



CTL.SC1x - Supply Chain Fundamentals

Key Concepts Document

V1.1

This document contains the Key Concepts for the SC1x course.
These are meant to complement, not replace, the lesson videos and slides. They are intended to be references for you to use going forward and are based on the assumption that you have learned the concepts and completed the practice problems.

The second draft was created by Dr. Alexis Bateman in the Winter of 2017.

This is a draft of the material, so please post any suggestions, corrections, or recommendations to the Discussion Forum under the topic thread “Key Concept Documents Improvements.
Thanks,

Chris Caplice, Eva Ponce and the SC0x Teaching Community
Winter 2017 v2

V2 Winter 2017 | This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

Table of Contents

Core Supply Chain Concepts **3**

Demand Forecasting **9**

Time Series Analysis *11*

Exponential Smoothing *13*

Exponential Smoothing with Holt-Winter *15*

Special Cases *16*

Inventory Management **22**

Economic Order Quantity (EOQ) *24*

Economic Order Quantity (EOQ) Extensions *27*

Single Period Inventory Models *31*

Single Period Inventory Models-Expected Profitability *34*

Probabilistic Inventory Models *36*

Inventory Models for Multiple Items & Locations *43*

Inventory Models for Class A & C Items *47*

Fundamentals of Freight Transportation **54**

Lead Time Variability & Mode Selection *55*

Core Supply Chain Concepts

Summary

Virtually all supply chains are a combination of push and pull systems. A push system is where execution is performed ahead of an actual order so that the forecasted demand, rather than actual demand, has to be used in planning. A pull system is where execution is performed in response to an order so that the actual demand is known with certainty. The point in the process where a supply chain shifts from being push to pull is sometimes called the push/pull boundary or push/pull point. In manufacturing, the push/pull point is also known as the decoupling point (DP) or customer order decoupling point (CODP). The CODP coincides with an important stock point, where the customer order arrives (switching inventory based on a forecast to actual demand), and also allows to differentiate basic production systems: make-to-stock, assemble-to-order, make-to-order, or engineer-to-order.

Postponement is a common strategy to combine the benefits of push (product ready for demand) and pull (fast customized service) systems. Postponement is where the undifferentiated raw or components are "pushed" through a forecast, and the final finished and customized products are then "pulled".

Segmentation is a method of dividing a supply chain into two or more groupings where the supply chains operate differently and more efficiently and/or effectively. While there are no absolute rules for segmentation, there are some rules of thumb, such as: items should be homogenous within the segment and heterogeneous across segments; there should be critical mass within each segment; and the segments need to be useful and communicable. A segment only makes sense if it does something different (planning, inventory, transportation etc.) from the other segments. The most common segmentation is for products using an ABC classification.

In an ABC segmentation, the products driving the most revenue (or profit or sales) are Class A items (the important few). Products driving very little revenue (or profit or sales) are Class C items (the trivial many), and the products in the middle are Class B. A common breakdown is the top 20% of items (Class A) generate 80% of the revenue, Class B is 30% of the products generating 15% of the revenue, and the Class C items generate less than 5% of the revenue while constituting 50% of the items.

Supply chains operate in uncertainty. Demand is never known exactly, for example. In order to handle and be able to analyze systems with uncertainty, we need to capture the distribution of the variable in question. When we are describing a random situation, say, the expected demand for pizzas on a Thursday night, it is helpful to describe the potential outcomes in terms of the central tendency (mean or median) as well as the dispersion (standard deviation, range). We will often characterize the distribution of potential outcomes as following a well-known function such as Normal and Poisson.

Key Concepts

Pull vs. Push Process

- Push—work performed in anticipation of an order (forecasted demand)
- Pull—execution performed in response to an order (demand known with certainty)
- Hybrid or Mixed—push raw products, pull finished product (postponement or delayed differentiation)
- Push/pull boundary point — point in the process where a supply chain shifts from being push to pull
- In manufacturing, also known as “decoupling point” (DP) or “customer order decoupling point” (CODP) — the point in the material flow where the product is linked to a specific customer
- Mass customization / Postponement — to delay the final assembly, customization, or differentiation of a product until as late as possible

Segmentation

- Differentiate products in order to match the right supply chain to the right product
- Products typically segmented on
    - Physical characteristics (value, size, density, etc.)
    - Demand characteristics (sales volume, volatility, sales duration, etc.)
    - Supply characteristics (availability, location, reliability, etc.)
- Rules of thumb for number of segments
    - Homogeneous—products within a segment should be similar
    - Heterogeneous—products across segments should be very different
    - Critical Mass—segment should be big enough to be worthwhile
    - Pragmatic—segmentation should be useful and communicable
- Demand follows a power law distribution, meaning a large volume of sales is concentrated in few products

Power Law

The distribution of percent sales volume to percent of SKUs (Stock Keeping Units) tends to follow a Power Law distribution ($y=ax^k$) where y is percent of demand (units or sales or profit), x is percent of SKUs, and a and k are parameters. The value for k should obviously be less than 1 since if k=1 the relationship is linear. In addition to segmenting according to products, many firms segment by customer, geographic region, or supplier. Segmentation is typically done using revenue as the key driver, but many firms also include variability of demand, profitability, and other factors, to include:

- Revenue = average sales*unit sales price;
- Profit = average sales*margin;

- Margin = unit sales price–unit cost.

Handling Uncertainty

Uncertainty of an outcome (demand, transit time, manufacturing yield, etc.) is modeled through a probability distribution. We discussed two in the lesson: Poisson and Normal.

Normal Distribution ~N(μ,σ)

This is the Bell Shaped distribution that is widely used by both practitioners and academics. While not perfect, it is a good place to start for most random variables that you will encounter in practice such as transit time and demand. The distribution is both continuous (it can take any number, not just integers or positive numbers) and is symmetric around its mean or average. Being symmetric additionally means the mean is also the median and the mode. The common notation that we will use to indicate that some value follows a Normal Distribution is ~N(μ, σ) where mu, μ, is the mean and sigma, σ, is the standard deviation. Some books use the notation ~N(μ, σ^2) showing the variance, σ^2, instead of the standard deviation. Just be sure which notation is being followed when you consult other texts.

The Normal Distribution is formally defined as:

$$f_x(x_0) = \frac{e^{\frac{-(x_0-\mu)^2}{2\sigma_x^2}}}{\sigma_x\sqrt{2\pi}}$$

We will also make use of the Unit Normal or Standard Normal Distribution. This is ~N(0,1) where the mean is zero and the standard deviation is 1 (as is the variance, obviously). The chart below shows the standard or unit normal distribution. We will be making use of the transformation from any Normal Distribution to the Unit Normal (See Figure 1).

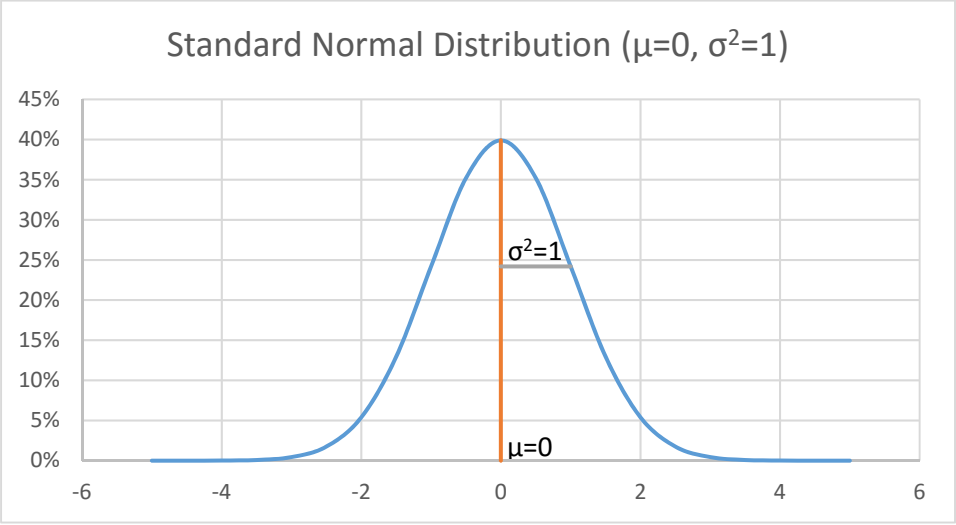


Figure 1. Standard Normal Distribution

We will make extensive use of spreadsheets (whether Excel or LibreOffice) to calculate probabilities under the Normal Distribution. The following functions are helpful:

- NORMDIST(x, μ, σ, true) = the probability that a random variable is less than or equal to x under the Normal Distribution ~N(μ, σ). So, that NORMDIST(25, 20, 3, 1) = 0.952 which means that there is a 95.2% probability that a number from this distribution will be less than 25.
- NORMINV(probability, μ, σ) = the value of x where the probability that a random variable is less than or equal to it is the specified probability. So, NORMINV(0.952, 20, 3) = 25.

To use the Unit Normal Distribution ~N(0,1) we need to transform the given distribution by calculating a k value where k=(x- μ)/σ. This is sometimes called a z value in statistics courses, but in almost all supply chain and inventory contexts it is referred to as a k value. So, in our example, k = (25 – 20)/3 = 1.67. Why do we use the Unit Normal? Well, the k value is a helpful and convenient piece of information. The k is the number of standard deviations the value x is above (or below if it is negative) the mean. We will be looking at a number of specific values for k that are widely used as thresholds in practice, specifically:

- Probability (x ≤ 0.90) where k = 1.28
- Probability (x ≤ 0.95) where k = 1.62
- Probability (x ≤ 0.99) where k = 2.33

Because the Normal Distribution is symmetric, there are also some common confidence intervals:

- μ ± σ 68.3% — meaning that 68.3% of the values fall within 1 standard deviation of the mean,
- μ ± 2σ 95.5% — 95.5% of the values fall within 2 standard deviations of the mean, and
- μ ± 3σ 99.7% — 99.7% of the values fall within 3 standard deviations of the mean.

In a spreadsheets you can use the functions:

- **NORMSDIST(k)** = the probability that a random variable is less than k units above (or below) mean. For example, NORMSDIST(2.0) = 0.977 meaning the 97.7% of the distribution is less than 2 standard deviations above the mean.
- **NORMSINV(probability)** = the value corresponding to the given probability. So that NORMSINV(0.977) = 2.0. If I then wanted to find the value that would cover 97.7% of a specific distribution, say where ~N(279, 46) I would just transform it. Since k=(x- μ)/σ for the transformation, I can simply solve for x and get: x = μ + kσ = 279 + (2.0)(46) = 371. This means that the random variable ~N(279, 46) will be equal or less than 371 for 97.7% of the time.

Poisson distribution ~Poisson (λ)

We will also use the Poisson (pronounced pwa-SOHN) distribution for modeling things like demand, stock outs, and other less frequent events. The Poisson, unlike the Normal, is discrete (it can only be integers ≥ 0), always positive, and non-symmetric. It is skewed right – that is, it

has a long right tail. It is very commonly used for low value distributions or slow moving items. While the Normal Distribution has two parameters (mu and sigma), the Poisson only has one, lambda, λ.

Formally, the Poisson Distribution is defined as shown below:

$$p[x_0] = \text{Prob}\left[x = x_0\right] = \frac{e^{-\lambda}\lambda^{x_0}}{x_0!} \quad for\ x_0 = 0,1,2,\ldots$$

$$F[x_0] = \text{Prob}\left[x \le x_0\right] = \sum_{x=0}^{x_0} \frac{e^{-\lambda}\lambda^{x}}{x!}$$

The chart below (Figure 2) shows the Poisson Distribution for λ=3. The Poisson parameter λ is both the mean and the variance for the distribution! Note that λ does not have to be an integer.

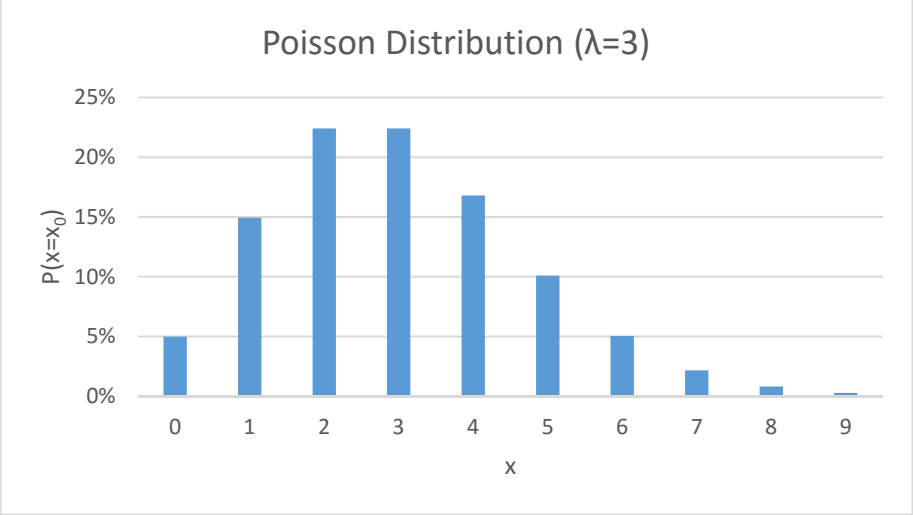


Figure 2. Poisson Distribution

In spreadsheets, the following functions are helpful:

- **POISSON(x_0, λ, false)** => P(x = x_0) = the probability that a random variable is equal to x_0 under the Poisson Distribution ~P(λ). So, that POISSON(2, 1.56, 0) = 0.256 which means that there is a 25.6% probability that a number from this distribution will be equal to 2.
- **POISSON(x_0, λ, true)** => P(x ≤ x_0) = the probability that a random variable is less than or equal to x_0 under the Poisson Distribution ~P(λ). So, that POISSON(2, 1.56, 1) = 0.793 which means that there is a 79.3% probability that a number from this distribution will be less than or equal to 2. This is simply just the cumulative distribution function.

Uniform distribution ~U (a,b)

We will sometimes use the Uniform distribution, which has two parameters: a minimum value *a* and a maximum value *b*. Each point within this range is equally likely to occur. To find the cumulative probability for some value C, the probability that x≤c = (c-a) / (b-a), that is, the area from *a* to *c* minus the total area from a to b. The expected value or the mean is simply (a+b)/2 while the standard deviation is = $(b-a)/\sqrt{12}$.

Learning Objectives

- Identify and understand differences between push and pull systems.
- Understand why and how to segment supply chains by products, customers, etc.
- Ability to model uncertainty in supply chains, primarily, but not exclusively, in demand uncertainty.

References

Push/Pull Processes: Chopra & Meindl Chpt 1; Nahmias Chpt 7;
Segmentation: Nahmias Chpt 5; Silver, Pyke, & Peterson Chpt 3; Ballou Chpt 3
Probability Distributions: Chopra & Meindl Chpt 12; Nahmias Chpt 5; Silver, Pyke, & Peterson App B

Fisher, M. (1997) "What Is the Right Supply Chain for Your Product?," Harvard Business Review.

Olavson, T., Lee, H. & DeNyse, G. (2010) "A Portfolio Approach to Supply Chain Design," Supply Chain Management Review.

Demand Forecasting

Summary

Forecasting is one of three components of an organization's Demand Planning, Forecasting, and Management process. Demand Planning answers the question "What should we do to shape and create demand for our product?" and concerns things like promotions, pricing, packaging, etc. Demand Forecasting then answers "What should we expect demand for our product to be given the demand plan in place?" The final component, Demand Management, answers the question, "How do we prepare for and act on demand when it materializes?" This concerns things like Sales & Operations Planning (S&OP) and balancing supply and demand.

Within the Demand Forecasting component, you can think of three levels, each with its own time horizon and purpose. Strategic forecasts (years) are used for capacity planning, investment strategies, etc. Tactical forecasts (weeks to months to quarters) are used for sales plans, short-term budgets, inventory planning, labor planning, etc. Finally, operations forecasts (hours to days) are used for production, transportation, and inventory replenishment decisions. The time frame of the action dictates the time horizon of the forecast.

Forecasting methods can be divided into being subjective (most often used by marketing and sales) or objective (most often used by production and inventory planners). Subjective methods can be further divided into being either Judgmental (someone somewhere knows the truth), such as sales force surveys, Delphi sessions, or expert opinions, or Experimental (sampling local and then extrapolating), such as customer surveys, focus groups, or test marketing. Objective methods are either Causal (there is an underlying relationship or reason) such as leading indicators, etc. or Time Series (there are patterns in the demand) such as exponential smoothing, moving average, etc. All methods have their place and their role. We will spend a lot of time on the objective methods but will also discuss the subjective ones as well.

Regardless of the forecasting method used, you will want to measure the quality of the forecast. The two major dimensions of quality are bias (a persistent tendency to over- or under-predict) and accuracy (closeness to the actual observations). No single metric does a good job capturing both dimensions, so it is worth having multiple.

Key Concepts

Forecasting is both an art and a science. There are many "truisms" concerning forecasting including:

Forecasting Truisms

1. *Forecasts are always wrong* – Yes, point forecasts will never be completely perfect. The solution is to not rely totally on point forecasts. Incorporate ranges into your forecasts. Also you should try to capture and track the forecast errors so that you can sense and measure any drift or changes.
2. *Aggregated forecasts are more accurate than dis-aggregated forecasts* – The idea is that combining different items leads to a pooling effect that will in turn lessen the variability. The peaks balance out the valleys. The coefficient of variation (CV) is commonly used to measure variability and is defined as the standard deviation over the mean ($CV = \sigma/\mu$). Forecasts are generally aggregated by SKU (a family of products versus an individual one), time (demand over a month versus over a single day), or location (demand for a region versus a single store).
3. *Shorter horizon forecasts are more accurate than longer horizon forecasts* – Essentially this means that forecasting tomorrow's temperature (or demand) is easier and probably more accurate than forecasting for a year from tomorrow. This is not the same as aggregating. It is all about the time between making the forecast and the event happening. Shorter is always better. This is where postponement and modularization helps. If we can somehow shorten the forecasting time for an end item, we will generally be more accurate.

Forecasting Metrics

There is a cost trade-off between cost of errors in forecasting and cost of quality forecasts that must be balanced. Forecast metric systems should capture bias and accuracy.

Notation

A_t: Actual value for observation t
F_t: Forecasted value for observation t
e_t: Error for observation t, $e_t = A_t - F_t$
n: number of observations
μ: mean
σ: standard deviation
CV: Coefficient of Variation – a measure of volatility – $CV = \frac{\sigma}{\mu}$

Formulas:

Mean Deviation: $MD = \frac{\sum_{t=1}^{n} e_t}{n}$

Mean Absolute Deviation: $MAD = \frac{\sum_{t=1}^{n} |e_t|}{n}$

Mean Squared Error: $MSE = \frac{\sum_{t=1}^{n} e_t^2}{n}$

Root Mean Squared Error: $RMSE = \sqrt{\frac{\sum_{t=1}^{n} e_t^2}{n}}$

Mean Percent Error: $MPE = \frac{\sum_{t=1}^{n} \frac{e_t}{A_t}}{n}$

Mean Absolute Percent Error: $MAPE = \frac{\sum_{t=1}^{n} \frac{|e_t|}{A_t}}{n}$

Statistical Aggregation:

$$\sigma_{agg}^2 = \sigma_1^2 + \sigma_2^2 + \sigma_3^2 + \cdots + \sigma_n^2$$

$$\sigma_{agg} = \sqrt{\sigma_1^2 + \sigma_2^2 + \sigma_3^2 + \cdots + \sigma_n^2}$$

$$\mu_{agg} = \mu_1 + \mu_2 + \mu_3 + \cdots + \mu_n$$

Statistical Aggregation of n Distributions of Equal Mean and Variance:

$$\sigma_{agg} = \sqrt{\sigma_1^2 + \sigma_2^2 + \sigma_3^2 + \cdots + \sigma_n^2} = \sigma_{ind}\sqrt{n}$$

$$\mu_{agg} = \mu_1 + \mu_2 + \mu_3 + \cdots + \mu_n = n\mu_{ind}$$

$$CV_{agg} = \frac{\sigma\sqrt{n}}{\mu n} = \frac{\sigma}{\mu\sqrt{n}} = \frac{CV_{ind}}{\sqrt{n}}$$

Time Series Analysis

Time Series is an extremely widely used forecasting technique for mid-range forecasts for items that have a long history or record of demand. Time series is essentially pattern matching of data that are distributed over time. For this reason, you tend to need a lot of data to be able to capture the components or patterns. Business cycles are more suited to longer range, strategic forecasting time horizons.

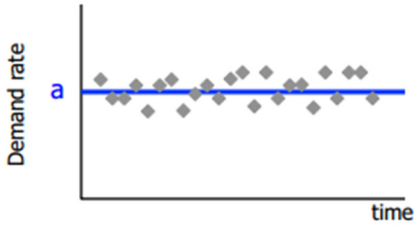
Three important time series models:

- Cumulative – where everything matters and all data are included. This results in a very calm forecast that changes very slowly over time – thus it is more stable than responsive.
- Naïve – where only the latest data point matters. This results in very nervous or volatile forecast that can change quickly and dramatically – thus it is more responsive than stable.
- Moving Average – where we can select how much data to use (the last M periods). This is essentially the generalized form for both the Cumulative (M = ∞) and Naïve (M=1) models.

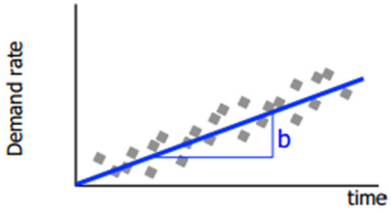
All three of these models are similar in that they assume stationary demand. Any trend in the underlying data will lead to severe lagging. These models also apply equal weighting to each piece of information that is included. Interestingly, while the M-Period Moving Average model requires M data elements for each SKU being forecast, the Naïve and Cumulative models only require 1 data element each.

Components of time series

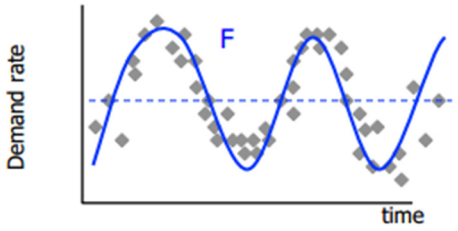
- Level (a)
  - Value where demand hovers (mean)
  - Captures scale of the time series
  - With no other pattern present, it is a constant value



- Trend (b)
  - Rate of growth or decline
  - Persistent movement in one direction
  - Typically linear but can be exponential, quadratic, etc.



- Season Variations (F)
  - Repeated cycle around a known and fixed period
  - Hourly, daily, weekly, monthly, quarterly, etc.
  - Can be caused by natural or man-made forces



- Random Fluctuation (e or ε)
  - Remainder of variability after other components
  - Irregular and unpredictable variations, noise

Notation

x_t: Actual demand in period t
$\hat{x}_{t,t+1}$: Forecast for time t+1 made during time t
a: Level component
b: Linear trend component
F_t: Season index appropriate for period t
e_t: Error for observation t, $e_t = A_t - F_t$
t: Time period (0, 1, 2,…n)
Level Model: $x_t = a + e_t$
Trend Model: $x_t = a + bt + e_t$
Mix Level-Seasonality Model: $x_t = aF_t + e_t$
Mix Level-Trend-Seasonality Model: $x_t = (a + bt)F_t + e_t$

Formulas

Time Series Models (Stationary Demand only):

Cumulative Model: $\hat{x}_{t,t+1} = \frac{\sum_i^t x_i}{t}$

Naïve Model: $\hat{x}_{t,t+1} = x_t$

M-Period Moving Average Forecast Model: $\hat{x}_{t,t+1} = \frac{\sum_{i=t+1-M}^t x_i}{M}$

- If M=t, we have the cumulative model where all data is included
- If M=1, we have the naïve model, where the last data point is used to predict the next data point

Exponential Smoothing

Exponential smoothing, as opposed to Cumulative, Naïve, and Moving Average, treats data differently depending on its age. The idea is that the value of data degrades over time so that newer observations of demand are weighted more heavily than older observations. The weights decrease exponentially as they age. Exponential models simply blend the value of new and old information.

The alpha factor (ranging between 0 and 1) determines the weighting for the newest information versus the older information. The "α" value indicates the value of "new" information versus "old" information:

- As α → 1, the forecast becomes more nervous, volatile, and naïve
- As α → 0, the forecast becomes more calm, staid, and cumulative
- α can range from 0 ≤ α ≤ 1, but in practice, we typically see 0 ≤ α ≤ 0.3

The most basic exponential model, or Simple Exponential model, assumes stationary demand. Holt's Model is a modified version of exponential smoothing that also accounts for trend in

addition to level. A new smoothing parameter, β, is introduced. It operates in the same way as the α.

We can also use exponential smoothing to dampen trend models to account for the fact that trends usually do not remain unchanged indefinitely as well as for creating a more stable estimate of the forecast errors.

Notation

- x_t: Actual demand in period t
- $\hat{x}_{t,t+1}$: Forecast for time t+1 made during time t
- α: Exponential smoothing factor for level ($0 \le \alpha \le 1$)
- β: Exponential smoothing factor for trend ($0 \le \beta \le 1$)
- ϕ: Exponential smoothing factor for dampening ($0 \le \phi \le 1$)
- ω: Mean Square Error trending factor ($0.01 \le \omega \le 0.1$)

Forecasting Models

Simple Exponential Smoothing Model (Level Only) – This model is used for stationary demand. The "new" information is simply the latest observation. The "old" information is the most recent forecast since it encapsulates the older information.

$$\hat{x}_{t,t+1} = \alpha x_t + (1-\alpha)\hat{x}_{t-1,t}$$

Damped Trend Model with Level and Trend – We can use exponential smoothing to dampen a linear trend to better reflect the tapering effect of trends in practice.

$$\hat{x}_{t,t+\tau} = \hat{a}_t + \sum_{i=1}^{\tau} \varphi^i \hat{b}_t$$
$$\hat{a}_t = \alpha x_t + (1-\alpha)\left(\hat{a}_{t-1} + \varphi \hat{b}_{t-1}\right)$$
$$\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\varphi \hat{b}_{t-1}$$

Exponential Smoothing for Level & Trend – also known as Holt's Method, assumes a linear trend. The forecast for time t+τ made at time t is shown below. It is a combination of the latest estimates of the level and trend. For the level, the new information is the latest observation and the old information is the most recent forecast for that period – that is, the last period's estimate of level plus the last period's estimate of trend. For the trend, the new information is the difference between the most recent estimate of the level minus the second most recent estimate of the level. The old information is simply the last period's estimate of the trend.

$$\hat{x}_{t,t+\tau} = \hat{a}_t + \tau \hat{b}_t$$
$$\hat{a}_t = \alpha x_t + (1-\alpha)(\hat{a}_{t-1} + \hat{b}_{t-1})$$
$$\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\hat{b}_{t-1}$$

Mean Square Error Estimate – We can also use exponential smoothing to provide a more robust or stable value for the mean square error of the forecast.

$$MSE_t = \omega(x_t - \hat{x}_{t-1,t})^2 + (1-\omega)MSE_{t-1}$$

Exponential Smoothing with Holt-Winter

Seasonality

- For multiplicative seasonality, think of the F_i as "percent of average demand" for a period i
- The sum of the F_i for all periods within a season must equal P
- Seasonality factors must be kept current or they will drift dramatically. This requires a lot more bookkeeping, which is tricky to maintain in a spreadsheet, but it is important to understand

Forecasting Model Parameter Initialization Methods

- While there is no single best method, there are many good ones
- Simple Exponential Smoothing
  - Estimate level parameter $\hat{a}_0$ by averaging demand for first several periods
- Holt Model (trend and level)—must estimate both $\hat{a}_0$ and $\hat{b}_0$
  - Find a best fit linear equation from initial data
  - Use least squares regression of demand for several periods
    - Dependent variable = demand in each time period = x_t
    - Independent variable = slope = β_1
    - Regression equation: $x_t = \beta_0 + \beta_1 t$
- Seasonality Models
  - Much more complicated, you need at least two season of data but preferably four or more
  - First determine the level for each common season period and then the demand for all periods
  - Set initial seasonality indices to ratio of each season to all periods

Notation

x_t: Actual demand in period t
$\hat{x}_{t,t+1}$: Forecast for time t+1 made during time t
α: Exponential smoothing factor ($0 \le \alpha \le 1$)
β: Exponential smoothing trend factor ($0 \le \beta \le 1$)
γ: Seasonality smoothing factor ($0 \le \gamma \le 1$)
F_t: Multiplicative seasonal index appropriate for period t
P: Number of time periods within the seasonality (note: $\sum_{i=1}^{P} \hat{F}_i = P$)

Forecasting Models

Double Exponential Smoothing (Seasonality and Level) – This is a multiplicative model in that the seasonality for each period is the product of the level and that period's seasonality factor. The new information for the estimate of the level is the "de-seasoned" value of the latest observation; that is, you are trying to remove the seasonality factor. The old information is simply the previous most recent estimate for level. For the seasonality estimate, the new information is the "de-leveled" value of the latest observation; that is, you try to remove the level factor to understand any new seasonality. The old information is simply the previous most recent estimate for that period's seasonality.

$$\hat{x}_{t,t+\tau} = \hat{a}_t \hat{F}_{t+\tau-P}$$

$$\hat{a}_t = \alpha\left(\frac{x_t}{\hat{F}_{t-P}}\right) + (1-\alpha)(\hat{a}_{t-1})$$

$$\hat{F}_t = \gamma\left(\frac{x_t}{\hat{a}_t}\right) + (1-\gamma)\hat{F}_{t-P}$$

Holt-Winter Exponential Smoothing Model (Level, Trend, and Seasonality) – This model assumes a linear trend with a multiplicative seasonality effect over both level and trend. For the level estimate, the new information is again the "de-seasoned" value of the latest observation, while the old information is the old estimate of the level and trend. The estimate for the trend is the same as for the Holt model. The Seasonality estimate is the same as the Double Exponential smoothing model.

$$\hat{x}_{t,t+\tau} = (\hat{a}_t + \tau\hat{b}_t)\hat{F}_{t+\tau-P}$$

$$\hat{a}_t = \alpha\left(\frac{x_t}{\hat{F}_{t-P}}\right) + (1-\alpha)(\hat{a}_{t-1} + \hat{b}_{t-1})$$

$$\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\hat{b}_{t-1}$$

$$\hat{F}_t = \gamma\left(\frac{x_t}{\hat{a}_t}\right) + (1-\gamma)\hat{F}_{t-P}$$

Normalizing Seasonality Indices – This should be done after each forecast to ensure the seasonality does not get out of synch. If the indices are not updated, they will drift dramatically. Most software packages will take care of this – but it is worth checking.

$$\hat{F}_t^{NEW} = \hat{F}_t^{OLD}\left(\frac{P}{\sum_{i=t-P}^{t}\hat{F}_t^{OLD}}\right)$$

Special Cases

There are different types of new products and the forecasting techniques differ according to their type. The fundamental idea is that if you do not have any history to rely on, you can look

for history of similar products and build one. Bass Diffusion models can be used to estimate market demand for new-to-world products.

When the demand is very sparse, such as for spare parts, we cannot use traditional methods since the estimates tend to fluctuate dramatically. Croston's method can smooth out the estimate for the demand.

New Product Types

- Not all new products are the same. We can roughly classify them into the following six categories (listed from easiest to forecast to hardest):
    - Cost Reductions: Reduced price version of the product for the existing market
    - Product Repositioning: Taking existing products/services to new markets or applying them to a new purpose (aspirin from pain killer to reducing effects of a heart attack)
    - Line Extensions: Incremental innovations added to complement existing product lines (Vanilla Coke, Coke Zero) or Product Improvements: New, improved versions of existing offering targeted to the current market—replaces existing products (next generation of product)
    - New-to-Company: New market/category for the company but not to the market (Apple iPhone or iPod)
    - New-to-World: First of their kind, creates new market, radically different (Sony Walkman, Post-it notes, etc.)
- While they are a pain to forecast and to launch, firms introduce new products all the time – this is because they are the primary way to increase revenue and profits (See Table 1)

Type of New Product	Percent of Introductions	Forecast Accuracy (1-MAPE)	Launch Cycle Length	Success Rate
New-to-World	8-10%	40%	104 weeks	38-65%
New-to-Company	17-20%	47%		
Line Extensions	21-26%	54-62%	62/29* weeks	55-77%
Product Improvements	26-36%	65%		
Product Re-Positioning	5-7%	54-65%	N/A	66-79%
Cost Reductions	10-11%	72%		

*Major revisions/incremental improvements about evenly split

Table 1. New product introductions. Source: Adapted from Cooper, Robert (2001) Winning at New Products, Kahn, Kenneth (2006) New Product forecasting, and PDMA (2004) New Product Development Report.)

New Product Development Process

All firms use some version of the process shown below to introduce new products. This is sometimes called the stage-gate or funnel process. The concept is that lots of ideas come in on the left and very few final products come out on the right. Each stage or hurdle in the process winnows out the winners from the losers and is used to focus attention on the right products. The scope and scale of forecasting changes along the process as noted in Figure 3.

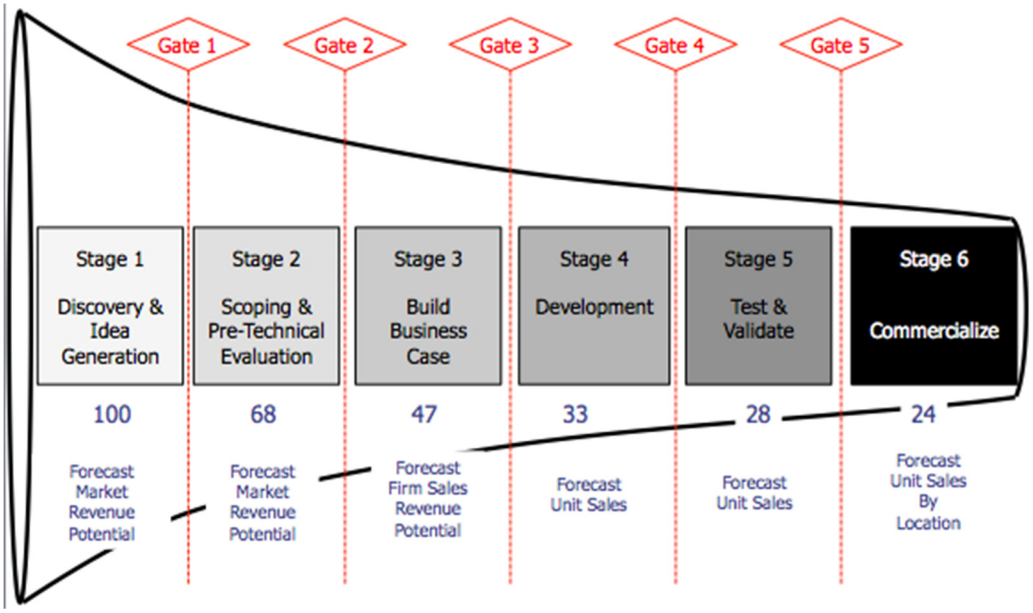


Figure 3. New product development process

Forecasting Models Discussed

New Product – "Looks-Like" or Analogous Forecasting

- Perform by looking at comparable product launches and create a week-by-week or month-by-month sales record.
- Then use the percent of total sales in each time increment as a trajectory guide.
- Each launch should be characterized by product type, season of introduction, price, target market demographics, and physical characteristics.

New-to-World Products - Bass Diffusion Model

- The idea is that there are two complementary effects at work, innovation and imitation:
  - Innovation effect: Innovators are early adopters that are drawn to technology regardless of who else is using it. Their demand peaks early in the product's lifecycle.
  - Imitation effect: Imitators hear about the product by word of mouth and are influenced by peers. Their demand lags behind the innovators.
- The values for p (innovation) and q (imitation) can be estimated using various techniques.

Intermittent or Sparse Demand – Croston's Method

- Used for products that are infrequently ordered in large quantities, irregularly ordered, or ordered in different sizes.
- Croston's Method separates out the demand and model—unbiased and has lower variance than simple smoothing.
- Cautions: infrequent ordering (and updating of model) induces a lag to responding to magnitude changes.

Notation

P: Innovation effect in Bass Diffusion Model; p ~ 0.03 and often <0.01
q: Imitation effect in Bass Diffusion Model; q ~ 0.38 and often 0.3 ≤ q ≤ 0.5
m: Total number of customers who will adopt
n(t): Number of customer adopting at time t
N(t-1): Cumulative number of sales by time t-1
t*: Peak sales time in Bass Diffusion Model
x_t: Demand in period t
y_t: 1 if transaction occurs in period t, =0 otherwise
z_t: Size (magnitude) of transaction in time t
n_t: Number of periods since last transaction
α: Smoothing parameter for magnitude
β: Smoothing parameter for transaction frequency

Formulas

Bass Diffusion Model

n(t) = p * [Remaining Potential] + q * [Adopters] * [Remaining Potential]

$$n(t) = p[m - N(t-1)] + q\left[\frac{N(t-1)}{m}\right][m - N(t-1)]$$

Peak sales time period using Bass Diffusion assuming a continuous value of time (not year by year).

$$t^* = \frac{\ln\left(\frac{q}{p}\right)}{p+q}$$

We can use regression to estimate Bass Diffusion parameters. The dependent variable, n_t, is a function of the previous cumulative sales, N_{t-1}, and its square and takes the form:

$$n_t = \beta_0 + \beta_1 N_{t-1} + \beta_2 (N_{t-1})^2$$

Then, in order to find the values of m (total number of customers), p (innovation factor), and q (imitation factor), we use the following formulas:

$$m = \frac{\beta_1 \pm \sqrt{\beta_1^2 - 4\beta_2\beta_0}}{-2\beta_2}$$

$$p = \frac{\beta_0}{m}$$

$$q = -\beta_2 m$$

Croston's Method

We can use Croston's method when demand is intermittent. It allows us to use the traditional exponential smoothing methods. We assume the Demand Process is
$x_t = y_t z_t$ and that demand is independent between time periods, so that the probability that a transaction occurs in the current time period is 1/n:

$$Prob(y_t = 1) = \frac{1}{n} \; and \; Prob(y_t = 0) = 1 - \frac{1}{n}$$

Updating Procedure:
If $x_t = 0$ (no transaction occurs), then

$$\hat{z}_t = \hat{z}_{t-1} \; and \; \hat{n}_t = \hat{n}_{t-1}$$

If $x_t > 0$ (transaction occurs), then

$$\hat{z}_t = \alpha x_t + (1 - \alpha)\hat{z}_{t-1}$$
$$\hat{n}_t = \beta n_t + (1 - \beta)\hat{n}_{t-1}$$

Forecast:

$$\hat{x}_{t,t+1} = \frac{\hat{z}_t}{\hat{n}_t}$$

Learning Objectives

- Forecasting is part of the entire Demand Planning and Management process.
- Range forecasts are better than point forecasts, aggregated forecasts are better than dis-aggregated, and shorter time horizons are better than longer.
- Forecasting metrics need to capture bias and accuracy.
- Understand how to initialize a forecast.
- Understand that Time Series is a useful technique when we believe demand follows certain repeating patterns.
- Recognize that all time series models make a trade-off between being naïve (using only the last most recent data) or cumulative (using all of the available data).
- Understand how exponential smoothing treats old and new information differently.
- Understand how changing the alpha or beta smoothing factors influences the forecasts.
- Understand how seasonality can be handled within exponential smoothing.
- Understand why demand for new products need to be forecasted with different techniques.
- Learn how to use basic Diffusion Models for new product demand and how to forecast intermittent demand using Croston's Method.
- Understand how the typical new product pipeline process (stage-.-gate) works and how forecasting fits in.

References

General Demand Forecasting

- Makridakis, Spyros, Steven C. Wheelwright, and Rob J. Hyndman. *Forecasting: Methods and Applications*. New York, NY: Wiley, 1998. ISBN 9780471532330.
- Hyndman, Rob J. and George Athanasopoulos. *Forecasting: Principles and Practice*. OTexts, 2014. ISBN 0987507109.
- Gilliland, Michael. *The Business Forecasting Deal: Exposing Bad Practices and Providing Practical Solutions*. Hoboken, NJ: Wiley, 2010. ISBN 0470574437.

Within the texts mentioned earlier: Silver, Pyke, and Peterson Chapter 4.1; Chopra & Meindl Chapter 7.1-7.4; Nahmias Chapter 2.1-2.6.
Also, I recommend checking out the Institute of Business Forecasting & Planning (https://ibf.org/) and their Journal of Business Forecasting.

For Time Series Analysis
Within the texts mentioned earlier: Silver, Pyke, and Peterson Chapter 4.2-5.5.1 & 4.6; Chopra & Meindl Chapter 7.5-7.6; Nahmias Chapter 2.7.
Also, I recommend checking out the Institute of Business Forecasting & Planning (https://ibf.org/) and their Journal of Business Forecasting.

- Makridakis, Spyros, Steven C. Wheelwright, and Rob J. Hyndman.*Forecasting: Methods and Applications*. New York, NY: Wiley, 1998. ISBN 9780471532330.
- Hyndman, Rob J. and George Athanasopoulos. *Forecasting: Principles and Practice*. OTexts, 2014. ISBN 0987507109.

For Exponential Smoothing

- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 4.
- Chopra, Sunil, and Peter Meindl. *Supply Chain Management: Strategy, Planning, and Operation*. 5th edition, Pearson Prentice Hall, 2013. Chapter 7.
- Nahmias, S. Production and Operations Analysis. McGraw-Hill International Edition. ISBN: 0-07-2231265-3. Chapter 2.

For Special Cases

- Cooper, Robert G. *Winning at New Products: Accelerating the Process from Idea to Launch*. Cambridge, MA: Perseus Pub., 2001. Print.
- Kahn, Kenneth B. *New Product Forecasting: An Applied Approach*. Armonk, NY: M.E. Sharpe, 2006.
- Adams, Marjorie. *PDMA Foundation NPD Best Practices Study: The PDMA Foundation's 2004 Comparative Performance Assessment Study (CPAS)*. Oak Ridge, NC: Product Development & Management Association, 2004.

Inventory Management

Summary

Inventory management is at the core of all supply chain and logistics management. There are many reasons for holding inventory including minimizing the cost of controlling a system, buffering against uncertainties in demand, supply, delivery and manufacturing, as well as covering the time required for any process. Having inventory allows for a smoother operation in most cases since it alleviates the need to create product from scratch for each individual demand. Inventory is the result of a push system where the forecast determines how much inventory of each item is required.

There is, however, a problem with having too much inventory. Excess inventory can lead to spoilage, obsolescence, and damage. Also, spending too much on inventory limits the resources available for other activities and investments. Inventory analysis is essentially the determination of the right amount of inventory of the right product in the right location in the right form. Strategic decisions cover the inventory implications of product and network design. Tactical decisions cover deployment and determine what items to carry, in what form (raw materials, work-in-process, finished goods, etc.), and where. Finally, operational decisions determine the replenishment policies (when and how much) of these inventories.

We seek the Order Replenishment Policy that minimizes these total costs and specifically the Total Relevant Costs (TRC). A cost component is considered relevant if it impacts the decision at hand and we can control it by some action. A Replenishment Policy essentially states two things: the quantity to be ordered, and when it should be ordered. As we will see, the exact form of the Total Cost Equation used depends on the assumptions we make in terms of the situation. There are many different assumptions inherent in any of the models we will use, but the primary assumptions are made concerning the form of the demand for the product (whether it is constant or variable, random or deterministic, continuous or discrete, etc.).

Key Concepts

Reasons to Hold Inventory

- Cover process time
- Allow for uncoupling of processes
- Anticipation/Speculation
- Minimize control costs
- Buffer against uncertainties such as demand, supply, delivery, and manufacturing.

Inventory Decisions

- Strategic supply chain decisions are long term and include decisions related to the supply chain such as potential alternatives to holding inventory and product design.
- Tactical are made within a month, a quarter or a year and are known as deployment decisions such as what items to carry as inventory, in what form to carry items and how much of each item to hold and where.
- Operational level decisions are made on daily, weekly or monthly basis and replenishment decisions such as how often to review inventory status, how often to make replenishment decisions and how large replenishment should be. The replenishment decisions are critical to determine how the supply chain is set up.

Inventory Classification

- Financial/Accounting Categories: Raw Materials, Work in Progress (WIP), Components/Semi-Finished Goods and Finished Goods. This category does not help in tracking opportunity costs and how one may wish to manage inventory.
- Functional (See Figure 4):
    - Cycle Stock – Amount of inventory between deliveries or replenishments
    - Safety Stock – Inventory to cover or buffer against uncertainties
    - Pipeline Inventory – Inventory when order is placed but has not yet arrived

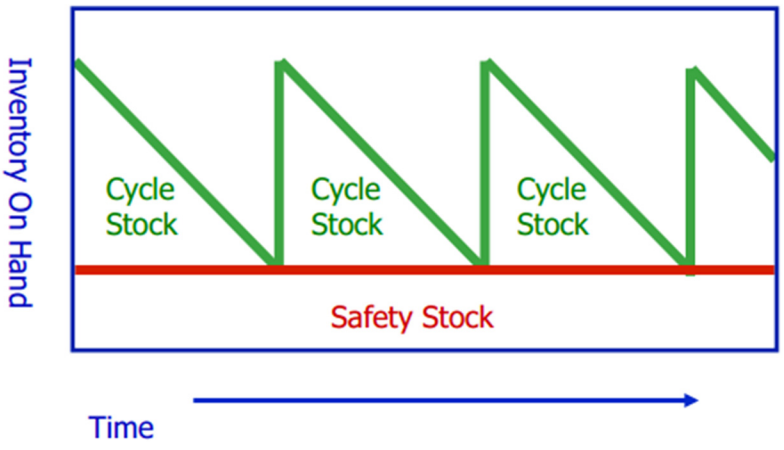


Figure 4. Inventory chart: Depiction of functional inventory classifications

Relevant Costs

The Total Cost (TC) equation is typically used to make the decisions of how much inventory to hold and how to replenish. It is the sum of the Purchasing, Ordering, Holding, and Shortage costs. The Purchasing costs are usually variable or per-item costs and cover the total landed cost for acquiring that product – whether from internal manufacturing or purchasing it from outside.

Total cost = Purchase (Unit Value) Cost + Order (Set Up) Cost + Holding (Carrying) Cost + Shortage (stock-out) Cost

- Purchase: Cost per item or total landed cost for acquiring product.
- Ordering: It is a fixed cost and contains cost to place, receive and process a batch of good including processing invoicing, auditing, labor, etc. In manufacturing this is the set up cost for a run.
- Holding: Costs required to hold inventory such as storage cost (warehouse space), service costs (insurance, taxes), risk costs (lost, stolen, damaged, obsolete), and capital costs (opportunity cost of alternative investment).
- Shortage: Costs of not having an item in stock (on-hand inventory) to satisfy a demand when it occurs, including backorder, lost sales, lost customers, and disruption costs. Also known as the penalty cost.

A cost is relevant if it is controllable and it applies to the specific decision being made.

Notation

c: Purchase cost ($/unit)
c_t: Ordering Costs ($/order)
h: Holding rate – usually expressed as a percentage ($/$ value/time)
c_e: Excess holding Costs ($/unit-time); also equal to ch
c_s: Shortage costs ($/unit)
TRC: Total Relevant Costs – the sum of the relevant cost components
TC: Total Costs – the sum of all four cost elements

Economic Order Quantity (EOQ)

The Economic Order Quantity or EOQ is the most influential and widely used (and sometimes misused!) inventory model in existence. While very simple, it provides deep and useful insights. Essentially, the EOQ is a trade-off between fixed (ordering) and variable (holding) costs. It is often called Lot-Sizing as well. The minimum of the Total Cost equation (when assuming demand is uniform and deterministic) is the EOQ or Q*. The Inventory Replenishment Policy becomes “Order Q* every T* time periods” which under our assumptions is the same as “Order Q* when Inventory Position (IP)=0”.

Like Wikipedia, the EOQ is a GREAT place to start, but not necessarily a great place to finish. It is a good first estimate because it is exceptionally robust. For example, a 50% increase in Q over the optimal quantity (Q*) only increases the TRC by ~ 8%!

While very insightful, the EOQ model should be used with caution as it has restrictive assumptions (uniform and deterministic demand). It can be safely used for items with relatively stable demand and is a good first-cut "back of the envelope" calculation in most situations. It is helpful to develop insights in understanding the trade-offs involved with taking certain managerial actions, such as lowering the ordering costs, lowering the purchase price, changing the holding costs, etc.

EOQ Model

- Assumptions
  - Demand is uniform and deterministic.
  - Lead time is instantaneous (0) – although this is not restrictive at all since the lead time, L, does not influence the Order Size, Q.
  - Total amount ordered is received.
- Inventory Replenishment Policy
  - Order Q* units every T* time periods.
  - Order Q* units when inventory on hand (IOH) is zero.
- Essentially, the Q* is the Cycle Stock for each replenishment cycle. It is the expected demand for that amount of time between order deliveries.

Notation

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: Demand (units/time)
D_A: Actual Demand (units/time)
D_F: Forecasted Demand (units/time)
h: Carrying or holding cost ($/inventory $/time)
Q: Replenishment Order Quantity (units/order)
Q*: Optimal Order Quantity under EOQ (units/order)
Q^*_A: Optimal Order Quantity with Actual Demand (units/order)
Q^*_F: Optimal Order Quantity with Forecasted Demand (units/order)
T: Order Cycle Time (time/order)
T*: Optimal Time between Replenishments (time/order)
N: Orders per Time or 1/T (order/time)
TRC(Q): Total Relevant Cost ($/time)
TC(Q): Total Cost ($/time)

Formulas

Total Costs: TC = Purchase + Order + Holding + Shortage
This is the generic total cost equation. The specific form of the different elements depends on the assumptions made concerning the demand, the shortage types, etc.

$$TC(Q) = cD + c_t\left(\frac{D}{Q}\right) + c_e\left(\frac{Q}{2}\right) + c_s E[Units\ Short]$$

Total Relevant Costs: TRC = Order + Holding
The purchasing cost and the shortage costs are not relevant for the EOQ because the purchase price does not change the optimal order quantity (Q*) and since we have deterministic demand, we will not stock out.

$$TRC(Q) = c_t\left(\frac{D}{Q}\right) + c_e\left(\frac{Q}{2}\right)$$

Optimal Order Quantity (Q*)
Recall that this is simply the First Order condition of the TRC equation – where it is a global minimum.

$$Q^* = \sqrt{\frac{2c_t D}{c_e}}$$

Optimal Time between Replenishments
Recall that T* = Q*/D. That is, the time between orders is the optimal order size divided by the annual demand. Similarly, the number of replenishments per year is simply N* = 1/T* = D/Q*. Plugging in the actual Q* gives you the formula below.

$$T^* = \sqrt{\frac{2c_t}{Dc_e}}$$

Note: Be sure to put T into units that make sense (days, weeks, months, etc.). Don't leave it in years!*

Optimal Total Costs
Adding the purchase cost to the TRC(Q*) costs gives you the TC(Q*). We still assume no stock out costs.

$$TC(Q^*) = cD + \sqrt{2c_t c_e D}$$

Optimal Total Relevant Costs
Plugging the Q* back into the TRC equation and simplifying gives you the formula below.

$$TRC(Q^*) = \sqrt{2c_t c_e D}$$

Sensitivity Analysis

The EOQ is very robust. The following formulas provide simple ways of calculating the impact of using a non-optimal Q, an incorrect annual Demand D, or a non-optimal time interval, T.

EOQ Sensitivity with Respect to Order Quantity

The equation below calculates the percent difference in total relevant costs to optimal when using a non-optimal order quantity (Q):

$$\frac{TRC(Q)}{TRC(Q^*)} = \left(\frac{1}{2}\right)\left(\frac{Q^*}{Q} + \frac{Q}{Q^*}\right)$$

Note: If optimal quantity does not make sense, it is always better to order little more rather ordering little less.

EOQ Sensitivity with Respect to Demand

The equation below calculates the percent difference in total relevant costs to optimal when assuming an incorrect annual demand (D_F) when in fact the actual annual demand is D_A:

$$\frac{TRC(Q_F^*)}{TRC(Q_A^*)} = \left(\frac{1}{2}\right)\left(\sqrt{\frac{D_A}{D_F}} + \sqrt{\frac{D_F}{D_A}}\right)$$

EOQ Sensitivity with Respect to Time Interval between Orders

The equation below calculates the percent difference in total relevant costs to optimal when using a non-optimal replenishment time interval (T). This will become very important when finding realistic replenishment intervals. The Power of Two Policy shows that ordering in increments of 2^k time periods, we will stay within 6% of the optimal solution. For example, if the base time period is one week, then the Power of Two Policy would suggest ordering every week (2^0) or every two weeks (2^1) or every four weeks (2^2) or every eight weeks (2^3) etc. Select the interval closest to one of these increments.

$$\frac{TRC(T)}{TRC(T^*)} = \left(\frac{1}{2}\right)\left(\frac{T}{T^*} + \frac{T^*}{T}\right)$$

Economic Order Quantity (EOQ) Extensions

The Economic Order Quantity can be extended to cover many different situations, three extensions include: lead-time, volume discounts, and finite replenishment or EPQ.

We developed the EOQ previously assuming the rather restrictive (and ridiculous) assumption that lead-time was zero. That is, instantaneous replenishment like on Star Trek. However, including a non-zero lead time while increasing the total cost due to having pipeline inventory

will NOT change the calculation of the optimal order quantity, Q*. In other words, lead-time is not relevant to the determination of the needed cycle stock.

Volume discounts are more complicated. Including them makes the purchasing costs relevant since they now impact the order size. We discussed three types of discounts: All-Units (where the discount applies to all items purchased if the total amount exceeds the break point quantity), Incremental (where the discount only applies to the quantity purchased that exceeds the breakpoint quantity), and One-Time (where a one-time-only discount is offered and you need to determine the optimal quantity to procure as an advance buy). Discounts are exceptionally common in practice as they are used to incentivize buyers to purchase more or to order in convenient quantities (full pallet, full truckload, etc.).

Finite Replenishment is very similar to the EOQ model, except that the product is available at a certain production rate rather than all at once. In the lesson we show that this tends to reduce the average inventory on hand (since some of each order is manufactured once the order is received) and therefore increases the optimal order quantity.

- Lead time is greater than 0 (order not received instantaneously)
    - Inventory Policy:
        - Order Q* units when IP=DL
        - Order Q* units every T* time periods
- Discounts
    - All Units Discount—Discount applies to all units purchased if total amount exceeds the break point quantity
    - Incremental Discount—Discount applies only to the quantity purchased that exceeds the break point quantity
    - One-Time-Only Discount—A one-time-only discount applies to all units you order right now (no quantity minimum or limit)
- Finite Replenishment
    - Inventory becomes available at a rate of P units/time rather than all at one time
    - If Production rate approach infinity, model converges to EOQ

Notation

c: Purchase cost ($/unit)
c_i: Discounted purchase price for discount range i ($/unit)
c^e_i: Effective purchase cost for discount range i ($/unit) [for incremental discounts]
c_t: Ordering Costs ($/order)
c_e: Excess holding Costs ($/unit/time); Equal to ch
c_s: Shortage costs ($/unit)
c_g: One Time Good Deal Purchase Price ($/unit)

F_i: Fixed Costs Associated with Units Ordered below Incremental Discount Breakpoint i
D: Demand (units/time)
D_A: Actual Demand (units/time)
D_F: Forecasted Demand (untis/time)
h: Carrying or holding cost ($/inventory $/time)
L: Order Leadtime
Q: Replenishment Order Quantity (units/order)
Q*: Optimal Order Quantity under EOQ (units/order)
Q_i: Breakpoint for quantity discount for discount i (units per order)
Q_g: One Time Good Deal Order Quantity
P: Production (units/time)
T: Order Cycle Time (time/order)
T*: Optimal Time between Replenishments (time/order)
N: Orders per Time or 1/T (order/time)
TRC(Q): Total Relevant Cost ($/time)
TC(Q): Total Cost ($/time)

Formulas

Inventory Position
Inventory Position (IP) = Inventory on Hand (IOH) + Inventory on Order (IOO) – Back Orders (BO) – Committed Orders (CO)
Inventory on Order (IOO) is the inventory that has been ordered, but not yet received. This is inventory in transit and also knows as Pipeline Inventory (PI).

Average Pipeline Inventory
Average Pipeline Inventory (API), on average, is the annual demand times the lead time. Essentially, every item spends L time periods in transit.

$$API = DL$$

Total Cost including Pipeline Inventory
The TC equation changes slightly if we assume a non-zero lead time and include the pipeline inventory.

$$TC(Q) = cD + c_t\left(\frac{D}{Q}\right) + c_e\left(\frac{Q}{2} + DL\right) + c_s E[Units\ Short]$$

Note that as before, though, the purchase cost, shortage costs, and now pipeline inventory is not relevant to determining the optimal order quantity, Q*:

$$Q^* = \sqrt{\frac{2c_t D}{c_e}}$$

Discounts

If we include volume discounts, then the purchasing cost becomes relevant to our decision of order quantity.

All Units Discounts

Discount applies to all units purchased if total amount exceeds the break point quantity.
The procedure for a single range All Units quantity discount (where new price is c_1 if ordering at least Q_1 units) is as follows:

1. Calculate Q^*_{C0}, the EOQ using the base (non-discounted) price, and Q^*_{C1}, the EOQ using the first discounted price
2. If $Q^*_{C1} \geq Q_1$, the breakpoint for the first all units discount, then order Q^*_{C1} since it satisfies the condition of the discount. Otherwise, go to step 3.
3. Compare the $TRC(Q^*_{C0})$, the total relevant cost with the base (non-discounted) price, with $TRC(Q_1)$, the total relevant cost using the discounted price (c_1) at the breakpoint for the discount. If $TRC(Q^*_{C0}) < TRC(Q_1)$, select Q^*_{C0}, otherwise order Q_1.

Note that if there are more discount levels, you need to check this for each one.

$$c = c_0 \; for \; 0 \leq Q \leq Q_1 \; and \; c = c_1 \; for \; Q_1 \leq Q$$

$$TRC = Dc_0 + c_t\left(\frac{D}{Q}\right) + c_0\left(\frac{hQ}{2}\right) \; for \; 0 \leq Q \leq Q_1$$

$$TRC = Dc_1 + c_t\left(\frac{D}{Q}\right) + c_1\left(\frac{hQ}{2}\right) \; for \; Q_1 \leq Q$$

Note: All units discount tend to raise cycle stock in the supply chain by encouraging retailers to increase the size of each order. This makes economic sense for the manufacturer, especially when he incurs a very high fixed cost per order.

Incremental Discounts

Discount applies only to the quantity purchased that exceeds the break point quantity.
The procedure for a multi-range Incremental quantity discount (where if ordering at least Q_1 units, the new price for the $Q\text{-}Q_1$ units is c_1) is as follows:

1. Calculate the Fixed cost per breakpoint, F_i,
2. Calculate the Q^*_i for each discount range i (to include the F_i)
3. Calculate the TRC for all discount ranges where the $Q_{i-1} < Q^*_i < Q_{i+1}$, that is, if it is in range.
4. Select the discount that provides the lowest TRC.

The effective cost, c^e_i, can be used for the TRC calculations.

$$F_0 = 0 \; ; \; F_i = F_{i-1} + (c_{i-1} - c_i)Q_i$$

$$Q^* = \sqrt{\frac{2D(c_t + F_i)}{hc_i}}$$

$$c_i^e = c_i + \frac{F_i}{Q_i^*}$$

One Time Discount

This is a less common discount – but it does happen. A one time only discount applies to all units you order right now (no minimum quantity or limit).

Simply calculate the Q^*_g and that is your order quantity. If $Q^*_g = Q^*$ then the discount does not make sense. If you find that $Q^*_g < Q^*$, you made a mathematical mistake – check your work!

$$TC = (CycleTime)(TC^* + PurchaseCost) = \left(\frac{Q_g}{D}\right)\sqrt{2c_t hcD} + \left(\frac{Q_g}{D}\right)cD$$

$$Savings = TC - TC_{SP}$$

$$Savings = \left(\left(\frac{Q_g}{D}\right)\sqrt{2c_t hcD} + \left(\frac{Q_g}{D}\right)cD\right) - \left(c_g Q_g + hc_g\left(\frac{Q_g}{2}\right)\left(\frac{Q_g}{D}\right) + c_t\right)$$

$$Q_g^* = \frac{Q^* ch + D(c - c_g)}{hc_g}$$

Finite Replenishment or Economic Production Quantity

One can think of the EPQ equations as generalized forms where the EOQ is a special case where P=infinity. As the production rate decreases, the optimal quantity to be ordered increases. However, note that if P<D, this means the rate of production is slower than the rate of demand and that you will never have enough inventory to satisfy demand.

$$TRC[Q] = \frac{c_t D}{Q} + \frac{Q\left(1 - \frac{D}{P}\right)hc}{2}$$

$$EPQ = \sqrt{\frac{2c_t D}{hc\left(1 - \frac{D}{P}\right)}} = \frac{EOQ}{\sqrt{\left(1 - \frac{D}{P}\right)}}$$

Single Period Inventory Models

The single period inventory model is second only to the economic order quantity in its widespread use and influence. Also referred to as the Newsvendor (Newsboy) model, the single period model differs from the EOQ in three main ways. First, while the EOQ assumes uniform and deterministic demand, the single-period model allows demand to be variable and stochastic (random). Second, while the EOQ assumes a steady state condition (stable demand with essentially an infinite time horizon), the single-period model assumes a single period of time. All inventories must be ordered prior to the start of the time period and they cannot be replenished during the time period. Any inventory left over at the end of the time period is scrapped and cannot be used at a later time. If there is extra demand that is not satisfied

during the period, it too is lost. Third, for EOQ we are minimizing the expected costs, while for the single period model we are actually maximizing the expected profitability.

A planned backorder is where we stock out on purpose knowing that customers will wait, although we do incur a penalty cost, c_s, for stocking out. From this, we develop the idea of the critical ratio (CR), which is the ratio of the c_s (the cost of shortage or having too little product) to the ratio of the sum of c_s and c_e (the cost of having too much or an excess of product). The critical ratio, by definition, ranges between 0 and 1 and is a good metric of level of service. A high CR indicates a desire to stock out less frequently. The EOQ with planned backorders is essentially the generalized form where c_s is essentially infinity, meaning you will never ever stock out. As c_s gets smaller, the Q^*_{PBO} gets larger and a larger percentage is allowed to be backordered – since the penalty for stocking out gets reduced.

The critical ratio applies directly to the single period model as well. We show that the optimal order quantity, Q*, occurs when the probability that the demand is less than Q* = the Critical Ratio. In other words, the Critical Ratio tells me how much of the demand probability that should be covered in order to maximize the expected profits.

Marginal Analysis: Single Period Model

Two costs are associated with single period problems

- Excess cost (c_e) when D<Q ($/unit) i.e. too much product
- Shortage cost (c_s) when D>Q ($/unit) i.e. too little product

If we assume continuous distribution of demand

- c_e P[X≤Q] = expected excess cost of the Qth unit ordered
- c_s (1-P[X≤Q]) = expected shortage cost of the Qth unit ordered

This implies that if E[Excess Cost] < E[Shortage Cost] then increase Q and that we are at Q* when E[Shortage Cost] = E[Excess Cost]. Solving this gives us: $P[x \leq Q] = \frac{c_s}{(c_e+c_s)}$

In words, this means that the percentage of the demand distribution covered by Q should be equal to the Critical Ratio in order to maximize expected profits.

Notation

B: Penalty for not satisfying demand beyond lost profit ($/unit)
b: Backorder Demand (units)
b*: Optimal units on backorder when placing an order (unit)
c: Purchase cost ($/unit)
c_t: Ordering Costs ($/order)
c_e: Excess holding Costs ($/unit/time); Equal to ch
c_s: Shortage Costs ($/unit)
D: Average Demand (units/time)
g: Salvage value for excess inventory ($/unit)
h: Carrying or holding cost ($/inventory $/time)

L: Replenishment Lead Time (time)
Q: Replenishment Order Quantity (units/order)
Q^*_{PBO}: Optimal Order Quantity with Planned backorders
T: Order Cycle Time (time/order)
TRC(Q): Total Relevant Cost ($/time)
TC(Q): Total Cost ($/time)

Formulas

EOQ with Planned Backorders

This is an extension of the standard EOQ with the ability to allow for backorders at a penalty of c_s.

$$TRC(Q,b) = c_t\left(\frac{D}{Q}\right) + c_e\left(\frac{(Q-b)^2}{2Q}\right) + c_s\left(\frac{b^2}{2Q}\right)$$

$$Q^*_{PBO} = \sqrt{\frac{2c_tD}{c_e}}\sqrt{\frac{c_sc_e}{c_s}} = Q^*\sqrt{\frac{(c_s+c_e)}{c_s}} = Q^*\sqrt{\frac{1}{CR}}$$

$$b^* = \frac{c_eQ^*_{PBO}}{(c_s+c_e)} = \left(1-\frac{c_s}{(c_s+c_e)}\right)Q^*_{PBO}$$

$$T^*_{PBO} = \frac{D}{Q^*_{PBO}}$$

Order Q^*_{PBO} when IOH = -b*; Order Q^*_{PBO} every T^*_{PBO} time periods

Single Period (Newsvendor) Model

To maximize expected profitability, we need to order sufficient inventory, Q, such that the probability that the demand is less than or equal to this amount is equal to the Critical Ratio. Thus, the probability of stocking out is equal to 1 – CR.

$$P[x \le Q] = \frac{c_s}{(c_e+c_s)}$$

For the simplest case where there is neither salvage value nor extra penalty of stocking out, these become:

c_s = p – c, that is the lost margin of missing a potential sale and,
c_e = c, that is, the cost of purchasing one unit.

The Critical Ratio becomes: $CR = \frac{c_s}{c_s+c_e} = \frac{(p-c)}{(p-c+c)} = \frac{p-c}{p}$ which is simply the margin divided by the price!

When we consider also salvage value (g) and shortage penalty (B), these become:

c_s = p – c + B, that is the lost margin of missing a potential sale plus a penalty per item short and

c_e = c – g, that is, the cost of purchasing one unit minus the salvage value I can gain back.

Now the critical ratio becomes

$$CR = \frac{c_s}{c_s + c_e} = \frac{(p - c + B)}{(p - c + B + c - g)} = \frac{(p - c + B)}{(p + B - g)}$$

Single Period Inventory Models-Expected Profitability

We expand our analysis of the single period model to be able to calculate the expected profitability of a given solution. In the previous lesson, we learned how to determine the optimal order quantity, Q*, such that the probability of the demand distribution covered by Q* is equal to the Critical Ratio, which is the ratio of the shortage costs divided by the sum of the shortage and excess costs.

In order to determine the profitability for a solution, we need to calculate the expected units sold, the expected cost of buying Q units, and the expected units short, E[US]. Calculating the E[US] is tricky, but we show how to use the Normal Tables as well as spreadsheets to determine this value.

Notation

B: Penalty for not satisfying demand beyond lost profit ($/unit)
c: Purchase cost ($/unit)
c_t: Ordering Costs ($/order)
c_e: Excess holding Costs ($/unit); For single period problems this is not necessarily equal to ch, since that assumes that you can keep the inventory for later use.
c_s: Shortage Costs ($/unit)
D: Average Demand (units/time)
g: Salvage value for excess inventory ($/unit)
k: Safety Factor
x: Units Demanded
E[US]: Expected Units Short (units)
Q: Replenishment Order Quantity (units/order)
TRC(Q): Total Relevant Cost ($/period)
TC(Q): Total Cost ($/period)

Formulas

Profit Maximization

In words, the expected profit for ordering Q units is equal to the sales price, p, times the expected number of units sold, E[x]), minus the cost of purchasing Q units, cQ, minus the expected number of units I would be short times the sales price. The difficult part of this equation is the expected units short, or the E[US].

$$E[Profit(Q)] = pE[x] - cQ - pE[UnitsShort]$$

Expected Profits with Salvage and Penalty

If we include a salvage value, g, and a shortage penalty, B, then this becomes:

$$P(Q) = \begin{cases} -cQ + px + g(Q - x)\ if\ x \le Q \\ -cQ + pQ - B(x - Q) if\ x \ge Q \end{cases}$$

$$E[P(Q)] = (p - g)E[x] - (c - g)Q - (p - g + B)E[US]$$

Rearranging this becomes:

$$E[P(Q)] = p(E[x] - E[US]) - cQ + g\big(Q - (E[x] - E[US])\big) - B(E[US])$$

In words, the expected profit for ordering Q units is equal to four terms. The first term is the sales price, p, times the expected number of units sold, E[x]), minus the expected units short. The second term is simply the cost of purchasing Q units, cQ. The third term is the expected number of items that I would have left over for salvage, times the salvage value, g. The fourth and final term is the expected number of units short times the shortage penalty, B.

Expected Values

E[Units Demanded]

Continuous: $\int_{x=0}^{\infty} x f_x(x)dx = \hat{x}$ Discrete: $\sum_{x=0}^{\infty} xP[x] = \hat{x}$

E[Units Sold]

Continuous: $\int_{x=0}^{Q} x f_x(x)dx + Q \int_{x=Q}^{\infty} f_x(x)dx$ Discrete: $\sum_{x=0}^{Q} xP[x] + Q \sum_{x=Q+1}^{\infty} P[x]$

E[Units Short]

Continuous: $\int_{x=Q}^{\infty} (x - Q) f_x(x)dx$ Discrete: $\sum_{x=Q+1}^{\infty} (x - Q)P[x]$

Expected Units Short E[US]
This is a tricky concept to get your head around at first. Think of the E[US] as the average (mean or expected value) of the demand ABOVE some amount that we specify or have on hand. As my Q gets larger, then we expect the E[US] to get smaller, since I will probably not stock out as much.

Luckily for us, we have a nice way of calculating the E[US] for the Normal Distribution. The Expected Unit Normal Loss Function is noted as G(k). To find the actual units short, we simply multiply this G(k) times the standard deviation of the probability distribution.

$$E[US] = \int_{x=Q}^{\infty} (x - Q) f_x(x) dx = \sigma G\left(\frac{Q - \mu}{\sigma}\right) = \sigma G(k)$$

You can use the Normal tables to find the G(k) for a given k value or you can use spreadsheets with the equation below:

$$G(k) = NORMDIST(k, 0,1,0) - k * (1 - NORMSDIST(k))$$

Probabilistic Inventory Models

We 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 = μ_{DL}
- **Single Period / Newsvendor** — Probabilistic Demand with finite (single period) horizon
  - Order Q* at start of period where P[x ≤ 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 ≤ 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 covered are 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 Stock Out 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.

Notation

- B_1: Cost associated with a stock out 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 four items, I order four 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

Here are 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 = NORMINVDIST(CSL, Mean, StandardDeviation) and k=NORMSINV(CSL).

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

Note that as k increases, it gets difficult to improve CSL and it will require enormous amount of inventory to cover the extreme limits.

Cost Per Stock out 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).

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

$$G(k) = \frac{Q}{\sigma_{DL}}(1 - IFR)$$

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

$$G(k) = NORMDIST(k, 0,1,0) - k * (1 - NORMSDIST(k))$$

Once we find the k using unit normal tables, we can plug the values in $s = \mu_{DL} + k\sigma_{DL}$ to frame the policy.

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 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 $\frac{Qc_e}{Dc_s} \leq 1$, then we should find k such that:

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

Otherwise, we should set k as low as management allows. In a spreadsheet, this becomes k=NORMSINV(1 - $\frac{Qc_e}{Dc_s}$)

Summary of the Metrics Presented

	Metric	**How to find k**
% Service Based	Cycle Service Level (CSL)	K = NORMSINV(1 − P[X > s])
% Service Based	Item Fill Rate (IFL)	Find k from $G(k) = \frac{Q}{\sigma_{DL}}(1 - IFR)$
$ Cost Based	Cost per Stock Out Event (CSOE)	$k = \sqrt{2 \ln\left(\frac{B_1 D}{c_e \sigma_{DL} Q \sqrt{2\pi}}\right)}$
$ Cost Based	Cost per Item Short (CIS)	K = NORMSINV(1 - $\frac{Qc_e}{Dc_s}$)

Table 2. Summary of metrics presented

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):

$$E[D_S] = E[D_L]/n$$
$$VAR[D_S] = VAR[D_L]/n$$
$$\sigma_s = \sigma_L/\sqrt{n}$$

Converting from short to long:

$$E[D_L] = nE[D_S]$$
$$VAR[D_L] = nVAR[D_S]$$
$$\sigma_L = \sqrt{n}\sigma_s$$

Periodic Review Policies

There are trade-offs between the different performance metrics (both cost- and service-based). We demonstrate that once one of the metrics is determined (or explicitly set) then the other three are implicitly set. Because they all lead to the establishment of a safety factor, k, they are

dependent on each other. This means that once you have set the safety stock, regardless of the method, you can calculate the expected performance implied by the remaining three metrics.

Periodic Review policies are very popular because they fit the regular pattern of work where ordering might occur only once a week or once every two weeks. The lead-time and the review period are related and can be traded-off to achieve certain goals.

Notation

B_1: Cost associated with a stock out 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)
c_g: One Time Good Deal Purchase Price ($/unit)
D: Average Demand (units/time)
h: Carrying or holding cost ($/inventory $/time)
L: Replenishment Lead Time (time)
Q: Replenishment Order Quantity (units/order)
T: Order Cycle Time (time/order)
μ_{DL}: Expected Demand over Lead Time (units/time)
σ_{DL}: Standard Deviation of Demand over Lead Time (units/time)
μ_{DL+R}: Expected Demand over Lead Time plus Review Period (units/time)
σ_{DL+R}: Standard Deviation of Demand over Lead Time plus Review Period (units/time)
k: Safety Factor
s: Reorder Point (units)
S: Order up to Point (units)
R: Review Period (time)
N: Orders per Time or 1/T (order/time)
IP: Inventory Position = Inventory on Hand + Inventory on Order – Backorders
IOH: Inventory on Hand (units)
IOO: Inventory on Order (units)
IFR: Item Fill Rate (%)
CSL: Cycle Service Level (%)
CSOE: Cost of Stock Out Event ($ / event)
CIS: Cost per Item Short
E[US]: Expected Units Short (units)
G(k): Unit Normal Loss Function

Formulas

Inventory Performance Metrics

Safety stock is determined by the safety factor, k. So that: $s = \mu_{DL} + k\sigma_{DL}$ and the expected cost of safety stock = $c_e k\sigma_{DL}$.

Two ways to calculate k: Service based or Cost based metrics:

- Service Based Metrics—set k to meet expected level of service
  - Cycle Service Level ($CSL = P[x \leq k]$)
  - Item Fill Rate ($IFR = 1 - \frac{\sigma_{DL} G[k]}{Q}$)

Note: IFR is always higher than CSL for the same safety stock level.

- Cost Based Metrics—find k that minimizes total costs
  - Cost per Stock out Event ($\mathrm{E[CSOE]} = (B_1)P[x \geq k]\left(\frac{D}{Q}\right)$)
  - Cost per Items Short ($E[CIS] = c_s \sigma_{DL} G(k)\left(\frac{D}{Q}\right)$)

Safety Stock Logic – relationship between performance metrics

The relationship between the four metrics (2 cost and 2 service based) is shown in the flowchart below (Figure 5). Once one metric (CSL, IFR, CSOE, or CIS) is explicitly set, then the other three metrics are implicitly determined.

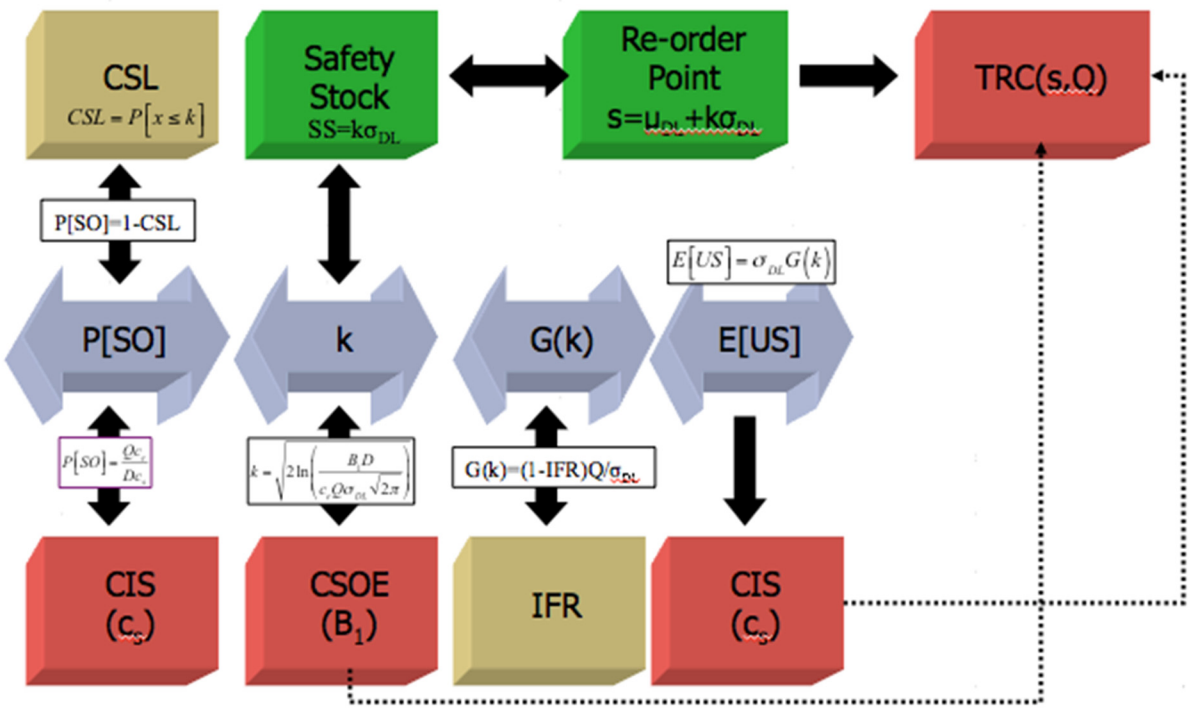


Figure 5. Relationship among the four metrics

Periodic Review Policy (R,S)

This is also known as the Order Up To policy and is essentially a two-bin system. The policy is "Order Up To S* units every R time periods". This means the order quantity will be S*-IP. The order up to point, S*, is the sum of the expected demand over the lead-time and the

replenishment time plus the RMSE of the forecast error over lead plus replenishment time multiplied by some safety factor k.

- Order Up To Point: $S = \mu_{DL+R} + k\sigma_{DL+R}$

Periodic (R,S) versus Continuous (s, Q) Review

- There is a convenient transformation of (s, Q) to (R, S)
  - (s,Q) = Continuous, order Q when IP ≤ s
  - (R, S) = Periodic, order up to S every R time periods
- Allows for the use of all previous (s, Q) decision rules
  - Reorder point, s, for continuous becomes Order Up To point, S, for periodic system
  - Q for continuous becomes D*R for periodic
  - L for a continuous becomes R+L for periodic
- Approach
  - Make transformations
  - Solve for (s, Q) using transformations
  - Determine final policy such that $S = x_{DL+R} + k\sigma_{DL+R}$

(s, Q)		(R, S)
s	↔	S
Q	↔	D*R
L	↔	R+L

Relationship Between L & R

The lead time, L, and the review period, R, both influence the total costs. Note that the average inventory costs for a (R,S) system is $= c_e[\frac{DR}{2} + k\sigma_{DL+R} + LD]$. This implies that increasing Lead Time, L, will increase Safety Stock non-linearly and Pipeline Inventory linearly while increasing the Review Period, R will increase the Safety Stock non-linearly and the Cycle Stock linearly.

Inventory Models for Multiple Items & Locations

There are several problems with managing items independently, including:

- Lack of coordination—constantly ordering items
- Ignoring of common constraints such as financial budgets or space
- Missed opportunities for consolidation and synergies
- Waste of management time

Managing Multiple Items

There are two issues to solve in order to manage multiple items:

1. Can we aggregate SKUs to use similar operating policies?
   a. Group using common cost characteristics or break points
   b. Group using Power of Two Policies
2. How do we manage inventory under common constraints?

a. Exchange curves for cycle stock
b. Exchange curves for safety stock

Aggregation Methods

When we have multiple SKUs to manage, we want to aggregate those SKUs where we can use the same policies.

Grouping Like Items—Break Points

- Basic Idea: Replenish higher value items faster
- Used for situations with multiple items that have
  - Relatively stable demand
  - Common ordering costs, c_t, and holding charges, h
  - Different annual demands, D_i, and purchase cost c_i
- Approach
  - Pick a base time period, w_0, (typically a week)
  - Create a set of candidate ordering periods (w_1, w_2, etc.)
  - Find D_ic_i values where $TRC(w_j)=TRC(w_{j+1})$
  - Group SKUs that fall in common value (D_ic_i) buckets

Power of Two Formula

- Order in time intervals of powers of two
- Select a realistic base period, T_{base} (day, week, month)
- Guarantees that TRC will be within 6% of optimal!

Managing Under Common Constraints

There is typically a budget or space constraint that limits the amount of inventory that you can actually keep on hand. Managing each inventory item separately could lead to violating this constraint. Exchange curves are a good way to use the managerial levers of holding charge, ordering cost, and safety factor to set inventory policies to meet a common constraint.

Exchange Curves: Cycle Stock

- Helps determine the best allocation of inventory budget across multiple SKUs
- Relevant Cost parameters
  - Holding Charge (h)
    - There is no single correct value
    - Cost allocations for time and systems differ between firms
    - Reflection of management's investment and risk profile
  - Order Cost (c_t)
    - Not know with precision
    - Cost allocations for time and systems differ between firms
- Exchange Curve
  - Depicts trade-off between total annual cycle stock (TACS) and number of replenishments (N)
  - Determines the c_t/h value that meets budget constraints

Exchange Curves: Safety Stock

- Need to trade-off cost of safety stock and level of service
- Key parameter is safety factor (k) – usually set by management
- Estimate the aggregate service level for different budgets
- The process is as follows:
    1. Select an inventory metric to target
    2. Starting with a high metric value calculate:
        a. The required k_i to meet that target for each SKU
        b. The resulting safety stock cost for each SKU and the total safety stock (TSS)
        c. The other resulting inventory metrics of interest for each SKU and total
    3. Lower the metric value, go to step 2
    4. Chart resulting TSS versus Inventory Metrics

Managing Multiple Locations

Managing the same item in multiple locations will lead to a higher inventory level than managing them in a single location. Consolidating inventory locations to a single common location is known as inventory pooling. Pooling reduces the cycle stock needed by reducing the number of deliveries required and reduces the safety stock by risk pooling that reduces the CV of the demand. This is also called the square root "law" – which is insightful and powerful, but also makes some restrictive assumptions, such as uniformly distributed demand, use of EOQ ordering principles, and independence of demand in different locations.

Notation

c_i: Purchase cost for item i ($/unit)
c_t: Ordering Costs ($/order)
c_e: Excess holding Costs ($/unit/time); Equal to ch
c_s: Shortage costs ($/unit)
D_i: Average Demand for item i (units/time)
h: Carrying or holding cost ($/inventory $/time)
Q: Replenishment Order Quantity (units/order)
T: Order Cycle Time (time/order)
$T_{Practical}$: Practical Order Cycle Time (time/order)
k: Safety Factor
w_0: Base Time Period (time)
s: Reorder Point (units)
R: Review Period (time)
N: Number of Inventory Replenishment Cycles
TACS: Total Annual Cycle Stock
TSS: Total Value of Safety Stock
TVIS: Total Value of Items Short
G(k): Unit Normal Loss Function

Formulas

Power of Two Policy

The process is as follows:

1. Create table of SKUs
2. Calculate T* for each SKU
3. Calculate $T_{practical}$ for each SKU

$$T^* = \frac{Q^*}{D} = \frac{\sqrt{\frac{2c_t D}{c_e}}}{D} = \sqrt{\frac{2c_t}{Dc_e}}$$

$$T_{practical} = 2^{\ln\left(\frac{T^*}{\sqrt{2}}\right)/\ln(2)}$$

In a spreadsheet this is: Tpractical = 2^(ROUNDUP(LN($T_{optimal}$ / SQRT(2)) /LN(2)))

Exchange Curves: Cycle Stock

$$TACS = \sum_{i=1}^{n} \frac{Q_i c_i}{2} = \sqrt{\frac{c_t}{h}} \frac{1}{\sqrt{2}} \sum_{i=1}^{n} \sqrt{D_i c_i} \qquad N = \sum_{i=1}^{n} \frac{D_i}{Q_i} = \sqrt{\frac{h}{c_t}} \frac{1}{\sqrt{2}} \sum_{i=1}^{n} \sqrt{D_i c_i}$$

Process

1. Create a table of SKUs with "Annual Value" ($D_i c_i$) and $\sqrt{D_i c_i}$
2. Find the sum of $\sqrt{D_i c_i}$ term for SKUs being analyzed
3. Calculate TACS and N for range of (c_t/h) values
4. Chart N vs TACS

Exchange Curves: Safety Stock

$$TSS = \sum_{i=1}^{n} k_i \sigma_{DLi} c_i \qquad TVIS = \sum_{i=1}^{n} \left(\frac{D_i}{Q_i}\right) c_i \sigma_{DLi} G(k_i))$$

Process:

1. Select an inventory metric to target
2. Starting with a high metric value calculate:
   a. The required k_i to meet that target for each SKU
   b. The resulting safety stock cost for each SKU and the total safety stock (TSS)
   c. The other resulting inventory metrics of interest for each SKU and total
3. Lower the metric value, go to step 2
4. Chart resulting TSS versus Inventory Metrics

Pooled Inventories

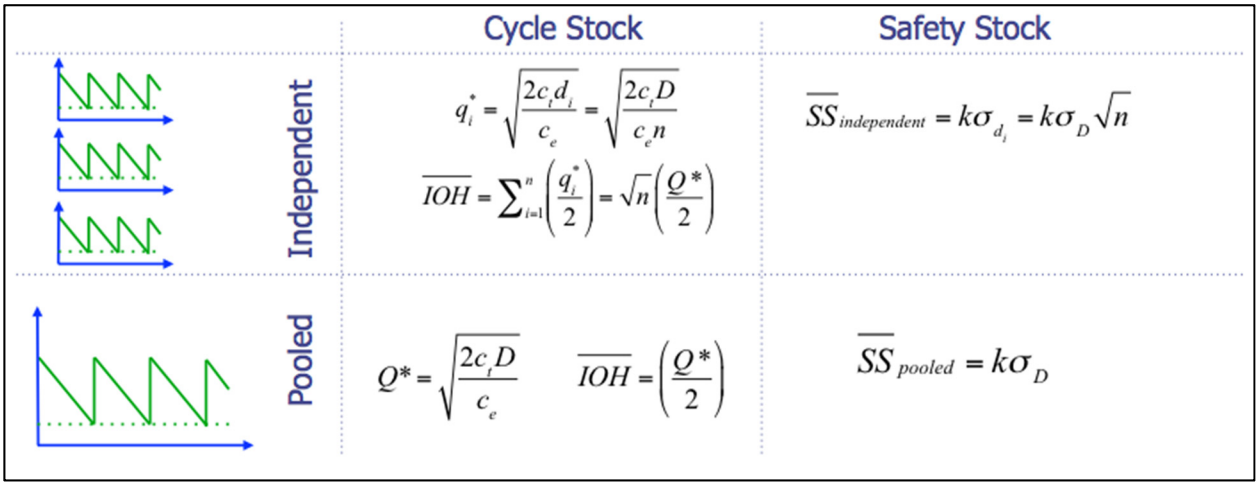


Chart 1. Comparison between independent and pooled inventories

Inventory Models for Class A & C Items

Inventory Management by Segment

	A Items	B Items	C Items
Type of Records	Extensive, Transactional	Moderate	None-use a rule
Level of Management Reporting	Frequent (Monthly or More)	Infrequently—Aggregated	Only as Aggregate
Interaction with Demand	Direct Input, High Data Integrity, Manipulate (pricing, etc.)	Modified Forecast (promotions, etc.)	Simple Forecast at Best
Interaction with Supply	Actively Manage	Manage by Exception	Non
Initial Deployment	Minimize Exposure (high v)	Steady State	Steady State
Frequency of Policy Review	Very Frequent (Monthly or More)	Moderate (Annually/Event Based)	Very Infrequent
Importance of Parameter Precision	Very High—Accuracy Worthwhile	Moderate—Rounding and Approximation ok	Very Low
Shortage Strategy	Actively Manage (Confront)	Set Service Level & Manage by Exception	Set & Forget Service Levels
Demand Distribution	Consider Alternatives to Normal as Situation Fits	Normal	N/A
Management Strategy	Active	Automatic	Passive

Table 3. Inventory management by segment

Inventory Policies (Rules of Thumb)

Type of Item	Continuous Review	Periodic Review
A Items	(s, S)	(R, s, S)
B Items	(s, Q)	(R, S)
C Items		Manual ~(R, S)

Table 4. Inventory policies (rules of Thumb)

Managing Class A Items

There are two general ways that items can be considered Class A:

- Fast Moving but Cheap (Large D, Small c → Q > 1)
- Slow Moving but Expensive (Large c, Small D → Q = 1

This dictates which Probability Distribution to use for modeling the demand

- Fast Movers
    - Normal or Lognormal Distribution
    - Good enough for B items
    - OK for A items if μ_{DL} or $\mu_{DL+R} \geq 10$
- Slow Movers
    - Poisson Distribution
    - More complicated to handle
    - OK for A items if μ_{DL} or $\mu_{DL+R} < 10$

Managing Class C Items

Class C items have low cD values but comprise the lion-share of the SKUs. When managing them we need to consider the implicit & explicit costs. The objective is to minimize management attention. Regardless of policy, savings will most likely not be significant, so try to design simple rules to follow and explore opportunities for disposing of inventory. Alternatively, try to set common reorder quantities. This can be done by assuming common c_t and h values and then finding $D_i c_i$ values for ordering frequencies.

Disposing of Excess Inventory

- Why does excess inventory occur?
    - SKU portfolios tend to grow
    - Poor forecasts - Shorter lifecycles
- Which items to dispose?
    - Look at DOS (days of supply) for each item = IOH/D
    - Consider getting rid of items that have DOS > x years
- What actions to take?
    - Convert to other uses
    - Ship to more desired location
    - Mark down price
    - Auction

Real World Inventory Challenges

While models are important, it is also important to understand where there are challenges implementing models in real life.

- Models are not used exactly as in textbooks
- Data is not always available or correct
- Technology matters
- Business processes matter even more
- Inventory policies try to answer three questions:
    - How often should I check my inventory?
    - How do I know if I should order more?
    - How much to order?
- All inventory models use two key numbers
    - Inventory Position
    - Order Point

Notation

B_1: Cost Associated with a Stock out 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)
c_g: One Time Good Deal Purchase Price ($/unit)
D: Average Demand (units/time)
h: Carrying or Holding Cost ($/inventory $/time)
$L[X_i]$: Discrete Unit Loss Function
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)
μ_{DL+R}: Expected Demand over Lead Time plus Review Period (units/time)
σ_{DL+R}: Standard Deviation of Demand over Lead Time plus Review Period (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 (IOO) – Backorders
IOH: Inventory on Hand (units)
IOO: Inventory on Order (units)
IFR: Item Fill Rate (%)
CSL: Cycle Service Level (%)
E[US]: Expected Units Short (units)
G(k): Unit Normal Loss Function

Formulas

Fast Moving A Items

$$TRC = c_t\left(\frac{D}{Q}\right) + c_e\left(\frac{Q}{2} + k\sigma_{DL}\right) + B_1\left(\frac{D}{Q}\right)P[x > k]$$

$$Q^* = EOQ\sqrt{1 + \frac{B_1 P[x > k]}{c_t}}$$

$$k^* = \sqrt{2\ln\left(\frac{DB_1}{\sqrt{2\pi}Qc_e\sigma_{DL}}\right)}$$

- Iteratively solve the two equations
- Stop when Q* and k* converge within acceptable range

Slow Moving A Items

Use a Poisson distribution to model sales

- Probability of x events occurring within a time period
- Mean = Variance = λ

$$p[x_0] = Prob[x = x_0] = \frac{e^{-\lambda}\lambda^{x_0}}{x_0!}\ for\ x_0$$

$$F[x_0] = Prob[x \le x_0] = \sum_{x=0}^{x_0}\frac{e^{-\lambda}\lambda^{x}}{x!}$$

For a discrete function, the loss function $L[X_i]$ can be calculated as follows (Cachon & Terwiesch)

$$L[X_i] = L[X_{i-1}] - (X_i - X_{i-1})(1 - F[X_{i-1}])$$

Learning Objectives

- Understand the reasons for holding inventory and the different types of inventory.
- Understand the concepts of total cost and total relevant costs.
- Identify and quantify the four major cost components of total costs: Purchasing, Ordering, Holding, and Shortage.
- Able to estimate the Economic Order Quantity (EOQ) and to determine when it is appropriate to use.
- Able to estimate sensitivity of EOQ to underlying changes in the input data and understanding of its underlying robustness.

- Understand how to determine the EOQ with different volume discounting schemes.
- Understand how to determine the Economic Production Quantity (EPQ) when the inventory becomes available at a certain rate of time instead of all at once.
- Ability to use the Critical Ratio to determine the optimal order quantity to maximize expected profits.
- Ability to established inventory policies for EOQ with planned backorders as well as single period models.
- Ability to determine profitability, expected units short, expected units sold of a single period model.
- 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.
- Able to establish a periodic review, Order Up To (S,R) Replenishment Policy using any of the four performance metrics.
- Understand relationships between the performance metrics (CSL, IFR, CSOE, and CIS) and be able to calculate the implicit values.
- Able to use the inventory models to make trade-offs and estimate impacts of policy changes.
- Understand how to use different methods to aggregate SKUs for common inventory policies.
- Understand how to use Exchange Curves.
- Understand how inventory pooling impacts both cycle stock and safety stock.
- Understand how to use different inventory models for Class A and C items.

References

For General Inventory Management

There are more books that cover the basics of inventory management than there are grains of sand on the beach! Inventory management is also usually covered in Operations Management and Industrial Engineering texts as well. A word of warning, though. Every textbook uses different notation for the same concepts. Get used to it. Always be sure to understand what the nomenclature means so that you do not get confused.

- Nahmias, S. Production and Operations Analysis. McGraw-Hill International Edition. ISBN: 0-07-2231265-3. Chapter 4.
- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 1
- Ballou, R.H. Business Logistics Management. ISBN: 978-0130661845. Chapter 9.

For EOQ

- Schwarz, Leroy B., "The Economic Order-Quantity (EOQ) Model" in Building intuition : insights from basic operations management models and principles, edited by Dilip Chhajed, Timothy Lowe, 2007, Springer, New York, (pp 135-154).
- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 5
- Ballou, R.H. Business Logistics Management. ISBN: 978-0130661845. Chapter 9.

For EOQ Extensions

- Nahmias, S. Production and Operations Analysis. McGraw-Hill International Edition. ISBN: 0-07-2231265-3. Chapter 4.
- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 5.
- Ballou, R.H. Business Logistics Management. ISBN: 978-0130661845. Chapter 9.
- Schwarz, Leroy B., "The Economic Order-Quantity (EOQ) Model" in Building intuition : insights from basic operations management models and principles, edited by Dilip Chhajed, Timothy Lowe, 2007, Springer, New York, (pp 135-154).
- Muckstadt, John and Amar Sapra "Models and Solutions in Inventory Management"., 2006, Springer New York, New York, NY. Chapter 2 & 3.

For Single Period Inventory Models

- Nahmias, S. Production and Operations Analysis. McGraw-Hill International Edition. ISBN: 0-07-2231265-3. Chapter 5.
- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 10.
- Porteus, Evan L., "The Newsvendor Problem" in Building intuition: insights from basic operations management models and principles, edited by Dilip Chhajed, Timothy Lowe, 2007, Springer, New York, (pp 115-134).
- Muckstadt, John and Amar Sapra "Models and Solutions in Inventory Management"., 2006, Springer New York, New York, NY. Chapter 5.
- Ballou, R.H. Business Logistics Management. ISBN: 978-0130661845. Chapter 9.

For Probabilistic Inventory Models

- Nahmias, S. Production and Operations Analysis. McGraw-Hill International Edition. ISBN: 0-07-2231265-3. Chapter 5.
- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 7.
- Ballou, R.H. Business Logistics Management. ISBN: 978-0130661845. Chapter 9.
- Muckstadt, John and Amar Sapra "Models and Solutions in Inventory Management"., 2006, Springer New York, New York, NY. Chapter 9, 10

For Inventory Models with Multiple Items and Locations

- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 7 & 8.

For Inventory Models for Class A & Class C Items

- Cachon, Gérard, and Christian Terwiesch. *Matching Supply with Demand: An Introduction to Operations Management*. Boston, MA: McGraw-Hill/Irwin, 2005.
- Silver, E.A., Pyke, D.F., Peterson, R. Inventory Management and Production Planning and Scheduling. ISBN: 978-0471119470. Chapter 8 & 9.

Fundamentals of Freight Transportation

Summary

The fundamentals of freight transportation provides and overview of different modes of transportation and some different ways to make decisions of the mode choice, analyzing the trade-offs between cost and level of service.

There are different levels of transportation networks (from strategic to physical). Physical network represents how the product physically moves, the actual path from origin to destination. Costs and distances calculations are made based on this level. Decisions from nodes (decision points) and arcs (a specific mode) are made in the Operational network. The third network, the strategic or service network, represents individual paths from end-to-end, and those decisions that tie into the inventory policies are made in the Strategic or Service network level.

Freight transportation also includes the important component of packaging. The Primary packaging, has direct contact with the product and is usually the smallest unit of distribution (e.g. a bottle of wine, a can, etc.). The Secondary packaging contains product and also a middle layer of packaging that is outside the primary packaging, mainly to group primary packages together (e.g. a box with 12 bottle of wines, cases, cartons, etc.). The Tertiary packaging is designed thinking more on transport shipping, warehouse storage and bulk handling (e.g. pallets, containers, etc.).

Key Concepts

Trade-offs between Cost and Level of Service (LOS):

- Provides path view of the Network
- Summarizes the movement in common financial and performance terms
- Used for selecting one option from many by making trade-offs

Packaging

- Level of packaging mirrors handling needs
- Pallets—standard size of 48 x 40 in in the USA (120 x 80 cm in Europe)
- Shipping Containers
    - TEU (20 ft) 33 m^3 volume with 24.8 kkg total payload
    - FEU (40 ft) 67 m^3 volume with 28.8 kkg total payload
    - 53 ft long (Domestic US) 111 m^3 volume with 20.5 kkg total payload

Transportation Networks

- Physical Network: The actual path that the product takes from origin to destination including guide ways, terminals and controls. Basis for all costs and distance calculations – typically only found once.
- Operational Network: The route the shipment takes in terms of decision points. Each arc is a specific mode with costs, distance, etc. Each node is a decision point. The four primary components are loading/unloading, local-routing, line-haul, and sorting.
- Strategic Network: A series of paths through the network from origin to destination. Each represents a complete option and has end-to-end cost, distance, and service characteristics.

Notation

TL: Truckload
TEU: Twenty Foot Equivalent (cargo container)
FEU: Forty Foot Equivalent (cargo container)

Lead Time Variability & Mode Selection

Variability in transit time impacts the total cost equation for inventory. There are important linkages between transportation reliability, forecast accuracy, and inventory levels. Mode selection is heavily influenced not only by the value of the product being transported, but also the expected and variability of the lead-time.

Impact on Inventory

Transportation affects total cost via

- Cost of transportation (fixed, variable, or some combination)
- Lead time (expected value as well as variability)
- Capacity restrictions (as they limit optimal order size)
- Miscellaneous factors (such as material restrictions or perishability)

Transportation Cost Functions

Transportation costs can take many different forms, to include:

- Pure variable cost / unit
- Pure fixed cost / shipment
- Mixed variable & fixed cost
- Variable cost / unit with a minimum quantity
- Incremental discounts

Lead/Transit Time Reliability

There are two different dimensions of reliability that do not always match:

- Credibility (reserve slots are agreed, stop at all ports, load all containers, etc.)

- Schedule consistency (actual vs. quoted performance)

Contract reliability in procurement and operations do not always match as they are typically performed by different parts of an organization. Contract reliability differs dramatically across different route segments (origin port dwell vs. port-to-port transit time vs. destination port dwell for instance). For most shippers, the most transit variability occurs in the origin inland transportation legs and at the ports.

Mode Selection

Transportation modes have specific niches and perform better than other modes in certain situations. Also, in many cases, there are only one or two feasible options between modes.

Criteria for Feasibility

- Geography
    - Global: Air versus Ocean (trucks cannot cross oceans!)
    - Surface: Trucking (TL, LTL, parcel) vs. Rail vs. Intermodal vs. Barge
- Required speed
    - >500 miles in 1 day—Air
    - <500 miles in 1 day—TL
- Shipment size (weight/density/cube, etc.)
    - High weight, cube items cannot be moved by air
    - Large oversized shipments might be restricted to rail or barge
- Other restrictions
    - Nuclear or hazardous materials (HazMat)
    - Product characteristics

Trade-offs within the set of feasible choices

Once all feasible modes (or separate carrier firms) have been identified, the selection within this feasible set is made as a trade-off between costs. It is important to translate the "non-cost" elements into costs via the total cost equation. The typical non-cost elements are:

- Time (mean transit time, variability of transit time, frequency)
- Capacity
- Loss and Damage

Notation

- c_i: Purchase cost for item i ($/unit)
- c_t: Ordering Costs ($/order)
- c_e: Excess holding Costs ($/unit/time); Equal to ch
- c_s: Shortage costs ($/unit)
- D: Average Demand (units/time)
- h: Carrying or holding cost ($/inventory $/time)
- Q: Replenishment Order Quantity (units/order)
- T: Order Cycle Time (time/order)
- μ_D: Expected Demand (Items) during One Time Period
- σ_D: Standard Deviation of Demand (Items) during One Time Period
- μ_L: Expected Number of Time Periods for Lead Time (Unitless Multiplier)
- σ_L: Standard Deviation of Time Periods for Lead Time (Unitless Multiplier)
- μ_{DL}: Expected Demand (Items) over Lead Time
- σ_{DL}: Standard Deviation of Demand (Items) over Lead Time
- N: Random Variable Assuming Positive Integer Values (1, 2, 3…)
- x_i: Independent Random Variables such that $E[x_i] = E[X]$
- S: Sum of x_i from i = 1 to N

Formulas

Random Sums of Random Variables

$$E[S] = E\left[\sum_{i=1}^{N} X_i\right] = E[N]E[X]$$

$$Var[S] = Var\left[\sum_{i=1}^{N} X_i\right] = E[N]Var[X] + (E[X])^2 Var[N]$$

Lead Time Variability

$$\mu_{DL} = \mu_L \mu_D$$

$$\sigma_{DL} = \sqrt{\mu_L \sigma_D^2 + (\mu_D)^2 \sigma_L^2}$$

Learning Objectives

- Understand common terminology and concepts of global freight transportation.
- Understanding of physical, operational, and strategic networks.
- Ability to select mode by trading off Level of Service (LOS) and cost.
- Understand the impact of transportation on cycle, safety, and pipeline stock.
- Understand how the variability of transportation transit time impacts inventory
- Able to use continuous approximation to make quick estimates of costs using a minimal amount of data.

References

- Ballou, Ronald H., Business Logistics: Supply Chain Management, 3rd edition, Pearson Prentice Hall, 2003. Chapter 6.
- Chopra, Sunil and Peter Meindl, Supply Chain Management, Strategy, Planning, and Operation, 5th edition, Pearson Prentice Hall, 2013. Chapter 14.