Inventory Strategies and Management

Grover V. Jeung S.

1 Inventory Functionality

From an inventory perspective, the ideal situation would be response capability to manufacture products to customer specification. While a zero-inventory manufacturing/ distribution system is typically not attainable, it is important to remember that each dollar invested in inventory is a trade-off with an alternative use of assets that may provide a better return.

Inventory is a major asset that should provide return for the capital invested. The return on inventory investment is the marginal profit on sales that would not occur without inventory. Accounting experts have long recognized that measuring the true cost and benefits of inventory on the corporate profit-and-loss is difficult. Lack of measurement sophistication makes it difficult to evaluate the trade-off among service levels, operating efficiencies, and inventory levels. While aggregate inventory levels have decreased, many enterprises still carry an average inventory that exceeds their basic requirements. This generalization can be understood better through a review of the four prime functions of inventory.

These four functions-geographical specialization, decoupling, balancing supply and demand, and buffering uncertainty-require inventory investment to achieve managerial operating objectives. While lean logistics has made significant progress in reducing overall supply chain inventory, inventory properly deployed can create value and reduce total cost. Given a specific manufacturing/marketing strategy, inventories planned and committed to operations can only be reduced to a level consistent with performing the four inventory functions. All inventories exceeding the minimum level are excess commitments.

At the minimum level, inventory invested to achieve geographical specialization and decoupling can only be modified by changes in facility location and operational process of the enterprise. The minimum level of inventory required to balance supply and demand depends on the difficult task of estimating seasonal requirements. With accumulated experience over a number of seasonal periods, the inventory required to achieve marginal sales during periods of high demand can be projected fairly well. A seasonal inventory plan can be formulated based upon this experience.

Inventories committed to safety stocks represent the greatest potential for improved logistical performance. These commitments are operational in nature and can be adjusted rapidly in the event of an error or policy change. A variety of techniques are available to assist management in planning safety stock commitments.

2 Inventory Cost

Before we can suggest actions that a manager can take to reduce cycle inventory, it is important to understand why stages of a supply chain produce or purchase in large lots and how lot size reduction affects supply chain perform-

ance.

Cycle inventory is primarily held to take advantage of economies of scale and reduce cost within the supply chain. Increasing the lot size or cycle inventory often decreases the cost incurred by different stages of a supply chain. To understand how the supply chain achieves these economies of scale, we must first identify supply chain costs that are influenced by the lot size.

The average price paid per unit purchased is a key cost in the lot sizing decision. A buyer may increase the lot size if this action results in a reduction in the price paid per unit purchased. For example, if the jean manufacturer charges $20 per pair of jeans for orders under 500 pairs and $18 per pair of jeans for larger orders, the store manager at Jean-Mart may order in lots of at least 500 pairs to get the lower price. The price paid per unit is referred to as the material cost and is denoted by C. It is measured in $/unit. In many practical situations, material cost displays economies of scale and increasing the lot size decreases the material cost.

The fixed ordering cost includes all costs that do not vary with the size of the order but are incurred each time an order is placed. For example, a fixed administrative cost may be incurred to place an order, a trucking cost may be incurred to transport the order, and a labor cost incurred to receive the order. Jean-Mart for example, incurs a cost of $400 for the truck that brings the shipment of jeans from the manufacturer. Assume that the truck can hold up to 2,000 pairs of jeans. The $400 cost is incurred irrespective of the number of jeans shipped on the truck. A lot size of 100 pairs of jeans results in a transportation cost of $4 per pair, whereas a lot size of 1,000 pairs results in a transportation cost of $0.40 per pair. Given the fixed transportation cost per batch, the store manager is inclined to increase to size to reduce transportation cost per pair of jeans. The fixed ordering cost per lot or batch is denoted by s and is measured in $/lot. The ordering cost also displays economies of scale, and increasing the lot size decreases the fixed ordering cost per unit purchased.

Holding cost is the cost of carrying one unit in inventory for a specified period to time, usually one year. It is a combination of the cost of capital, the cost of physically storing the inventory, and the cost that results from the product becoming obsolete. The holding cost is denoted by H and is measured in $/unit/year. It may also be obtained as a fraction, where h is the cost of holding $1 in inventory for one year. Given a unit cost of C, the holding cost H is given by

$$H=hC$$

The total holding cost increases with an increase in lot size and cycle inventory.

To summarize, the costs that must be considered in any lot sizing decision are as follows:

- Average price per unit purchased, $C/unit
- Fixed ordering cost incurred per lot, $S/lot
- Holding cost incurred per unit per year, $/z/unit/year = hC

The primary role of cycle inventory is to allow different stages in the supply chain to purchase product in lot sizes that minimize the sum of the material, ordering, and holding cost. If a manager was considering the holding cost alone, he or she would reduce the lot size and cycle inventory. However, economies of scale in purchasing or ordering motivate a manager to increase the lot size and cycle inventory. A manager must make the trade-off that minimizes the total cost when making the lot sizing decision.

Ideally, cycle inventory decisions should be made considering the total cost across the entire supply chain. However, in practice, each stage makes its cycle inventory decisions independently.

3 Planning Inventory

Safety inventory is inventory carried for the purpose of satisfying demand that exceeds the

amount forecasted for a given period. Safety inventory is carried because demand forecasts are uncertain and a product shortage may result if actual demand exceeds the forecast demand. Consider, for example, Bloomingdales, a high-end department store. Bloomingdales sells purses for women and purchases them from Gucci, an Italian manufacturer. Given the high transportation cost from Italy, the store manager at Bloomingdales orders in lots of 600 purses. Demand for purses at Bloomingdales averages 100 per week. Gucci takes three weeks to respond to a Bloomingdales order and deliver the purses to the store. If there is no demand uncertainty and exactly 100 purses are sold each week, the store manager at Bloomingdales can place an order when the store has exactly 300 purses remaining. In the absence of demand uncertainty, such a policy ensures that the new lot arrives just as the last purse is being sold at the store.

However, demand forecasts are unlikely to be completely accurate. Given forecast errors, actual demand over the three weeks may be higher or lower than the 300 purses forecasted. If the actual demand at Bloomingdales is higher than 300, some customers will be unable to purchase purses, resulting in a loss of potential margin for Bloomingdales. The store manager thus decides to place an order with Gucci when the store still has 400 purses. This policy allows the store manager to improve product availability for the customer, because the store now runs out of purses only if the demand over the three weeks exceeds 400. Given an average weekly demand of 100 purses, the store will have an average of 100 purses remaining when the replenishment lot arrives. Safety inventory is the average inventory remaining when the replenishment lot arrives. Thus, Bloomingdales carries a safety inventory of 100 purses.

The appropriate level of safety inventory is determined by the following two factors:

- The uncertainty of demand or supply
- The desired level of product availability

As the uncertainty of supply or demand grows, the required level of safety inventories increases. Consider the sale of Palm Pilots at B&M Office Supplies, a distributor of Palm Pilots. When a new Palm Pilot model is introduced, demand is highly uncertain. B&M thus carries a much higher level of safety inventory relative to demand. As the market accepts the model, demand stabilizes and uncertainty reduces. At that point, B&M can carry a lower level of safety inventory relative to demand.

Next we discuss some measures of demand uncertainty.

Measuring Demand Uncertainty

Demand has a systematic as well as a random component. The goal of forecasting is to predict the systematic component and estimate the random component. The estimate of the random component is a measure of demand uncertainty. The random component is usually estimated as the standard deviation demand. We assume the following inputs for demand:

R: Average demand per period

O: Standard deviation of demand per period

For now, we assume that demand is normally distributed. In the case of B&M, weekly demand for the Palm Pilot is normally distributed with a mean of R and a standard deviation of O.

Measuring Product Availability

Product availability reflects a firm's ability to fill a customer order out of available inventory. A stock out results if a customer order arrives when product is not available. There are several ways to measure product availability. All availability measures are defined on average over a given time frame, which can range from hours to a year. Some of the important measures follow:

1. Product fill rate (fr) is the fraction of product demand that is satisfied from product in inventory. It is equivalent to the probability that product demand is supplied from available inventory. Assume that B&M provides Palm Pilots to 90 percent of its customer from inventory, with the remaining 10 percent lost to a neighboring

competitor because of a lack of available inventory. In this case, B&M achieves a fill rate of 90 percent.

2. Order fill rate is the fraction of orders that are filled from available inventory. In a multi-product scenario, an order is filled from inventory only if all products in the order can be supplied from the available inventory. In the case of B&M, a customer may order a Palm Pilot along with a calculator. The order is filled from inventory only if both the Palm Pilot and the calculator are available through the store. Order fill rates tend to be lower than product fill rates.

3. Cycle service level (CSL) is the fraction of replenishment cycles that end with all the customer demand being met. A replenishment cycle is the interval between two successive replenishment deliveries. The cycle service level is equal to the probability of not having a stock out in replenishment cycle. If B&M orders replenishment lots of 600 Palm Pilots, the interval between the arrivals of two successive replenishment lots is a replenishment cycle. If the manager at B&M manages inventory such that the store does not run out of inventory in 6 out of 10 replenishment cycles, the store achieves a cycle service level of 60 percent. Observe that a cycle service level of 60 percent will typically result in much higher fill rate. In the 60 percent of cycles in which B&M does not run out of inventory, all customer demand is satisfied from available inventory. In the 40 percent of cycles in which a stock out does occur most of the customer demand is satisfied from inventory. Only the small fraction toward the end of the cycle that arrives after B&M is out of inventory is lost. As a result, the fill rate will be much higher than 60 percent.

The distinction between product fill rate and order fill rate is not very significant in a single product situation. However, when a firm is selling multiple products, this difference may be significant. For example, if most orders include 10 or more different products that are to be shipped, an out-of-stock situation of one product results in the order not being filled from stock. The firm in this case may have a poor order fill rate even though it has good product fill rates. Tracking order fill rates is important when customers place a high value on the entire order being filled simultaneously.

4 Inventory Management Practices

In this section, we raise some issues that must be considered when putting into practice the ideas discussed:

Account for the fact that supply chain demand is lumpy. In practice, a manufacturer or distributor does not order one unit at a time but instead often orders in a large lot. Thus, demand observed by different stages of the supply chain tends to be lumpy. Lumpiness adds to the variability of demand, Lumpiness of demand is not a significant issue when using periodic review policies. However, when using a continuous review policy, it may lead to inventory dropping far below the reorder point before a replenishment order is placed. On average, inventory will drop below the reorder point by about half the average size of an order. The lumpiness can be accounted for in practice by raising the safety inventory suggested by the models discussed earlier by half the average size of an order.

Adjust inventory policies if demand is seasonal. In practice, demand is often seasonal, with the mean and the standard deviation of demand varying by the time of year. Thus, a given reorder point or order up to level may correspond to 10 days of demand during the low-demand season and only 2 day of demand during the peak demand season. If the lead time is a week, stock outs are certain to occur during the peak season. In the presence of seasonality, it is not appropriate to select an average demand and standard deviation over the year to evaluate fixed reorder points and order up to levels. Both the mean and the standard deviation of demand must be adjusted by the time of year to reflect changing demand. Corresponding adjustments in the reorder points, order up to levels, and safety inventories must be made over the year. Adjustments for

changes in the mean demand over the year are generally more significant than adjustments for changes in variability.

Use simulation to test inventory policies. Given that demand may not be perfectly normal and may be seasonal, it is a good idea to test and adjust inventory policies using a computer simulation before they are implemented. The simulation should use a demand pattern that truly reflects actual demand, including any lumpiness as well as seasonality. The inventory policies obtained using some models can then be tested and adjusted if needed to obtain the desired service levels. Very powerful simulations can be built using Excel or other simulation packages. Identifying problems in a simulation can save a lot of time and money compared with facing these problems once the inventory policy is in place.

Start with a pilot. Even a simulation cannot identify all problems that may arise when using an inventory policy. Once an inventory policy has been selected and tested using simulation, it is often a good idea to start implementation with a pilot program of products representative of the entire set of products in inventory. By starting with a pilot, many of the problems (both in the inventory policies themselves and in the process of applying the policies) can be solved. Getting these problems solved before the policy is rolled out to all the products can save a lot of time and money.

Monitor service levels. Once an inventory policy has been implemented, it is important that its performance be tracked and monitored. Monitoring is crucial because it allows a supply chain to identify when a policy is not working well and make adjustments before supply chain performance is significantly affected. Monitoring requires not just tracking the inventory levels but also tracking any stock outs that may result. Historically, firms have not tracked stock outs very well, partly because they are difficult to track, and partly because there is the perception that they affect the customer and not the firm itself. Stock outs can be difficult to measure in a situation like a supermarket, where the customer simply does not buy the product when it is not on the shelf. However, there are simple ways to estimate stock outs. At a supermarket, the fraction of time that a shelf does not contain a product may be used to estimate the fill rate. Stock outs are in fact easier to estimate on-line, where the number of clicks on an out-of-stock product can be measured. Given the fraction of clicks that turn into orders and the average size of an order, demand during a stock out can be estimated.

Focus on reducing safety inventories. Given that safety inventory is often a large fraction of the total inventory in a supply chain, the ability to reduce safety inventory without hurting product availability can significantly increase supply chain profitability. This is particularly important in the high-tech industry, where product life cycles are short.

Regents Publishing Company, Inc., 2005