The Impact of Integrated Business Planning on Safety Stock Calculation Methods

by Bill Tonetti

Introduction

Approaches for calculating safety stocks have evolved over the years, keeping step with systems and methodological improvements. Before forecasting systems became widely used, methods for calculating safety stocks centered on demand history statistics like the standard deviation. With the wide availability of automatic forecasting systems, businesses turned to fitted statistics such as the standard error (or RMSE). Today, as holistic business planning processes like Sales and Operations Planning (S&OP) and Integrated Business Planning (IBP) have emerged, businesses are developing planning and data retention competencies that can drive a major evolution in inventory optimization practices.

S&OP involves the routine periodic development of consensus demand, supply and financial plans, typically beginning with a “demand review” step. A key output from the demand review process is a published forecast with documented assumptions. For better accountability and ongoing process improvement, previous forecasts are retained and systematically compared with actual demands. These previous forecasts can be leveraged to support better safety stock calculation methods and wiser corporate inventory investments.

Why Businesses Invest in Safety Stocks?

Inventory is expensive. It consumes scarce working capital, sometimes becomes obsolete, consumes valuable warehouse space, and can detract from the agility of a business. So, why do companies invest in safety stocks?

The answer is simple: Businesses hold safety stocks because there are circumstances where they need to provide goods to their customers in more immediate time frames than the sum of the time that it takes to source and produce those goods. Safety stocks are a buffer that enables them to provide better service.

Sometimes buffer inventories are held in the form of finished goods inventory. Other times, businesses are able to postpone some of the manufacturing steps and store safety stocks in the form of less expensive and more versatile raw materials or intermediates. But, in all cases, rational businesses invest in safety stocks because they need to make sourcing or production decisions before they are able to gain perfect information about demand.

Basic Safety Stock Calculation Methods

The most widely published safety stock formula is actually rarely used today. It is based on the series standard deviation (SD), a descriptive statistic that quantifies the dispersion of observed values relative to the series mean (average). The basic safety stock formula is shown below:

SS = z * SD

Where:

z = A one-tailed lookup from a z table (normal distribution), given the desired service level. (For example, a z score of 2.33 represents a distance 2.33 standard deviations from the series mean which translates into a 99% service level. A z of 1.65 represents a 95% service level.)

SD = Series standard deviation, usually estimated from the sample series.

$$= \sqrt{\frac{\sum (x - \bar{x})^2}{n - 1}}$$

There are a few things to like about this technique. Most importantly, it considers the variability of demand and does not require a forecast. That makes it easy to calculate. Forecasts can sometimes be difficult to calculate and this method simplifies that problem away. Secondly, it considers the desired service level which is a key consideration in buffer inventory sizing. And lastly, it's easy to understand, clearly encompassing the essential notions of unpredictability and service.

However, there are also many weaknesses. One weakness is that it assumes that the values will be normally distributed, which is not the case since the demand is bound by zero on the low end, but not bound at all on the high side. A second weakness is that the standard deviation of demand must be calculated using time slices that correspond to the lead time. Since this is often not possible, the above formula can be modified to include lead time as well as a mechanism to scale it to align with the time bucketing that is used to calculate the standard deviation. The modified formula is as follows:

SS = z * SD * SQRT(LT / DFP)

Where:

LT = Lead Time (in days)
SQRT = The mathematical operator "Square Root"
DFP = Days per forecasting period. For example, 30 if using monthly time periods to calculate the standard deviation of demand, or 7 if using weekly time periods.

The above technique is a meaningful improvement over the first formula, but it still has some crippling shortcomings.

Using Unexplained Variation

One major problem with the basic approach described above is that it considers all demand variation and not just the unpredictable variation. It is possible or even likely that some of that variation may be predictable, so the use of the historical standard deviation will lead to excessive safety stocks. For example, if the demand is significantly trended or seasonal then a two- or three-parameter exponential smoothing method could explain much of the variation. A simple method like using the series SD would clearly produce unsatisfactory results in those cases. However, we would have to fit a model to the series in order to consider only the unexplained variability.

A host of useful statistics become available once a model is fit to the historical series and a forecast is produced. To filter out explainable variation, many companies have chosen to use the fitted Root Mean Squared Error (RMSE) in place of SD. RMSE is calculated similarly to SD, but it uses the squared deviations of the fitted values from the actual. If the model is able to represent or track the trends or seasonality then RMSE will be considerably lower than SD, resulting in lower required safety stocks.

Replacing SD with Fitted RMSE, we get the following formula:

SS = z * RMSE * SQRT(LT / DFP)

One common elaboration on this technique is to convert the safety stock to days using the series mean. With the minimum converted to coverage days, safety stocks can vary over seasons, increasing the requirement during peak seasons and lessening inventories during times when they aren't needed. Now, we've got a pretty good formula, and it is indeed the one that is used by many companies and commercial inventory planning software solutions.

The "Fit Problem"

Time-series forecasting methods begin by fitting a model to history, then projecting the fitted model into the future to produce a forecast. Complex models such as trended and/or seasonal ones will typically fit the history better than simpler linear models. However, that doesn't necessarily mean that they are better at predicting future values. Several academic studies, in fact, have found that simplistic approaches such as a random walk (next period equals previous period) often outperform other methods in predicting short-term estimates such as stock price fluctuations. Most world class forecast engines deliberately discard complex models in cases where tests for seasonality or trends indicate a reasonable probability that the better "fit" of the complex model is less likely to yield better predictions. Due to the difference between fit and prediction, it is very common for fitted statistics such as the fitted RMSE to understate unexplained variability. Hence, the "fit problem."

There are other reasons why fitted RMSE may be a poor measure of unexplained variability. Most of them have their roots in the fact that business forecasts are rarely purely statistical. For example, many companies communicate demand expectations with their suppliers and such communications can significantly improve forecast accuracy. Or, sometimes collaboration can reduce forecast accuracy. Gamesmanship can result from constrained physical or financial capacity or incentive plans that can cause deliberate under- or over-forecasting. Another structural issue is bias caused by forecasts that are bound by zero on the low side but unbounded on the high side. Bias can also result from over-forecasting following large promotions or pipeline fills, though automatically ruling out periods with outliers can cause other problems such as underestimating true error.

In summary, there are a variety of reasons why RMSE will be a less than perfect indicator of unexpected differences between forecasts and demand. Not only are there many potential causes, but the causes are widely present so the chances are excellent that RMSE will misrepresent true unexplained variability.

RMSE Based on Archived Forecasts

What we need is a measure that is based on actual final forecasts. Not fitted ones, but the real ones that businesses actually use when they make manufacturing and distribution decisions. Just as forecasting systems enabled the advancement from SD to fitted RMSE, modern integrated demand management systems are preparing businesses for the next major step forward.

As businesses have progressed from planning silos to integrated sales and operations planning, demand management systems are increasingly incorporating forecast performance measurement and reporting capabilities. Sales and operations planning processes typically begin with a root cause analysis of recent forecast performance and conclude with a consensus plan including basic assumptions that are embedded within it. To support these processes, demand management systems have deepened their capabilities for retaining, comparing and visualizing current and previous forecast scenarios. With the storage of actual previous forecasts, we are now ready to take the next step forward in optimizing inventories.

As long as we have archived forecasts, then the solution is actually quite simple. All we need to do is to calculate a new RMSE based on prior forecasts instead of fitted ones. The only remaining question is which archived forecast to use. In theory, it would be best to use the archived forecast with a lag that approximates the lead time. In any event, an actual RMSE based safety stock calculation method should outperform fitted calculation methods.

Supply Variability

In addition to considering the variability of demand, some companies find that they have supply variability issues such as changing lead times or unexpected shortages. While demand variability is by far the prevailing cause for buffer inventories in our modern economy, the variability of supply can be an issue as well.

There are two relatively simple ways to accommodate supply variability in the calculations for safety stock. One is to extend the lead time. If the average lead time is 30-days, but sometimes it takes 45-days, then one conservative and practical approach would be to use the longer lead time. The second option is more precise, but it requires data that may not be readily available. It is possible to combine the variance of demand with the supply variance. Taking the square root of the summed variances, you would then have a combined standard deviation that could be used to calculate safety stocks.

About the Author

Bill Tonetti is the President of Demand Works. Bill's career began in consumer product sales, then moved into advertising. When he was ready for a change, he obtained an MBA from the University of Virginia and took a job working in operations management for a large paper company. While working in the paper industry as a supply chain manager, he implemented his first demand management solution as well as various other solutions including an ERP system and optimizations for production and transportation. He became passionate about supply chain management and he's spent the last two decades working, consulting, and providing demand and supply chain solutions for hundreds of consumer and industrial companies worldwide.

Bill has presented and published a variety of works in the demand management area and he holds a seat on the practitioner advisory board of the International Institute of Forecasters (IIF).

About Demand Works

Originally founded in 1993 as Advanced Planning Systems, Demand Works provides demand and supply planning solutions that deliver large-scale improvements in forecast accuracy, coordination and asset deployment. Demand Works is exclusively focused on the areas of forecasting, demand management and supply planning and S&OP, and we support hundreds of corporations throughout the United States, Canada, Latin America, Europe and other parts of the world. Our customers include manufacturers and distributors of many household brands as well as countless industrial and other manufactured goods. The company's principals are respected leaders in the sales and operations planning area, and the solutions reflect their innovative thinking.

Demand Works has taken a holistic approach toward eliminating unnecessary cost and complexity without sacrificing functionality or business results. The engineering is top-notch, leveraging the most exciting forecasting, OLAP database and collaborative Internet technologies. Demand Works invented many leading practices in the industry, most notably a technology called Pivot Forecasting® that makes it possible for businesses to forecast and automatically synchronize plans at any level of product, manufacturing or distribution aggregation in real-time. By combining innovative solutions engineering and business practices, Demand Works is changing the way people think about best-in-class demand and supply planning software.