SAFETY INVENTORY ANALYSIS

WHY AND HOW?

BY SCOTT W. HADLEY

WOULDN'T YOU LIKE TO KNOW how the stock market was going to behave tomorrow, whether or not there's going to be an accident on your way to work, and if your flight is going to be delayed—and for how long? It goes without saying—we live in a world fraught with uncertainty!

To deal with uncertainty, more often than not you build some buffer into your life. For instance, you keep extra money in the bank just in case the car needs major repairs. You plan to arrive at the airport 90 minutes early just in case there's a traffic jam, and you schedule your meeting for two hours after you are due to land, just in case.

Companies also buffer against uncertainty. Perhaps the most common buffer companies use is safety inventory—carrying extra inventory just in case demand exceeds the forecast, just in case manufacturing has a breakdown, and just in case a supplier delivery is late or short in quantity.

"A Modern View of Inventory" in the July 2004 issue of *Strategic Finance* outlined the various roles of inventory and described techniques for computing target inventory levels. This article zeroes in on actual computation techniques to determine safety inventory levels. I'll explain a new one that addresses inventory uncertainty and compare it with common calculation techniques. A sample data set will demonstrate that the new technique far outperforms other methods, so you can use it to drive business value by achieving your customer service goals with less inventory. Let's begin.

THE ROLE OF SAFETY INVENTORY

Safety inventory protects against inventory uncertainty by ensuring there is enough product available to maintain desired service levels. Many factors contribute to inventory uncertainty, but three main elements give rise to the differences between planned inventory and actual inventory levels: demand deviations, supply deviations, and inventory accuracy. Traditionally, the analysis of safety inventory levels has focused primarily on demand and demand uncertainty, while supply uncertainty and inventory accuracy aren't typically addressed explicitly in the analysis of safety inventory.

What causes the deviations? Actual demand differing from forecast demand is the most common source of demand deviations. Supply deviations, however, arise for a variety of reasons, including:

- Late delivery,
- Short shipments,
- Production delays,
- Production yield differing from planned, and
- Substandard materials and production.

Finally, discrepancies between actual inventory and inventory records in corporate systems occur frequently and are often significant. Since computer systems generally initiate the operations and activities impacting inventory (e.g., procurement, production planning, logistics), it's necessary to account for inventory accuracy when computing safety inventory levels.

WHAT DO WE MEAN BY SERVICE LEVEL?

Most companies measure and manage service levels using a number of different definitions. Each of the service-level definitions focuses on how well you are able to satisfy specific demand as requested by the customer (e.g., meet the delivery dates, deliver the quantity requested). You can specify a number of different targets for each definition, but it's important to realize that each target must have a specified time frame, such as per hour, per day, per week, per month, per quarter, per year. The longer the time period, the easier it is to achieve a given level.

Some common service-level measures include line fill rate, order fill rate, demand fill rate, and 100% coverage rate, which are described in Table 1.

Table 2 shows how a company may set the following service-level targets for a specific product. Notice that there are two different demand fill-rate targets: one based on a monthly fill rate and the other on an annual fill rate.

Carrying safety inventory helps companies improve in each of these measures. The challenge, however, is to determine the lowest amount of safety inventory required to achieve all of the service-level targets. From an inventory-management perspective, the easiest measures to conduct safety inventory analysis on are demand fill rate and 100% coverage rate. The other measures are complicated by additional factors such as ordering behavior (e.g., the number of lines per order and the order quantity per line).

Companies often do one analysis for determining safe-

TABLE 1

MEASURE	DESCRIPTION
Line fill rate	Percentage of order lines that are filled perfectly
Order fill rate	Percentage of complete orders (all lines) that are filled perfectly
Demand fill rate	Percentage of total demand that is filled
100% coverage rate	Percentage of periods that have 100% demand fill rate

TABLE 2

MEASURE	PERIOD	LEVEL
Demand fill rate	Monthly	95%
Demand fill rate	Yearly	98%
Line fill rate	Monthly	90%
Order fill rate	Yearly	90%
100% coverage rate (measured monthly)	Ongoing	95%

ty inventory levels. But safety inventory analysis is so poorly understood that the analysis is rarely based on the correct assumptions—hence it's no wonder the resulting safety inventory levels aren't necessarily valid. In fact, many common techniques rely on statistical analysis, in particular on the normal distribution or bell curve. For more information on the normal distribution, see the sidebar, "Statistics 101 Review."

SOME COMMON MISCONCEPTIONS

To illustrate some misconceptions about safety inventory calculations, consider a scenario where a company tracks demand and the demand forecast for a given product on a monthly basis. It wants a 95% service level. Given this information, the inventory analyst applies one of three methods for calculating target safety inventory levels—demand variation, forecast error, and absolute error.

To begin, some critical information is missing, such as the service-level definition and the period for the service. The analyst knew the company wanted to achieve a 95% service level and then assumed that the requirement is to satisfy 95% of customer demand on a monthly basis or 95% demand fill rate on a monthly basis. For the rest of this discussion, we'll take the analyst's assumptions as correct, but of course you would have to validate them in the real world.

Method 1: Demand Variation

The analyst uses statistical techniques and historical data to analyze the variation of demand by computing:

Step 1: mean = AVERAGE (demand)
Step 2: standard deviation = STDEVPA (demand)
Step 3: service level, which uses NORMINV (0.95, mean, standard deviation)
Step 4: target inventory = level (The target inventory level includes forecast and safety inventory.)

Method 2: Forecast Error

The analyst uses the more sophisticated technique of analyzing the monthly forecast error (i.e., the difference between actual demand and forecast demand). The goal is to reduce demand uncertainty by improving forecast accuracy. This method includes forecast error and uses the same statistical techniques as the demand variation method to analyze the forecast error by computing:

Step 1: forecast error = actual demand – forecast
Step 2: mean = AVERAGE (forecast error)
Step 3: standard deviation = STDEVPA (forecast error)
Step 4: level = NORMINV (0.95, mean, standard deviation)
Step 5: safety inventory = level
Step 6: target inventory = forecast + safety inventory

Method 3: Absolute Error

Although the following calculation is the same as the previous method, the data are quite different. The analyst uses a technique similar to the forecast-error method to analyze the magnitude of forecast error, then uses the same statistical techniques to analyze the forecast error. The magnitude of the forecast error is computed using the absolute value function from Excel ABS (forecast error).

Step 1: forecast error = actual demand – forecast
Step 2: mean = AVERAGE (ABS [forecast error])
Step 3: standard deviation = STDEVPA (ABS [forecast error])
Step 4: level = NORMIN V (0.95, mean, standard deviation)
Step 5: safety inventory = level
Step 6: target inventory = forecast + safety inventory

WHAT ARE THE FLAWS?

Let's now examine why these methods fall short of the new approach I describe next.

Demand-Variation Shortcomings

The analysis in this method determines the amount of safety inventory for a "100% coverage" measure. The output of this method, target inventory, is correctly interpreted as being the amount of inventory that will cover 100% of demand 95% of the time. That is, 95% of the time actual demand will be less than target inventory. This will result in excessive inventory to achieve the monthly demand fill-rate objective of 95%—since in 95% of the months there's a 100% demand fill rate. Second, this method assumes the demand pattern is stable, so it doesn't take into account trends (e.g., growth or decline) or seasonality.

Forecast-Error Shortcomings

The analysis in this method actually determines the amount of inventory required to exceed the forecast error 95% of the time. First notice that the error can be either positive (actual demand exceeds forecast) or negative (forecast exceeds actual demand). You need to realize that only the positive errors should drive safety inventory requirements because negative errors result in your holding excess inventory and don't have a negative impact on product availability. The errors associated with most commonly used forecasting techniques can be assumed to

Figure 1: *DISTRIBUTION OF FORECAST ERROR*

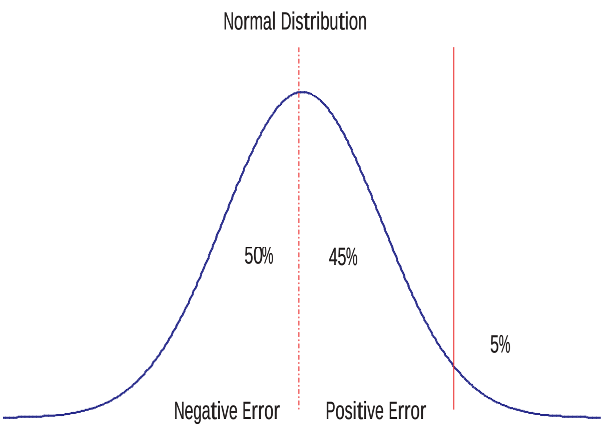

follow a normal distribution with mean zero where, on average, the positive and negative errors balance out.

The forecast-error method calculates the point that covers 95% of all errors, which, in reality, covers all of the negative errors and 90% of the positive errors (see Figure 1).

But this still results in excessive safety inventory. If you assume demand coverage, the percentage of demand covered by the inventory was 90%, and the forecast error would be 10%. In order to satisfy 95% of demand, it's only necessary to cover half of the positive error—not 95% of the error.

Absolute-Error Shortcomings

The analysis in Method 3 is faulty for several reasons. Most significantly, it assumes that the magnitude of the forecast error follows a normal distribution, the bell curve. This is almost certainly wrong. It's likely that the actual forecast error (e.g., when demand exceeds forecast the error is positive, and when forecast exceeds demand the error is negative) follows a normal distribution. Figure 2 shows the plot of a random sample of 20,000 numbers generated from a normal (0,1) distribution. Figure 2a shows the density function and takes on the familiar bell shape of a normally distributed sample. The density function of the magnitude (e.g., the absolute value) of the same sample is shown in Figure 2b and obviously doesn't represent a normal distribution.

As with the forecast-error method, this approach attempts to determine how much safety inventory is required to cover 95% of the forecast error rather than what's required to achieve a 95% demand fill rate.

The Biggest Oversight

Safety inventory buffers against stockouts due to uncertainty. To do so, it needs to address supply uncertainty as well as demand uncertainty, but none of the methods above takes supply uncertainty into consideration. They all implicitly assume that supply is 100% certain when, in reality, we know that isn't the case. In fact, in many environments, supply uncertainty is on a par with demand uncertainty. For example, have you ever attempted to make a purchase where the vendor wouldn't commit to when the product would be available? Examples range from the expensive (new homes, vehicles) to the moderate (computers, electronics) to the everyday (clothes, books).

Figure 2: *DISTRIBUTION OF ABSOLUTE VALUES*

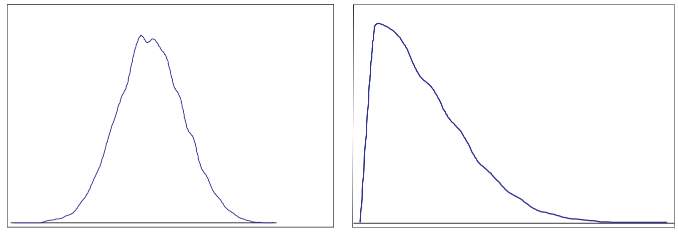

Figure 2a: Sample from N (0,1) | Figure 2b: Density Function of Absolute Values

A NEW APPROACH

Since inventory uncertainty arises primarily from demand uncertainty, supply uncertainty, and inventory accuracy, why don't we directly analyze inventory uncertainty rather than just uncertainty in demand or supply? Historically that's been difficult to do because inventory-planning applications weren't very robust. Over the past several years, however, there have been significant advances in the area of supply chain planning, including inventory planning. (See "A Modern View of Inventory" in the July 2004 issue of *Strategic Finance* for more details.)

I'll now introduce a new approach—coverage. This method calculates the safety inventory level based on the following service-level question: How much inventory is required so that at least z% of the demand is met in y% of the time periods? For instance, how much inventory is required to meet 95% of monthly demand 99% of the time?

Computing the amount of inventory required to close the gap between the demand coverage and the desired service level, coverage can be expressed mathematically as:

$$Coverage = \frac{Actual\ Inventory}{Actual\ Demand}$$

STATISTICS 101 REVIEW

I'M SURE YOU RECALL from your college statistics courses that the normal distribution or bell curve is used to model the distribution of a population of numbers.

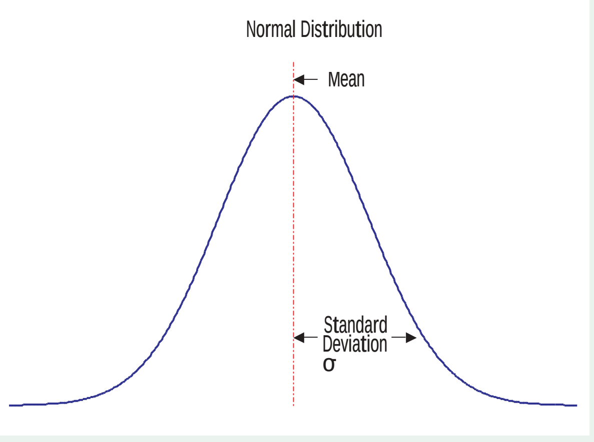

The key parameters involved in a normal distribution are the mean and the standard deviation. The mean is the average value of the set of numbers. The standard deviation is a measure of how spread out (or clustered) the set of numbers is.

The mean of a set of numbers is the same as the average and can be computed using the AVERAGE function in Excel.

The standard deviation describes the spread of a set of numbers and is calculated as follows:

- ◆ Determine the mean,
- ◆ Determine the difference between each number and the mean,
- ◆ Square each difference,
- ◆ Calculate the average of the squares, and
- ◆ Calculate the square root of the average.

The standard deviation can be computed using the function STDEVPA in Excel.

Of particular interest is determining the point x where you would expect a certain percentage, say level, of the population is greater than x. For instance, we may want to know the point x where 95% of the population has coverage greater than x. (This is equivalent to finding the point where (100 – level)% of the population is less than x.)

The point x can be computed using the EXCEL function NORMINV ((1-level), mean, standard deviation).

To take things one step farther, assume we have the ability to shift the population in some manner in order to meet a desired objective. For instance, the objective may be that a certain portion of the population is greater than some number, say target.

An Example

Consider a sample of six numbers from a larger population: 89, 76, 98, 82, 96, 77. Assume we want to determine how much we need to shift the population so that 95% of the population is greater than 80. In this case:

- ◆ level = 0.95
- ◆ target = 80

Stage 1: Compute x – the point such that 95% of the population is greater than x.

Step 1: Compute the mean.

AVERAGE (89,76,98,82,96,77)

(Mean = 86.33)

Step 2: Compute the standard deviation.

STDEVPA (89,76,98,82,96,77)

(Standard Deviation = 8.65)

Step 3: Compute x.

NORMINV ([1-0.95], mean, standard deviation)

(x = 72.10)

This means we expect 95% of the population to be greater than 72.10.

Stage 2: Compute the required population shift.

Let gap denote the required shift.

Gap = target – x

= 80 – 72.10

= 7.90

The Final Answer

If each member of the population is increased 7.90, then it's expected that 95% of the population will be greater than 80.

The challenge is in computing future coverage when neither future inventory nor demand is known. In order to estimate future coverage, we need to forecast future inventory as well as demand since the proposed approach works on the philosophy of ensuring there's adequate inventory to cover actual demand at the desired service level. The analysis is based on historical coverage of demand as well as forecast error, and safety stock is computed as a percentage of demand.

There are two key stages to the computation. The first stage analyzes historical demand coverage and determines the required safety stock as a function of actual demand. The second stage analyzes historical forecast error so that you can express safety stock as a function of forecast demand.

Figure 3: *ACTUAL DEMAND FROM SAMPLE DATA SET*

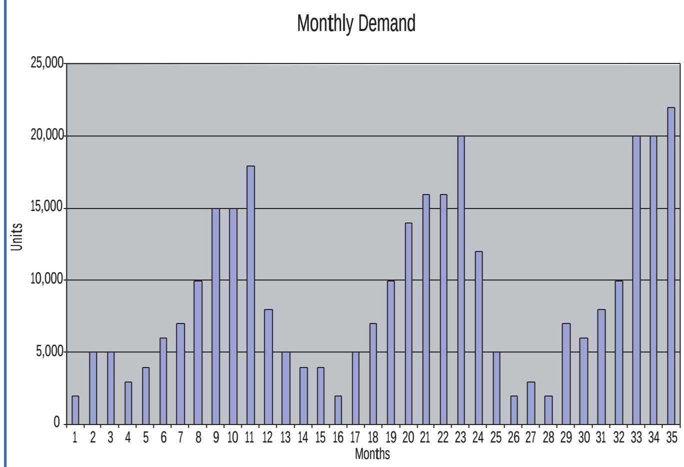

Stage 1

This stage analyzes historical coverage of demand. As a function of actual demand it computes the safety inventory required to achieve a desired service level.

Step 1: Compute demand coverage for each past period:

demand coverage = actual inventory/demand

Step 2: Compute mean and standard deviation of demand coverage:

mean = AVERAGE (demand coverage)

standard deviation = STDEVPA (demand coverage)

Step 3: Compute the gap

The gap is the percentage of demand needed to be covered by inventory in excess of the forecast demand. We need to compute the point p, where y% of the periods have demand coverage greater than or equal to p.

P is computed as:

p = NORMINV ([1-y], mean, standard deviation)

We need to add enough safety inventory to close the gap from p to x. That is, we need to increase the demand coverage from p to z, so the gap is the difference between the desired level z and p:

gap = z–p

By adding units of safety inventory, the new demand coverage becomes:

demand coverage (new) = (actual inventory + safety inventory)/demand = demand coverage (old) + (safety inventory/demand)

So, to increase demand coverage from p to z, we need:

z = p + (safety inventory/demand) or

safety inventory = (z – p) * demand

safety inventory = gap * demand

Unfortunately, when computing safety inventory targets for future periods, actual demand isn't known, so we need to express safety inventory as a function of forecast demand rather than actual demand.

Stage 2

In the second stage, a factor f is computed such that f * forecast > demand for y% of the periods. This computation is similar to the computations in Stage 1, but, instead of analyzing the coverage of actual demand by actual supply, we focus the analysis on the coverage of demand by forecast.

Step 1: Compute coverage-by-forecast for each past period:

coverage-by-forecast = forecast/demand

Step 2: Compute mean and standard deviation of demand coverage:

mean = AVERAGE (coverage-by-forecast)

standard deviation = TDEVPA (coverage-by-forecast)

Step 3: Compute forecast gap

The forecast gap is the shift required in the coverage-by-forecast so that the coverage is greater than 1.00 in y% of the periods. First compute the point p where y% of the periods have forecast coverage greater than or equal to p. Thus p is computed as:

p = NORMINV ([1-y], mean, standard deviation)

The forecast gap is the difference between the desired level of 1.00 and p:

forecast gap=1.00 – p

Thus we expect the following: forecast * (1 + forecast gap) > demand in y% of the periods.

Substituting forecast * (1 + forecast gap) for demand in the safety inventory calculation of stage 1, we now compute the required safety inventory as a function of

Figure 4: *DEMAND COVERAGE BY TECHNIQUE*

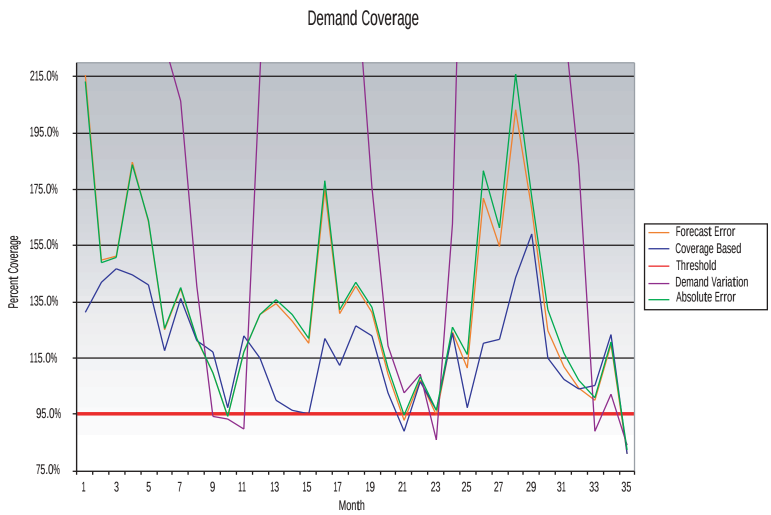

forecast demand. That is:

safety inventory = gap ∗ forecast ∗ (1 + forecast gap).

This calculation overestimates the required safety inventory in y% of the periods.

To recap, the first stage computed the required safety inventory as a function of actual demand. Stage 2 then determined a factor so that I could calculate safety inventory as a function of forecast demand. I haven't addressed the correlation between forecast error and demand coverage because this would require a deeper mathematical analysis, but the following example lends evidence that this isn't a significant oversight.

HOW DO THE TECHNIQUES COMPARE?

To compare and contrast the various techniques, let's look at a sample data set and the various safety inventory calculations to achieve 95% demand coverage in 19 out of 20 months, which would be 95% of the time.

Figure 3 shows the actual demand for each month. The demand is very cyclical with a year-over-year growth trend. Demand forecast error and supply variation factors represent best-in-class companies in a number of industries. Demand forecast error is less than 16.5% in 90% of the months (mean 0, standard deviation 0.1), and supply variation is less than 10% in 90% of the months (mean 0, standard deviation 0.06).

Figure 4 shows the results by graphing actual demand coverage for each of the 35 months. The horizontal red line indicates the threshold or desired service level.

Any point above the red line means that demand coverage exceeds the target of 95% coverage. The larger the space between threshold and the actual coverage equals excess inventory, and, as you can see, the new method—coverage—consistently carries the least amount of excess inventory. The two methods based on forecast error (forecast error and absolute error) are essentially equivalent in terms of demand coverage, and the demand variation technique clearly carries unacceptable amounts of excess inventory.

The target was to achieve 95% coverage in 95% of the months. Over 35 months this means we should have less than 95% coverage in two of the months. Of the four techniques, only the new coverage technique achieved this goal, while demand variation (six months with less than 95% coverage), forecast error (four months with less than 95% coverage), and absolute error (three months with less than 95% coverage) missed this performance target. Clearly the coverage technique outperforms the other techniques in terms of achieving the target—95% coverage 19 times out of 20—as well as in doing so with the least amount of inventory.

NOW IS THE TIME

My viewpoint of safety inventory is likely new for many of you. After all, the commonly held belief is that safety inventory exists to cover errors resulting from either variability in demand or inaccurate demand forecasts. You now understand that you should think about safety inventory from the perspective of ensuring there's adequate inventory to cover desired levels of demand. Not only does demand impact inventory, but so does supply—and neither is 100% predictable.

Don't wait to reassess how your company deals with safety inventory. Improving your safety stock policy will increase revenue through improved service levels while simultaneously reducing inventory carrying costs. ■

Scott W. Hadley, Ph.D., has delivered SCM knowledge and solutions around the world since 1991. At present, he teaches SCM at the Rotman School of Management, Canada's top-ranked MBA program. He spends much of his time and effort educating executives and senior managers on how to create value from their supply chains. He is also an associate of Toronto-based Focused Management Inc. You can reach Scott at (905) 681-8644 or shadley@sympatico.ca.