

Determining appropriate inventory levels is one of the most important and most challenging tasks faced by operations managers. If you carry too much inventory, you tie up money in working capital; if you don't carry enough inventory, you face stockouts. Fortunately, the cycle stock portion of the inventory equation is straightforward. What keeps people up at night is safety stock.

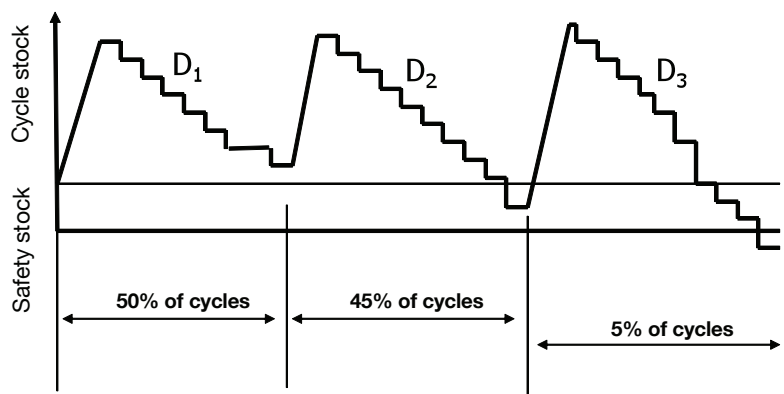
# CRACK THE CODE

## Understanding safety stock and mastering its equations

By Peter L. King, CSCP

Safety stock simply is inventory that is carried to prevent stockouts. Stockouts stem from factors such as fluctuating customer demand, forecast inaccuracy, and variability in lead times for raw materials or manufacturing. Some operations managers use gut feelings or hunches to set safety stock levels, while others base them on a portion of cycle stock level—10 or 20 percent, for example. While easy to execute, such techniques generally result in poor performance. A sound, mathematical approach to safety stock will not only justify the required inventory levels to business leaders, but also balance the conflicting goals of maximizing customer service and minimizing inventory cost.

Safety stock determinations are not intended to eliminate all stockouts—just the majority of them. For example, when designing for a 95 percent service level, expect that 50 percent of the time, not all cycle stock will be depleted and safety stock will not be needed. For another 45 percent of cycles, the safety stock will suffice. But in approximately 5 percent of replenishment cycles, expect a stockout. (See Figure 1.)



While designing for a higher service level—say, 98 percent—would result in fewer stockouts, this requires significantly more safety stock. There must be a balance between inventory costs and customer service. By using the methods and equations that follow, you can find safety stock levels to achieve your desired customer service levels.

### Variability in demand

Imagine the only variability you need to protect against is demand variability, and good historical data are available. The safety stock needed to give a certain level of protection simply is the standard deviation of demand variability multiplied by the Z-score—a statistical figure also known as standard score. For example, to satisfy demand with a 95 percent confidence level, according to statistical analysis, it's necessary to carry extra

— FIGURE 1: Inventory designed for a 95 percent service level —

inventory equal to 1.65 standard deviations of demand variability. This is equivalent to a Z-score of 1.65.

To further understand Z-score, imagine that no safety stock is carried. In this situation, the Z-score is zero. Even so, there will be enough inventory to meet demand in 50 percent of cycles. If Z-score equals 1, the safety stock will protect against one standard deviation; there will be enough inventory 84 percent of the time. This percentage of cycles where safety stock prevents stockouts is called the cycle service level.

Figure 2 shows the relationship of desired cycle service levels to Z-score. As illustrated, the relationship is nonlinear: Higher cycle service levels require disproportionately higher Z-scores and, thus, disproportionately higher safety stock levels. Typical goals fall between 90 and 98 percent, and—statistically speaking—a cycle service level of 100 percent is unattainable.

Rather than using a fixed Z-score for all products, set the Z-score independently for groups of products based on criteria such as strategic importance, profit margin, or dollar volume. Then, those stockkeeping

## Safety stock determinations are not intended to eliminate all stockouts—just the majority of them.

units with greater value to the business will have more safety stock, and vice versa.

So far, it has been assumed that the demand periods equal total lead time, including any review period. When this is not the case, instead, calculate standard deviation based on periods equal to the lead time. For example, if the standard deviation of demand is calculated from weekly demand data and the total lead time including review period is three weeks, the standard deviation of demand is the weekly standard deviation times the square root of the ratio of the time units, or  $\sqrt{3}$ .

Taking all this into consideration, the safety stock equation becomes:

$$\text{Safety stock} = Z \times \sqrt{\frac{PC}{T_1}} \times \sigma_D, \text{ where:}$$

Z = Z-score

PC = performance cycle, another term for total lead time

$T_1$  = time increment used for calculating standard deviation of demand

$\sigma_D$  = standard deviation of demand.

The performance cycle includes the time needed to perform functions such as deciding what to order or produce, communicating orders to the supplier, manufacturing and processing, and delivery and storage, as well as any additional time required to return to the start of the next cycle.

### Variability in lead time

In the previous equation, safety stock is used to mitigate demand variability. However, when variability in lead time is the primary concern, the safety stock equation becomes:

$$\text{Safety stock} = Z \times \sigma_{LT} \times D_{avg}, \text{ where:}$$

$\sigma_{LT}$  = standard deviation of lead time

$D_{avg}$  = average demand.

When both demand variability and lead time variability are present, statistical calculations can combine to give a lower total safety stock than the sum of the two individual calculations. In cases where demand and lead time variability are independent—that is, they are influenced by different factors—and both are normally distributed, the combined safety stock equation becomes:

$$\text{Safety stock} = Z \times \sqrt{\left(\frac{PC}{T_1} \times \sigma_D^2\right) + \left(\sigma_{LT} \times D_{avg}\right)^2}$$

In other words, the safety stock is Z-score times the square root of the sum of the squares of the individual variabilities.

But when demand and lead time variability are not independent of each other, this equation can't be used. In these cases, safety stock is the sum of the two individual calculations:

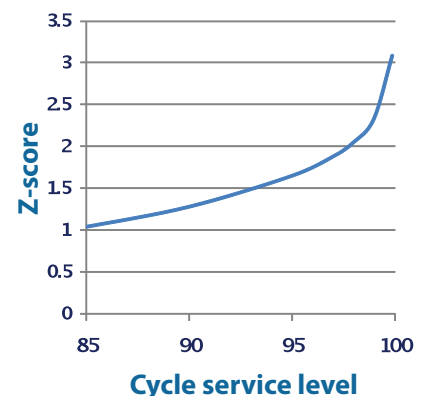
$$\text{Safety stock} = \left(Z \times \sqrt{\frac{PC}{T_1}} \times \sigma_D\right) + \left(Z \times \sigma_{LT} \times D_{avg}\right)$$

### Cycle service level and fill rate

The previous equations are useful for predicting the safety stock needed to attain a certain cycle service level—a percentage of replenishment cycles. Sometimes, business leaders instead wish to control the amount of volume ordered that is available to satisfy customer

demand—a quantity known as fill rate. Fill rate often is a better measure of inventory performance, as cycle service level merely indicates the frequency of stockouts without regard to their magnitudes. (See Figure 3.)

Desired cycle service level	Z-score
84	1
85	1.04
90	1.28
95	1.65
97	1.88
98	2.05
99	2.33
99.9	3.09



— FIGURE 2: Relationship between desired service level and Z-score —

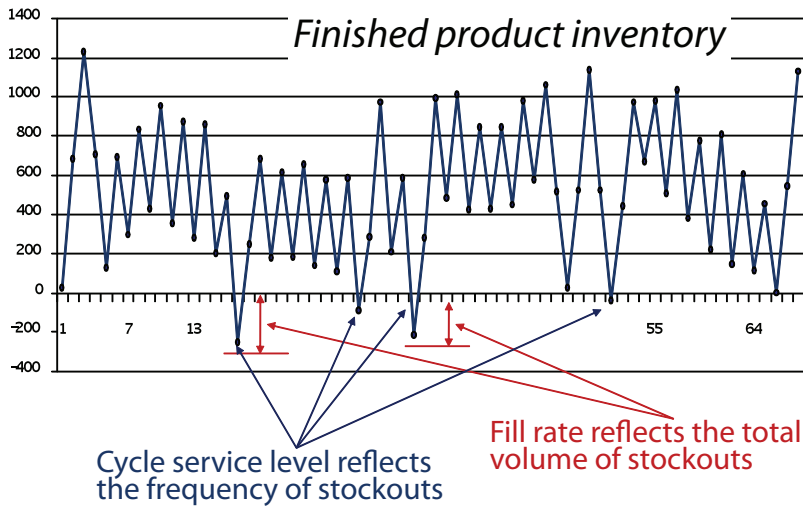


FIGURE 3: Cycle service level and fill rate

As illustrated in Figure 3, when supply and demand are relatively stable—that is, when standard deviations of demand and lead time are low—fill rate will tend to be higher than cycle service level. While stockouts still occur, the magnitude of stockouts tends to be small. Conversely, when demand or lead time variability is high, fill rate will be lower than cycle service level, and the volume of stockouts will be high.

To better understand these equations, consider the following example of a warehouse that holds large rolls of plastic film. The film gets sold to processors who cut it into shorter, narrower rolls for food packaging, pallet wrapping, or as a dielectric material in industrial capacitors.

The following are the relevant data for one unit type:

Weekly demand = 50 rolls

Standard deviation of weekly demand = 10 rolls

Standard deviation of lead time = 0

Production cycle = weekly

Production capacity = 496 rolls/week.

The lead time is very stable and predictable. Process reliability is high enough that the manufacturing lead time never exceeds seven days, and the lead time of transport from the manufacturing facility to the warehouse never exceeds one day. Because this gives a standard deviation of lead time of zero, the safety stock requirements can be calculated from the original equation:

$$\text{Safety stock} = Z \times \sqrt{\frac{PC}{T_1} \times \sigma_D}$$

The desired cycle service level is 95 percent; that is, the business can tolerate stockouts of this product on no more than 5 percent of the replenishment cycles, or slightly more than two per year. Using the chart in Figure 2, the Z-score is found to be 1.65. Performance cycle, which affects replenishment of the warehouse inventory, is the sum of the seven-day manufacturing time and the one day needed to arrive at the warehouse, for a total of eight days.  $T_1$ , the time increment used to calculate  $\sigma_D$ , is seven days. Thus:

$$\text{Safety stock} = 1.65 \times \sqrt{\frac{8}{7}} \times 10 \text{ rolls}$$

$$\text{Safety stock} = 18 \text{ rolls}$$

If lead time were variable, more safety stock would be required to meet performance goals, and the safety stock equation becomes:

$$\text{Safety stock} = Z \times \sqrt{(\frac{PC}{T_1} \times \sigma_D^2) + (\sigma_{LT} \times D_{avg})^2}$$

In this example, lead time varies with a standard deviation of half a day, or approximately 0.07 weeks. Thus:

$$\text{Safety stock} = 1.65 \times \sqrt{(8/7 \times 10^2) + (0.07 \times 50)^2}$$

$$\text{Safety stock} = 1.65 \times \sqrt{114.3 + 12.2}$$

$$\text{Safety stock} = 19 \text{ rolls}$$

These results demonstrate that demand variability is the dominant influence on safety stock requirements: Its effect is almost 10 times that of lead time variability. With the recognition of what factors dominate an equation, it becomes easier to focus improvement efforts. In this case, if a reduction in safety stock is desired, it is far more productive to reduce demand variability than lead time variability. Conversely, if a high level of customer service is not required, safety stock can be lowered to a more appropriate level.

Once safety stock levels have been established, inventory levels should be monitored on an ongoing basis to determine if the inventory profile is as expected. Is the safety stock being consumed in about half of the cycles? Are service level targets being realized? If not, before any adjustments are made, perform a root cause analysis to see if any special causes are responsible for the deviations from expected results.

### Safety stock alternatives

The previous calculations may result in safety stock recommendations higher than what business leaders feel they can carry. The good news is there are alternatives to mitigate variability other than safety stock.

## Demand variability is the dominant influence on safety stock requirements.

One is implementing an order-expediting process for preventing stockouts when safety stock is insufficient to cover all random variation. This practice is especially appropriate for products that cost more to produce (and thus cost more to carry in inventory).

In one example involving an expensive but relatively lightweight product, total supply chain costs were reduced significantly. The company carried small amounts of safety stock in overseas warehouses, and it relied on air freight to cover peaks in demand. The cost of shipping a small percentage of total demand via air was minimal compared to the cost of carrying large amounts of safety stock of the valuable material on an ongoing basis.

Another alternative to carrying safety stock is to consider a make-to-order (MTO) or finish-to-order (FTO) production environment. If lead times allow, MTO eliminates the need for most safety stock. Meanwhile, FTO allows for less differentiation in safety stock than finished-product inventory, which lowers demand variability and reduces safety stock requirements. FTO and MTO also are well suited for situations where customers are willing to accept longer lead times for highly sporadic purchases.

In the end, safety stock can be an effective way to mitigate demand uncertainty and lead time variability while still providing high service levels to customers. But before proceeding with a plan, first understand how to determine appropriate levels of safety stock—and what degree of protection they provide.

Peter L. King, CSCP, is founder and president of consulting firm Lean Dynamics LLC and author of *Lean for the Process Industries: Dealing with Complexity*. Previously, he was a principal consultant in the lean technology division of DuPont. He may be contacted at [peterking@leandynamics.us](mailto:peterking@leandynamics.us).

To comment on this article, send a message to [feedback@apics.org](mailto:feedback@apics.org).

## Network with APICS Online

### LinkedIn

Participate in the APICS LinkedIn group to network with nearly 18,000 of your peers in this active community. [apics.org/LinkedIn](http://apics.org/LinkedIn)

### Facebook

Connect to the APICS community on Facebook. Like us today! [apics.org/Facebook](http://apics.org/Facebook)

### APICS Blogs

Read and engage with the APICS International and the APICS Young Professionals blogs. [blogs.apics.org](http://blogs.apics.org)

### APICS Learning Communities

Connect with your peers and have your questions answered on the APICS Learning Communities. [apics.org/resources/learningcommunities](http://apics.org/resources/learningcommunities)

