BETTER PLANNING AND SCHEDULING WITH THE RIGHT CYCLE STOCK LEVELS

By Peter L. King, CSCP, and Courtney Bigler

Editor's Note: Peter L. King, CSCP, is the author of the 2011 article "Crack the Code," which is still one of our digital magazine's most popular stories. The following article is written as a response to the many questions readers have asked about the original article and to help take the topic of mastering inventory management a step further.

Setting the right inventory levels is always a challenge. If you carry too much inventory, then there's also too much money tied up in working capital and too much physical space dominated by unnecessary inventory. However, if there's not enough inventory, stockouts are inevitable.

Complicating the challenge even further is the need for two broad types of inventory: cycle stock and safety stock. Cycle stock is the amount of a specific product that is made during a production period to satisfy demand over the full production cycle, including the portion of the cycle when other products or processes are utilizing the asset. For example, if a production process is based on a total production cycle of seven days, the amount of cycle stock for material A must last seven days. If material A is only made on the first day of the production cycle, then the amount of cycle stock for that material needs to cover the amount that was consumed while the material was being produced as well as the six remaining days of the cycle until production resumes.

Safety stock is material held to satisfy demand in cases when actual demand is higher than expected or when the next cycle was late in starting. Figure 1 shows a profile of inventory levels for a single stock keeping unit (SKU) throughout three production cycles. In this case, both cycle stock and safety stock are present. Production period P1 raises the cycle stock level to A, the average demand level. Demand during the next demand cycle, D1, is equal to the average demand, so the cycle stock is consumed, but the safety stock is not. Production period P2 raises total inventory back to level A. Demand during cycle D2 is greater than average, so some safety stock is used. Alternatively, if P2 took longer than usual because of some interruption, safety stock would be needed to fulfill demand. Thus, safety stock protects flow against variations in both demand and supply lead time. Because safety stock is used in D2, more material than average must be produced in P3 to replenish both the cycle stock and the safety stock.

CALCULATING CYCLE STOCK

But how do you figure out the average demand? Typically, this number is based on demand history or a forecast. If previous demand is considered to be the best predictor of future demand, then the cycle stock level should match the demand history. If there is a forecast that is believed to be a more accurate indication of future demand — or if this is a totally new product — the amount of cycle stock should be based on the forecast. Because forecasts can vary by period, the cycle stock amount may be increased or decreased each period in accordance with the forecast.

Cycle stock can be replenished on a fixed-interval or on a fixed-order-quantity basis. As the name implies, fixed-interval replenishments occur on a regular schedule, although the quantity replaced may vary widely depending on the amount of material consumed during the most recent cycle. This method of replenishment is also known as a fixed-order-interval model, a fixed-reorder-cycle inventory model, a periodic review system and a time-based system.

Fixed order quantity replenishment behaves the opposite way. The quantity replenished is based on some specific criteria and does not vary. Instead, the replacement interval varies based on the rate of consumption since the last replenishment. This model also is known as a continuous review model, a reorder point model, a lot-size system and a quantity-based order system.

REGULAR REPLENISHMENT

Figure 1 shows the inventory profile for a single material in a fixed-interval strategy. This means that enough stock must be produced to last until the next production cycle, which, in this, case happens every 14 days. The amount of cycle stock will be equal to the average demand during the 14-day period, with safety stock making up the difference between average demand and peak demand. The amount produced during the next production cycle is based on material consumption during the previous demand cycle.

The standard equations governing this model are as follows:

$$\text{Peak inventory} = \text{Cycle stock} + \text{Safety stock}$$

$$\text{Average inventory} = \tfrac{1}{2}\,(\text{Cycle stock}) + \text{Safety stock}$$

The amount of inventory for any SKU peaks at the cycle stock plus safety stock level. Then over the course of the cycle, the inventory level drops to somewhere around the safety stock level. Therefore, the average inventory level over the cycle is safety stock plus half the cycle stock.

These equations are accurate for purchased materials or materials received as a complete lot equal to the cycle stock amount. They are approximations when applied to materials being produced within a process because some of the cycle stock is being consumed by downstream steps as it is being produced. This has a minor effect on products that occupy a small portion of the production cycle. However, the effect can be significant if a product occupies a large portion of the cycle.

The following equations apply to these situations:

$$\text{Peak inventory} = \text{Cycle stock}\left(1-\frac{D}{PR}\right) + \text{Safety stock}$$

$$\text{Average inventory} = \tfrac{1}{2}\,(\text{Cycle stock})\left(1-\frac{D}{PR}\right) + \text{Safety stock}$$

D is the demand for that material per unit of time, and PR is the production rate — or the total quantity produced during that same time. It is critical that both factors be in the same time units, whether hours, days, weeks or another time interval.

LEAD TIME CONSIDERATIONS

When a company needs to order raw materials, there generally will be a lead time before the material is received. In that case, the inventory profile will look more like the model shown in Figure 2. When the normal order interval begins at point A, enough material must be ordered to not only replenish cycle stock and safety stock to the target levels but also to cover demand during the lead time (DDLT). Thus, the amount to be ordered at point A is represented by the following equation:

$$\text{Order quantity} = \text{DDLT} + \text{Cycle stock} + \text{Safety stock} - \text{Current inventory}$$

The current inventory typically will be approximately DDLT plus safety stock, so the amount ordered will be approximately the cycle stock.

If the DDLT is greater than average, as shown in lead time C–D in Figure 2, safety stock will prevent a stockout. But when the new order arrives, the order quantity will not bring total inventory up to the cycle stock plus safety stock target. If safety stock has been calculated appropriately, that shortfall will be covered.

BUDGETING IN BULK

Fixed-quantity replenishment can be used when there is a benefit to buying or producing

FIGURE 1 | Cycle stock and safety stock levels throughout three production cycles

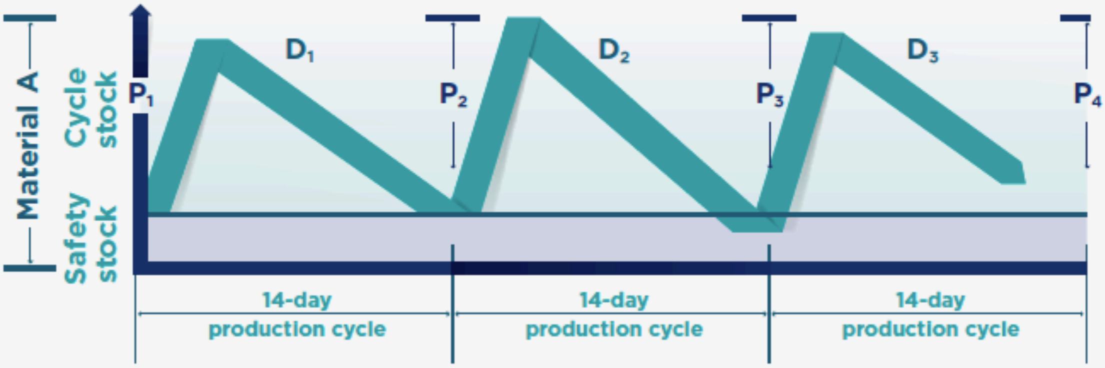

BETTER PLANNING AND SCHEDULING WITH THE RIGHT CYCLE STOCK LEVELS

DEALING WITH SEASONAL STOCK

If demand for a finished product exhibits predictable seasonal trends, cycle stock should be varied for the different seasons. Cycle stock of raw materials also should reflect the seasonal trends.

If demand for some seasonal products typically is very low in the off-season, it may make sense to put those products on a make-to-order strategy during those periods, if manufacturing cycle time is low enough to allow it.

materials in specific quantities. For example, in process industries, some materials are received in tank trucks, so transportation economics advises buying in truck quantities. Similarly, suppliers of cardboard packaging materials might require a minimum order quantity for custom print orders for branded packaging. In other cases, an economic order quantity calculation, which balances ordering costs with inventory carrying costs, will be used to optimize order quantity.

In the production process, there often is a specific campaign size that best balances changeover cost with inventory carrying cost, as determined by an economic production quantity calculation. That quantity would be used to replenish finished product inventory on a fixed-quantity basis.

An inventory profile for a single material replenished using a fixed-quantity model is illustrated in Figure 3. Because the order quantity Q already is known, the variable is when the next order should be placed. This usually depends on when inventory falls to or below the order point, which is when total inventory reaches the safety stock plus DDLT level. As shown in Figure 3, order interval B-C is slightly shorter than interval A-B because demand in D2 is greater than demand in D1.

In a perfect world in which no safety stock is needed, the new order would arrive just before a stockout would occur. However, in the real world, orders may be placed late, deliveries can be delayed, or DDLT could be higher than average. Safety stock is required to cover these and other scenarios. To account for this, the order point is set so that, on average, the new order will arrive just as inventory falls to the safety stock level.

This model is often called continuous review because it assumes that the inventory level is being continuously monitored and that the new order is placed immediately when the inventory falls to the reorder point. However, in some situations, inventory is not monitored continuously and instead is checked daily, weekly or even less often. Some enterprise resources planning systems only check inventory once per 24 hours. These review periods — or time between inventory checks — must be factored into the order point to cover the demand between review periods (DDRP) and prevent stockouts.

In this case, this is the order point calculation:

$$\text{Order point} = \text{DDRP} + \text{DDLT} + \text{Safety stock}$$

If the lead time for replenishment tends to be long, the order point will be very large, which can seem shocking at first. However, in this case, current inventory also can include

FIGURE 2 | Fixed-interval replenishment with lead time

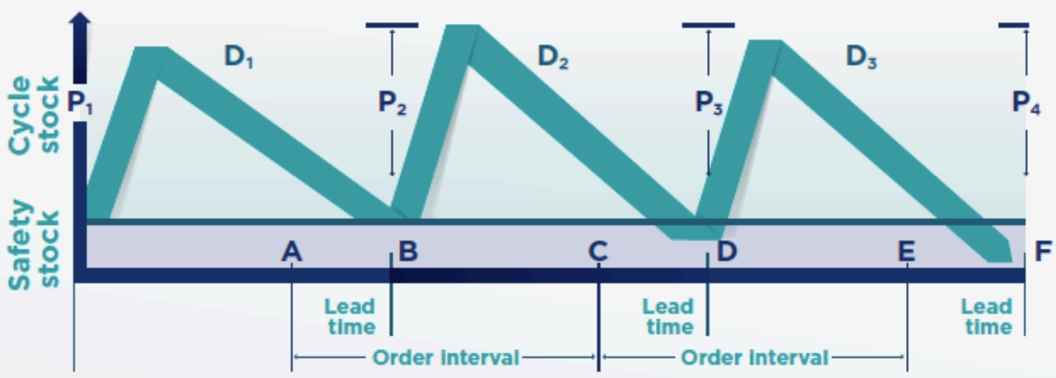

FIGURE 3 | Inventory profile with fixed-quantity replenishment

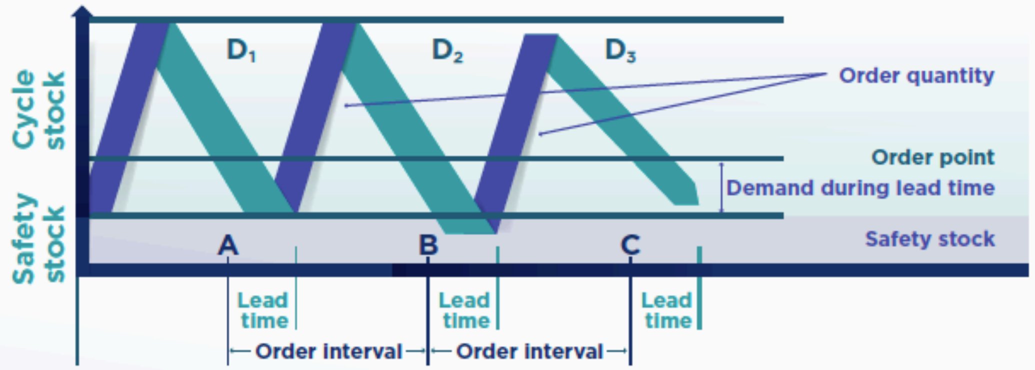

the inventory ordered but not yet received. So, an order actually gets placed when:

$$\text{Inventory on hand} + \text{Inventory in transit} + \text{Inventory ordered but not yet shipped} < \text{DDRP} + \text{DDLT} + \text{Safety stock}$$

In this replenishment model, the cycle stock amount is the order quantity, and peak inventory and average inventory are calculated in the same way as in the fixed-interval model. Also, as in the fixed-interval model, these average and peak inventory equations must be adjusted when a substantial portion of the cycle stock is consumed during production to factor in that consumption.

FINDING THE BEST FIT

Another inventory challenge is determining which strategy to use when. The fixed-quantity model generally requires less safety stock than the fixed-interval model when applied to the same situation. The fixed-quantity model only needs safety stock protection during the lead time; whereas, the fixed-interval model requires safety stock protection during lead time and the interval duration. However, if the fixed-quantity model has a review period, then additional protection will be needed during that period. Thus, the safety stock advantage diminishes as the length of the fixed-quantity review period approaches the fixed-interval duration.

Given this, the fixed-quantity model is most useful when

- the company is handling high-value materials
- there is an incentive to buy, produce or ship in specific quantities
- a system is in place to continuously or frequently monitor inventory levels.

If it is very difficult or costly to obtain frequent inventory-level updates, then a fixed-interval process may be preferable. This model's structure and predictability also enable better planning and scheduling of support activities, such as preventive maintenance or quality control checks.

Now that we've explained cycle stock and how to calculate the appropriate levels, it's time to turn our attention to the second inventory component, safety stock. This will be examined in an upcoming SCM Now magazine article.

Peter L. King, CSCP, is president of Lean Dynamics LLC. Prior to this, he spent 40 years with DuPont in a variety of manufacturing automation, project management and lean continuous improvement programs. He is also the author of several books about lean. King may be contacted at peterking@leandynamics.us.

Courtney Bigler is a senior demand planning analyst at The Boston Beer Company and a part-time business consultant for Lean Dynamics LLC. She may be contacted at courtneybigler@leandynamics.us.