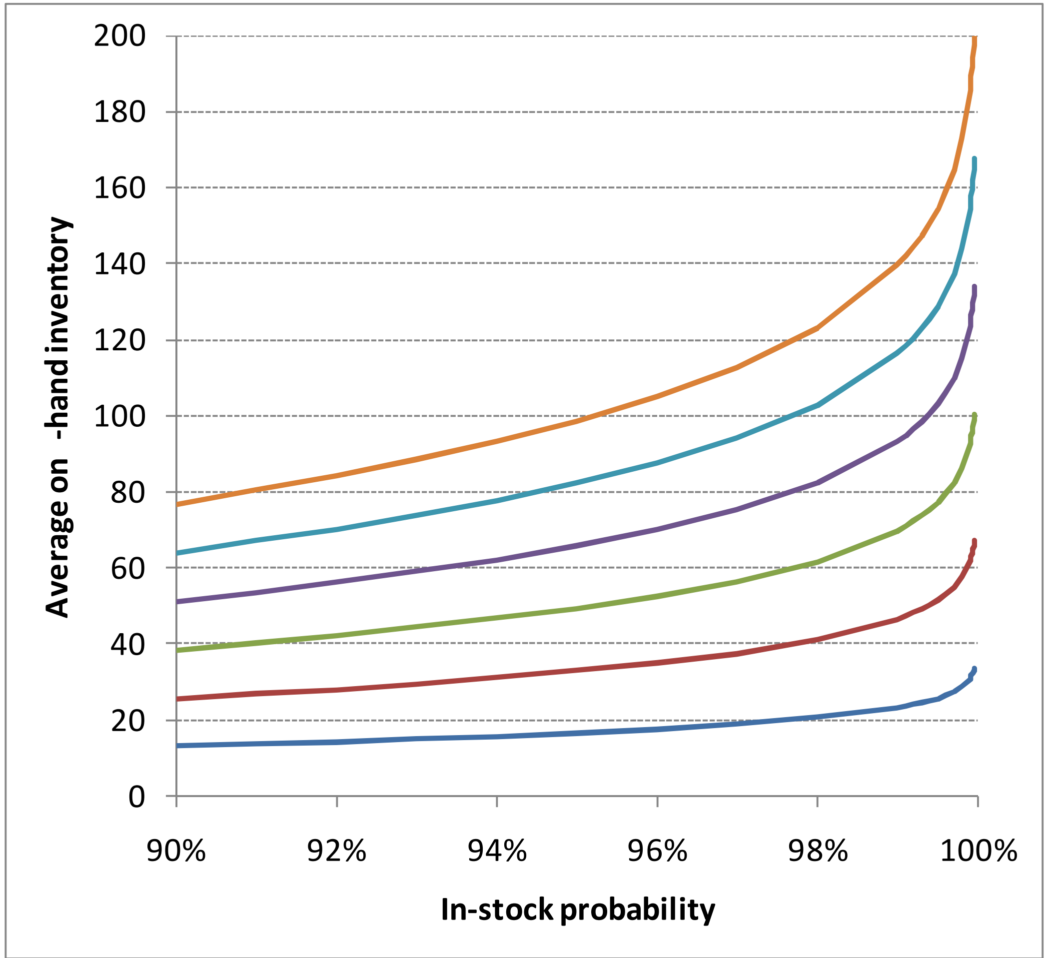

The tradeoff between inventory and in-stock probability with Normally distributed demand and a mean of 100 over (l+1) periods. The curves differ in the standard deviation of demand over (l+1) periods: 60,50,40,30,20,10 from top to bottom.

- **As the coefficient of variation of demand increases…**
  1. The more inventory is needed.
  2. More inventory needs to be added with higher service levels.

Order up-to model summary

- The order up-to model is appropriate for products with random demand but many replenishment opportunities.

- Expected inventory and service are controlled via the order up-to level:
  - The higher the order up-to level the greater the expected inventory and the better the in-stock probability.

- The key factors that determine the amount of inventory needed are…
  - The length of the replenishment lead time.
  - The desired in-stock probability.
  - Demand uncertainty.

- When inventory obsolescence is not an issue, the optimal in-stock probability is generally quite high.