BWH News Letter
Volume 3 Number 4

The Carburetor in Your Replenishment System

Replenishment systems use demand forecasts and current inventory levels to identify replenishment requirements. To do this, they must answer two questions: "When to replenish?" and "How much to acquire?" Both questions must be answered by analyzing the forecast demand and current inventory. This analysis depends very heavily on a single table maintained internally in your replenishment application. Much like the carburetor for an engine, this table determines how "rich" the recommendations will be coming from the replenishment application. And like the carburetor, some people will tweak this table to get better returns on their inventory.

Before I get too far into this, I must separate the replenishment of store inventories from warehouse inventories. This is necessary because of the need for "presentation stock" and the effect of case pack on the store replenishment decision. Store replenishment must consider the high labor cost associated with handling partial cases of merchandise. For that reason, the replenishment quantity for the stores is often determined from the case pack of the item. Especially for supermarkets, where 90% of the items sell less than a case a week, the presentation stock and case pack are the major determinants of inventory level. Rather than handle partial cases, the retailer will set presentation stock so that the shelf can accommodate a new case while there is still enough inventory to satisfy demand during the replenishment cycle. The smart retailer recognizes these items and their replenishment system refills them during the "mid-week lull" to avoid the heavy Friday and Monday delivery schedules. The point is that, the table I am about to discuss is a bigger factor in situations where minimum order quantities (e.g. case pack) cover less time than an order cycle and presentation stock is not a factor. That is, for warehouse replenishment.

Back to the Carburetor. The conversion of forecast and inventory to a replenishment quantity requires consideration of the actual values and their believability. We are used to seeing forecasts, such as the weather, reported with percentages of likelihood. This same information is available for demand forecasts. It is most often calculated by comparing past forecasts to actual demand. This comparison gives a "demand variability factor", often reported as an absolute deviation (ignoring whether the forecast miss was low or high). (It may also be calculated separately for low and high misses so that the replenishment application can consider them separately.) To handle this forecast error, the replenishment application factors in a "safety stock" factor into its recommendations.

The replenishment application needs to understand how critical an out of stock situation is for an item. This is probably the easiest and most often used method for adjusting replenishment recommendations. It is usually expressed as a "service level percentage". The buyer or inventory manager sets the desired service level based on the item's importance. "Importance" can be interpreted in a variety of ways. Most often, inventories will be stratified based on item demand. The top moving items will be put in the highest stratum while ranges will be set for each of the lower strata. Typically there will be five strata with service levels ranging from a 99% to 90%. There may also be a "do not change" category for items whose service level has been specified manually and should not be adjusted based on movement. Converting service level percentages to unit quantities is where our "replenishment carburetor" comes into play.

The desired service level must be converted to a number of units of safety stock. Forecast accuracy has a big impact. The best case would be forecasts and actual demand that match exactly and have zero deviation. The safety stock for these items would also be 0 and the replenishment quantity would be strictly for the expected demand. In reality, there is almost certainly going to be forecast error. Something will cause the actual demand to deviate from the forecast. The result will be the need to convert the past error rate into some type of unit quantity for safety stock.

The next chart illustrates the conversion of service level to k factor. It is only an illustration;

bill@hwbittner.com 201-906-6925

I did not copy any application's actual values. The general format of the curve is accurate, leveling off at the upper extremes to avoid high inventory levels as you approach 100% as a desired service level.

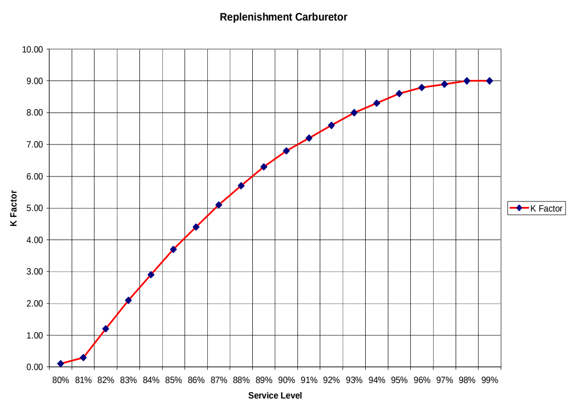

As you can see, the replenishment carburetor is a two dimensional table. The independent variable or x axis of the replenishment table is the service level requested. The dependent variable or the y axis of the replenishment table is the "k factor". The k factor is simply a value that is multiplied times the historical forecast error to determine the safety stock necessary to achieve the requested service level. There is nothing magic about the k factor values. Much like an engine carburetor, the only way to determine the right values for k is to monitor the actual results achieved and continue to adjust k until outcomes match the desired service levels. The problem is that most people don't even know this table exists in their replenishment application; let alone how to adjust it.

Unfortunately, determining the proper table values for k requires significant sample data. If the existing applications capture the data, the analysis is pretty simple. The problem is that when you begin the analysis, you may have to wait until enough historical data is captured. You need to capture past replenishment decisions, including the forecast, error rate, desired service level, order interval, units recommended, actual demand over the order interval, and units on hand at the end of the order interval. Since the same table is used for all items, you can analyze many items at once.

The result is that you will see how many cases are actually left at the end of each replenishment cycle and compare this to the desired service level.

The final analysis should compare several order cycles for all the items at a desired service level. For example, take all the items that were at a 95% service level. If 5% of the time, these items were out of inventory when the order cycle ended, your k factor is correct. If there is a significant difference and you are comfortable with raising or lowering inventories to achieve the desired result, you can change the k factors for various service levels in order to improve the replenishment recommendations. Increasing the k factor will raise both inventory and the actual service level. Lowering k will decrease both. In any case, this is not for the feint at heart and just like adjusting the carburetor on your car may invalidate the warranty. You need to discuss the decision to make these changes with your software vendor and monitor their impact on inventory and actual service level.

Most individuals never adjust the carburetor of their car, but the ones who do achieve the benefit of an fuel to air ratio customized for their circumstances. The same applies for replenishment application users. Most will never adjust their "k factor", but the ones who do will achieve superior results.