

HOW TO IMPROVE DEMAND FORECAST ACCURACY



A well established fact is that improved forecast accuracy leads to many downstream improvements not only in operations, but in a variety of business areas such as customer service and asset management. This white paper concerns itself with the improvement of forecasts that support the management of a company's supply chain. Inevitably, this leads to a focus on improving the product stocking levels or SKU forecast which drives procurement and manufacturing planning. These effects can even be seen in a lean environment where forecasts are needed for rough cut planning. The methods covered in this white paper will lend themselves to the classic make-to-stock (MTS) environment where operational and financial efficiency of a company is largely driven by forecast accuracy.

With the prime focus on product forecasting, understanding and predicting the elements that make up the final SKU forecast are the basis for improvement. The individual buying patterns of products at the customer level is normally too erratic or statistically random to be effective for supply chain planning, so this leads to the customer-aggregated product level. In some cases even the product-level forecast is too random to be useful, which leads to an aggregated level such as product sub group, group or category level. Even though forecasting at these summarized levels produces improved statistical results, the challenge of disaggregation to the prime procurement or manufacturing level (product) exists. Methods of breaking down aggregated forecasts will be covered in subsequent white papers, along with the consideration as to the use of units or currency.

The forecast accuracy improvement problem can be approached by examining the underlying elements that must be considered before the analysis and adjustment process is complete and a routine production forecasting process is operational.

There are important prerequisites that need to be in place before any real analysis, planning and improvement can begin.

- Cleansed data (data entry or transmission problems) that passes the sanity check of a demand planner.
- An easily accessible repository that contains both product-customer sales data with all appropriate dimensional hierarchies and facilities for storing forecasted results.
- A sophisticated forecasting platform that allows testing many combinations and data modifications to support improvement methods.
- A robust statistical forecasting engine that has "seamless" access to the data and modifications made to the data stored in the repository.

THE SHORT LIST OF FORECAST IMPROVEMENT METHODS

- Create at least 2 additional demand data measures. The first (DM1) should be a copy of actual unit sales, and the second (DM2) will be used as an option to deal with outliers discussed later. DM2 will likely be initially copied from the final iteration of DM1. From a supply chain perspective, product unit-level forecasting is the primary driver of the replenishment process. However, if higher levels of aggregation are chosen for any reason, then currency may be the only common dominator to use in the forecasting process. Additionally, it is sometimes appropriate to produce both unit-level forecasts and currency-denominated forecasts for business planning purposes.

- Filter out each customer's product sales from DM1 history where not purchases were made in the last 12 months. This can be done easily with a two-dimensional Stratum Planner view displaying all the customers by each of the top products using a User List as shown below. Simple use of a planning write back in Stratum Planner can set all the values to zeros in the shaded area as shown below. The result is a modified DM1 measure that has the history removed for those customers who will not be buying this product in the future. Product demand history should only reflect the active buying components (customers) in the future. Once the process of cleansing history is complete, the DM1 data measure should be copied to DM2 for further modification, as required, using events or demand adjustments.


Prod Nbr	UserList: top products		Values			
Bill To			*All Values*	Sales Units	Sales Units	Sales Units
Product Number	Description	Bill To	Description	FY2006	FY2007	FY2008
509	Widget1	26955	Customer1	120	0	0
509	Widget1	46220	Customer2	1,000	0	0
509	Widget1	59989	Customer3	40	0	0
509	Widget1	59990	Customer4	0	40	0
509	Widget1	59983	Customer5	40	80	40
509	Widget1	59988	Customer6	520	520	40
509	Widget1	59900	Customer7	0	800	160
509	Widget1	21385	Customer8	0	0	16,120
509	Widget1	48060	Customer9	13,080	320	31,080
509	Widget1	43970	Customer10	0	0	106,480
509	Widget1	90777	Customer11	29,480	62,680	116,920
509	Widget1	18394	Customer12	395,280	363,600	389,520
509	Widget1	18384	Customer13	2,149,240	2,253,560	2,400,680
			Totals for: Widget1	2,588,800	2,681,600	3,061,040

The same process described above could alternatively be applied to each of the top customers>products where products that have been discontinued or are planning to be discontinued can be removed from DM1 history by setting the values to zero. As in the previous approach using product>customer, this “normalizes” the product history data to only include the components of products that should be forecasted in the future.

- Produce a Baseline Forecast using the product elimination DM1 adjustments – Once DM1 has been product cleaned as shown above, a baseline forecast for all products will provide a means of estimating how much, if any, outlier or promotional demand adjustment will be required to meet the forecasting accuracy targets. This forecast run should be stored in a baseline measure in the repository and used as the benchmark for measuring additional improvements. An analysis view should be created to review all products using the following considerations:
 - o Ranking of SKUs by importance (\$, volume, Margin, etc)
 - o Individual MAD, MAPE, R-squared, mean and Standard Deviation
 - o Elimination of new products that have less than 2 years of history
 - o Review of period to period forecast output for “flat lines” indicating history is random without further adjustment
- Create a Stratum Planner Workbench view for analysis that provides the following
 - Gross Sales Units – forecast engine input
 - Fitted Units – model forecast of history
 - Average error for AVG which is MAD
 - o Absolute error = ABS(actual history – model fit)
 - o Average absolute error over total historical range
 - Average error for STD which is STD of MAD (dispersion)
 - o Standard Deviation of the absolute error of full range


- SD's per period – magnitude of outlier or promo
 - o Each week Absolute Error/Full range SD of MAD
- Outlier Ratio – magnitude of outlier or promo
 - o Each week Absolute Error/Full range absolute error (MAD)
- The key formulas for the Stratum Planner workbench are shown below
 - Average error for AVG a.k.a. MAD
 - o $\text{ABS}(\text{DM1 Sales Units}[\text{Wk start to Wk end} - \text{Model Fitted units}])$
 - o Summary Band - Average
 - Average error for STD
 - o $\text{ABS}(\text{DM1 Sales Units}[\text{Wk start to Wk end} - \text{Model Fitted units}])$
 - o Summary Band – StDev – Standard Deviation
 - SD's/error – Primary measure of outlier significance
 - o $\text{Avg Error for AVG} / \text{STDEV}(\text{Avg Error for STD DEV FOR ALL})$

Example Planner analysis View using the workbench formulas described above



This view is used to compare the baseline forecast to the event modified or demand (DM2) adjusted forecast. The view should only contain the historical period not the forecast periods since the statistical measures should be calculated with on history. Planner “traffic lights” as shown in **blue** above are used to highlight where the standard deviation of the error between the fitted value of the forecast and the historical demand exceeds a limit set by the user. It is suggested the limit be set at 4 Standard Deviations (SD) to detect outliers.

- Option 1 – Demand adjustment - after finding the individual time periods that have exceeded the limit above, modify DM2 to match the fitted value using a planning update process. This serves to remove the outlier from demand history and replaces it with the value produced by the forecast engine's best fit model. The advantage of this method is it allows the forecast engine to select other models rather than event-modified exponential smoothing. The disadvantage is that DM2 modifications are not easily tracked since DM2 may have other changes made to it.
- Option 2 – Event Marking - as in the demand adjustment option, after the limit-exceeded period has been detected, an event tag should be placed in the Event measure of the forecast category using the planning update process. The advantage of this method is it allows for easily tracking of modifications since the events are placed in each individual period of a separate Event measure for each SKU. The disadvantage is that the exponential smoothing model is always chosen by the forecast engine. In most applications ES will produce the best fit for reasonably forecastable data.
- Create an iterative process to analyze and tag/modify demand. It is suggested that optimum results can be obtained for outlier correction and promotional tagging in one to two analysis and tagging cycles. The process begins with the detection of outliers/promotions that exceed the limit followed by tagging of events and/or modifying DM2 to the fitted value and rerunning the forecast to compare improvement results. The primary metrics for understanding improvements in a given SKU are:
 - o MAD of base forecast vs. MAD of modified forecast (lower is better and means there is simply less error across all periods of history)
 - o Standard Deviation of MAD of base forecast vs. Standard Deviation of MAD of modified forecast (lower is better and means the error is more uniform with fewer random events)
 - o Typical improvements of 50% can be expected, resulting in a more accurate forecast assuming the outliers were truly outliers. If they were not and could appear again as with a promotion in the future then each should be marked in the expected future periods with the same indicator value.
- Inspect each SKU by comparing the yearly sales and forecast side-by-side in a period-based View to isolate possible future event to mark. By using both the data and a line graph as shown, future events can be differentiated from outliers. Events that are likely to repeat in the future may not occur in history at precisely the same period. Forecast engines look at precisely the same periods throughout history to determine future seasonality. Non-periodic events as described can be tagged for the forecast engine to analyze and use as predictors for future events that have been tagged with the same number in future periods.



- Simplified Seasonal Patterns (SSP) can be used to smooth out weekly data. The process involves setting the SSP check box in the Stratum forecast engine and then determining what level of smoothing should be applied – bi-weekly, monthly or quarterly. Conceptually, this deals with the issue of semi-seasonal demand patterns that occur in close time proximity but not precisely in the same period year to year. The process is simple to test, and it may produce the best results in the end if more time granular forecast is not required. Of course, the resulting forecast is not as granular (monthly) as the underlying weekly data.
- Become familiar with Forecast measurement methods – in addition to the basic measurements such as MAD, MAPE, Standard Deviation, R-Squared and Mean, there are several metrics that should be considered but not required in all situations.
 - o Forecast Value Add – each period statistical forecast improvement over the naive forecast (simply the last periods actual). Simply stated is the stat forecast better than using last periods actual.
 - o Coefficient of Variation (COV) – Standard Deviation over a historical period divided by the mean of the same historical period. Values greater than .7 indicate forecasting without adjustments to history may be problematic, but not impossible. A high value of the COV suggests more random data or more no seasonal events are in the data.
 - o Tracking Signal – rolling sum of the non-absolute forecast error over a user-selected period such as the average product lead time or the supply chain reaction time. This provides a cumulative measure of forecast error during a short planning cycle and could be used to trigger re-forecasting.

- o Upper and Lower Confidence limits of the forecast – With a user-defined probability such as 95%, two values are automatically produced by the forecast engine which represent the highest and lowest limits of the forecast at each period in the future for the user defined level of confidence (e.g. a typical band is 95% which results in an input to the engine of 2.5% on the low limit and 97.5% on the high limit). The upper confidence limit can be used as a rough approximation of demand + safety stock. The lower approximation could be used as a means of calculating worse cases for sales and manufacturing utilization. A Confidence range of 95% suggests there is a 95% chance that the forecast produced by the engine will fall between the upper and lower limits.
- Using custom forecasting models vs. expert model selection in the forecast engine – This approach requires a solid understanding of statistics and is not suggested for the casual user. In most cases, custom models are only used to deal with high-error forecasts produced by the expert selection and after appropriate adjustments have been applied and forecast iterations have been done. This can be an effective method for products with little history or products that are abnormally influenced by very near-term data not specifically accounted for in expert selection.
- Separate out new products to be forecasted using different special methods – This process will be discussed in more detail in subsequent white papers. At a summary level, new products lack one very important characteristic - historical demand data. When considering approaches for forecasting new products, the process will be different if the new product is a product line extension, new to the world or has a collection of similar mature product attributes that must be considered independently. In most situations new products (life cycle less than 2 years) are manually dealt with by the new product launch team. Once history begins to develop, the demand planning organization assumes responsibility and statistical forecasting methods can be effectively applied.

FINAL THOUGHTS

The customer-product and outlier cleansing processes described above can be time consuming. However, the process is performed only one time on historical data, providing an annuity for the invested time. A continuing process (quarterly or monthly) should be employed to “true up” the new period demand data which should be far less time-consuming and, in fact, is a good sanity checking review process for incoming data to be used in the forecasting process. While the automated forecasting process can be treated somewhat like a “black box” approach, a clear basic understanding of the process, measurement methods and adjustment methods is critical to the forecast improvement process.

The up front effort (time and thought) applied to the forecasting process can be a high return investment considering the downstream impact of improving forecast accuracy. Industry studies have found that an improvement in forecast accuracy produces significant downstream improvements in perfect order rates, reduced inventory levels, and higher profit margins. The return is clearly there, but the means of improving forecast accuracy is not always obvious or easy – if it were, forecasting would not be as important as it is today.

Send comments or questions to: john.hughes@silvon.com

About Silvon Software

Silvon Software is a global supply chain-focused business intelligence solutions provider headquartered in Chicago with more than 1,500 customers globally. The company's Stratum™ suite of operational planning, analysis and reporting applications is designed to help companies strategically plan, analyze and manage the performance of their enterprises and supply chains. The product suite features hundreds of pre-built analytical views, KPIs and reports; forecasting and collaborative planning functionality; exception management capabilities; plus, a number of flexible information delivery options for sharing data internally and with external supply chain partners. For more information, visit www.silvon.com



**NORTH AMERICA
CORPORATE HEADQUARTERS**
Silvon Software, Inc.
900 Oakmont Lane, Suite 400,
Westmont, IL 60559
Ph: (630) 655-3313
Fax: (630) 655-3377
Toll-Free: (800) 874-5866
E-mail: info@silvon.com

**EUROPE/MIDDLE EAST/AFRICA
EUROPEAN
HEADQUARTERS**
Silvon Software, Ltd. (UK)
Pinewood Studios
Pinewood Road
Iver Heath, Bucks SL00NH
Ph: +44 (0) 1753 631133
Fax: +44 (0) 1753 635192