Summary

• Traditional forecast error metrics, such as the MAE and MAPE, do not reveal the financial impact of forecast error. Better decisions can be made if we have an actual Cost of Forecast Error (CFE) metric.

• A CFE calculation should include both inventory costs and the costs of poor service (stock-outs). Importantly, it will incorporate the costs of holding safety stock and maintaining a desired service level.

• Applying a cost to forecast error helps to show the trade-offs inherent in varying the service level and also helps to determine optimal safety stocks.

Contact Details

Dr Peter Catt
Supply Chain Solution Manager

cell: 027 542 2053
ddi: 09 571 7105
peter.catt@soltius.co.nz

www.soltius.co.nz

soltius
TALENT AND ATTITUDE

THE COST OF SALES FORECAST ERROR: A PRACTICAL EXAMPLE

By Dr Peter Catt

This whitepaper is based on Peter's 2007 paper *Assessing the Cost of Forecast Error* published in *Foresight: The International Journal of Applied Forecasting.*

Introduction

Forecast accuracy can play a central role in increasing shareholder value, particularly when the enterprise is reliant on long-fulfillment lead-times. However, forecast error metrics, such as the Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE), do not reveal the financial impact of forecast error.

Flores et al. (1993, p. 140) state that "statistical measures of forecast accuracy are not designed to capture the economic implications associated with managing an inventory system." Roberts and Whybark (1974, p. 638) consider forecast cost implications as "probably the most essential measure." Hence, although measures of forecast accuracy serve an important purpose, it is also highly desirable to determine the financial costs associated with forecast error.

This paper will show how to calculate a cost of forecast error (CFE) and will examine the key issues and choices involved in the calculation. By applying the CFE one can compare the financial performance of different forecasting methods and also see the cost implications of specifying different service levels. The assessment of the costs of forecast error should include both inventory costs and the costs of poor service. Mitigating forecast error helps maintain desired levels of customer service while controlling the costs associated with excess inventory.

The balance between service level and inventory costs is achieved through safety stock, defined by Silver et al. (1998, p. 31) as "the amount of inventory kept on hand, on the average, to allow for the uncertainty of demand and the uncertainty of supply. . . ." Too much safety stock means that inventory costs will be high, while too little safety stock means that customer service levels may suffer. The right amount of safety stock strikes a balance between these two costs.

Inventory Costs and Safety Stocks

Inventory costs are the costs associated with procuring, producing, and carrying inventory. The procurement/production cost per unit is called the unit variable cost. For purchased items, it includes the purchase price plus freight costs. For manufactured items, it includes production and other costs associated with making the item available for sale.
The costs incurred by carrying inventory include storage, insurance, investment, obsolescence, damage, deterioration, and the opportunity cost of the funds tied up in inventory. Some perishable items have remarkably short life-cycles measured in days, while many fashion items and high-tech consumer products have life-cycles measured in months. Such items have extremely high carrying costs due to their inherent obsolescence.
In the illustrative examples, I will assume an inventory carrying charge equivalent to 2.5% per month (30% per annum) of the inventory value, e.g., a $5.00 item will incur a monthly carrying charge of $0.125 ($1.50 per annum). This may seem relatively high, but capital alone is typically worth 10%, and the life-cycle of high fashion/technology and perishable goods can frequently be measured in weeks. For a good introduction to calculating inventory carrying costs, I suggest the Timme and Williams-Timme (2003) paper "The Real Cost of Holding Inventory."

Exhibit 1 defines the traditional safety stock (SS) calculation. The key factors are:

- the safety factor (k), which increases with the desired service level,
- the variance (standard deviation) of the forecast error (σ), which is related to the size of the mean absolute error of the forecasts (MAE), and
- the lead-times in reviewing and replenishing product (R + L).

Exhibit 1. Safety stock calculation

$SS = k\,v\,\sqrt{R + L}$ where

SS = Safety stock in units

k = Safety factor from a table of $v \approx 1.25$ * MAE (mean absolute error)

v = An approximation of the standard deviation, or variability, of forecast error

$\sqrt{R + L}$ = The square root of the review period (R) and the replenishment leadtime (L), in months

The review period is the frequency with which the purchasing (replenishment) decision is made. Often this decision takes place monthly as part of the Sales & Operations Planning (S&OP) cycle.

The correspondence between the safety factor and the desired service level is traditionally determined from a table of the Normal Distribution, excerpts of which are shown in Table 1.

Table 1. Normal distribution values

Service Level % (P1)	*k*	Unit Normal Loss Function Gu(k)
99.98%	3.60	0.00003911
99.93%	3.20	0.0001852
99.74%	2.80	0.0007611
99.18%	2.40	0.00272
97.73%	2.00	0.008491
94.52%	1.60	0.02324
88.49%	1.20	0.0561
78.82%	0.80	0.1202
65.54%	0.40	0.2304
50.00%	0.00	0.3989

P1 is the probability of no stock-out occurring during the replenishment cycle. The table also provides the unit normal loss function, which we will later use to calculate the expected shortages per replenishment cycle. A complete table can be found in Silver, p. 724.

To illustrate the safety stock calculation, consider the 12 months of data plotted in Figure 1. Here we see the sales history of a product and the forecasts produced 2 months before the actual sales were known. By subtracting the forecast from the actual sales value for each month and ignoring any negative (minus) signs, we get the absolute error for each month. The average of the absolute forecast errors, the Mean Absolute Error (MAE) over the 12 months, is 10 units.

Contact Details

Dr Peter Catt
Supply Chain Solution Manager

cell: 027 542 2053
ddi: 09 571 7105
peter.catt@soltius.co.nz

www.soltius.co.nz

Figure 1. Sales Forecast and Absolute Errors

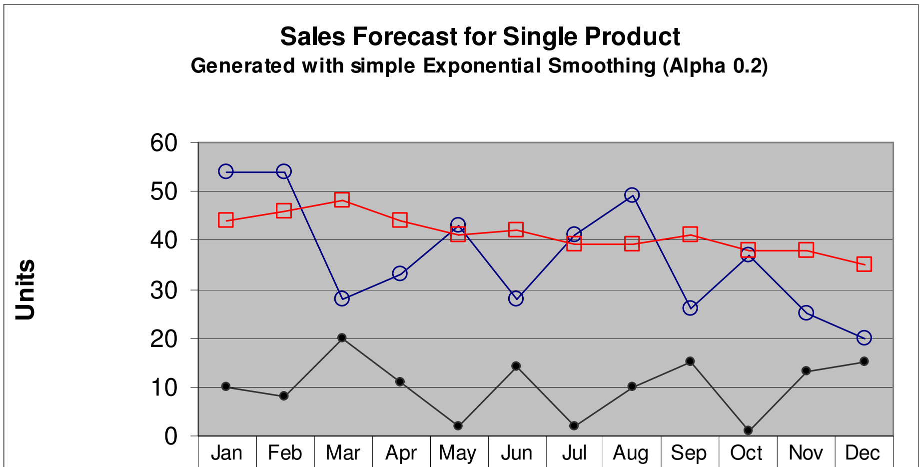

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Sales	54	54	28	33	43	28	41	49	26	37	25	20
Forecast	44	46	48	44	41	42	39	39	41	38	38	35
Absolute Error	10	8	20	11	2	14	2	10	15	1	13	15

The key inputs to the safety stock calculation are:

[1] Mean Absolute Error: MAE = 10 units/month
[2] Review Period + Lead-Time: R+L = 2 months
[3] Service Level, set at approx. 98%, implies a safety factor, k = 2.0 (See Table 1).
Using the formula for the safety stock in Exhibit 1,

SS = (safety factor * 1.25 * MAE * the square root of the (review period + lead-time))
SS = 2.0 * 1.25 * 10 * SQRT(2) = 35.36 units

So with a mean absolute forecast error of 10 units and a combined review period and lead-time of 2 months, we need to hold around 36 units (always round up) of stock to maintain a service level of 98%.

Stock-Outs and the Cost of Lost Sales

For retail companies, costs of inadequate stock can be punitive. In their paper "Stock-Outs Cause Walkouts," Corsten and Gruen (2004) report results of a survey of 71,000 retail consumers worldwide. They found that if the desired item is not in stock (i.e. a stock-out), 31% of the customers will leave the store to buy it elsewhere, while another 9% will choose not to make the purchase at all. Their study also found that world-wide stock-out rates sit at approximately 8%.

The longer-term effects of stock-outs are difficult to quantify; however, research has indicated that customer loyalty and hence the likelihood of repeat business, is also diminished (Schwartz, 1968). In the case of a retailer it appears unrealistic to assume that a stock-out will lead to a 100% loss in margin, as the customer may accept a substitute for the original product. It is more reasonable to assume that stock-outs have differing margin-loss impacts depending on situational factors, such as the ease of substitution for the product in question. The Corsten and Gruen study does provide some guidelines for retailers. In my illustrative example, I have set the percentage of lost product margin, B5 (using the notation of Silver), at 50%.This percentage reflects the sum of the 40% figure reported in the Corsten and Gruen study plus a further 10% to account for the ongoing loss of customer loyalty (future margin). I will also assume that

Contact Details

Dr Peter Catt
Supply Chain Solution Manager

cell: 027 542 2053
ddi: 09 571 7105
peter.catt@soltius.co.nz

www.soltius.co.nz

the product margin per unit is $2.50. So each unit of sales lost incurs an opportunity cost of 50% of $2.50 or $1.25.

Exhibit 2, presents the formula for the calculation of the volume of lost sales. The calculation has two components and these are multiplied together:

- The standard deviation of the forecast error over the review period and replenishment lead-time, = 1.25 * MAE * SQRT(R+L).
- A function that provides the statistical probability of lost sales at the desired service level (See Table 1, Unit Normal Loss Function).

The unit normal loss function, Gu(k), can either be calculated using standard Excel® functions (Silver, p. 735) or taken from a table of normal distribution functions (Silver, p. 724).

[1] Mean Absolute Error: MAE = 10 units/month
[2] Review Period + Lead-Time: R+L = 2 months
[3] Service Level, set at approx 98%, implies a safety factor, k = 2.0, and unit normal loss function, Gu(k) = 0.008491.

VLS = 1.25 * 10 * SQRT(2) * 0.008491 = 0.15 units per month

Table 1 shows that, at k = 2.0, the service level or probability of no stock-out is 97.73%. However, the standard deviation of forecast errors over the replenishment cycle (1.25 * 10 * SQRT(2) = 17.68) times the normal loss function (0.008491) indicates that only about 0.8% of the forecast error volume will be lost sales.

Exhibit 2. Volume of Lost Sales (VLS)

$VLS = \sigma \sqrt{R + L} * Gu(k)$ where

$\sigma \sqrt{R + L}$ = the standard deviation of forecast errors, in units, over the replenishment cycle (See Exhibit 1)

$Gu(k)$ = the unit normal loss function used to calculate shortages per replenishment cycle

We can now determine the lost sales margin (LSM) as the product of [1] percentage charge for lost margin (50%), [2] product margin ($2.50 per unit), and [3] lost sales per month (0.15).

LSM = (50%*$2.50) * 0.15 = $0.19 per month

Calculating the CFE

So far we have taken a single sales forecast with an MAE of 10 units, calculated the necessary safety stock to achieve our desired service level of 98%, and worked out our lost sales margin. Now we need to bring these elements together and add stockholding costs to get a comprehensive cost of forecast error.

My formulation for the cost calculation, shown in Exhibit 3, is a modification of that developed by Silver (p. 263). The CFE formula utilizes product margin to ascertain the cost of lost sales. Silver had used unit cost, but lost margin better reflects the opportunity cost of poor service.

To illustrate the CFE calculation, let us assume:

Contact Details

Dr Peter Catt
Supply Chain Solution Manager

cell: 027 542 2053
ddi: 09 571 7105
peter.catt@soltius.co.nz

www.soltius.co.nz

[1] Mean Absolute Error: MAE = 10 units/month

[2] Service Level = aprox 98%, implying a safety factor, k = 2.0, and unit normal loss function, Gu(k) = 0.008491

[3] Review Period + Lead-Time: R+L = 2 months

[4] Unit Cost: v = $5.00

[5] Inventory Carrying Charge: r = $0.125

[6] Lost Sales Margin per unit: B5 = 50%

[7] Product Margin in $/unit: mp = $2.50

[8] Period Multiplier to convert from monthly to annual cost of forecast error: P = 12

Here are the steps in performing the CFE calculation:

Step 1: Estimate the safety stock using the formula in Exhibit 1.
SS = 2.0 * 1.25 * 10 * SQRT(2) = 35.36 units

Step 2: Multiply the safety stock by the monthly inventory carrying charge (per unit of safety stock).
35.36 units * $0.125 = $4.42 per month (holding cost of safety stock)

Step 3: Calculate the expected volume of lost sales (VLS) due to stock-outs using the formula in Exhibit 2.
VLS = 1.25 * 10 * SQRT(2) * 0.00849 = 0.15 expected units per month of lost sales

Step 4: Calculate the lost sales margin as the product of percentage charge for lost margin (50%), product margin ($2.50 per unit), and lost sales per month (0.15) = $0.19 per month
LSM = (50% * $2.50) * 0.15 = $0.19 per month

Step 5: Add the holding cost of safety stock and lost sales margin, then multiply by 12 to annualize.
CFE = ($4.42 + $0.19) * 12 = $55.32 per annum

Exhibit 3. Annual Cost of Forecast Error

$$CFE = (SS * v * r + \frac{B_5 * m_p * \sigma\sqrt{R+L} * G_u(k)}{R}) * P \quad \text{where}$$

CFE = annual cost of forecast error in dollars

SS = safety stock in units

v = unit cost in \$/unit

r = inventory carrying charge in \$/month

B_5 = lost sales margin per unit short (the dollar penalty we apply for stockouts)

m_p = product margin in dollars

$\sigma\sqrt{R+L}$ = standard deviation of forecast errors, in units, over the replenishment cycle

$G_u(k)$ = the unit normal loss function used to calculate shortages per replenishment cycle

P = period multiplier to convert from months to year

Contact Details

Dr Peter Catt
Supply Chain Solution Manager

cell: 027 542 2053
ddi: 09 571 7105
peter.catt@soltius.co.nz

www.soltius.co.nz

The result in Step 5 is the cost of forecast error (CFE) for the item in question. By breaking down the two major components of the $55.32 CFE, we can see that $53.04 (96%) of the cost is attributed to the safety stock component while the remaining $2.28 (4%) consists of lost margin due to stock-outs. In this example, the cost of maintaining a relatively high service level is significant, particularly if one considers the total cost over all products.

Step 6: Repeat the process for every other item to derive the aggregate cost of forecast error (aggregate CFE).

Determination of Optimal Safety Stock

Our illustration of the CFE calculation assumed a 98% service level, which, given a particular replenishment cycle (review and lead-time equal to 2 months), resulted in a certain level of safety stock (36 units). But how can you decide what the appropriate safety stock should be? One way is to apply the CFE calculation across a range of service level parameters (from Table 1). A plot of the results will be a cost curve showing CFE as a function of service level. The minimum point on the curve tells us the service level associated with the minimum cost of forecast error, i.e. the lowest combined cost of safety stock and lost sales.

Figure 2 shows the cost curve resulting from our illustrative example. We can see that the lowest cost of forecast error is achieved at a service level of approximately 88%. Our arbitrarily chosen service level of 98% resulted in an annual CFE of $55.32. The optimum service level of 88% resulted in an annual CFE of $46.70, providing a potential cost reduction of $8.62 (a 16% decrease) on this one item.

However, the "strategic availability" of products should not be underestimated; for example, product families may contain some low-margin items that impact the sales of the entire family. In addition, poor product availability (relative to competitors) is likely to seriously harm customer loyalty.

Figure 2. Annual Cost of Forecast Error At Varying Service Levels

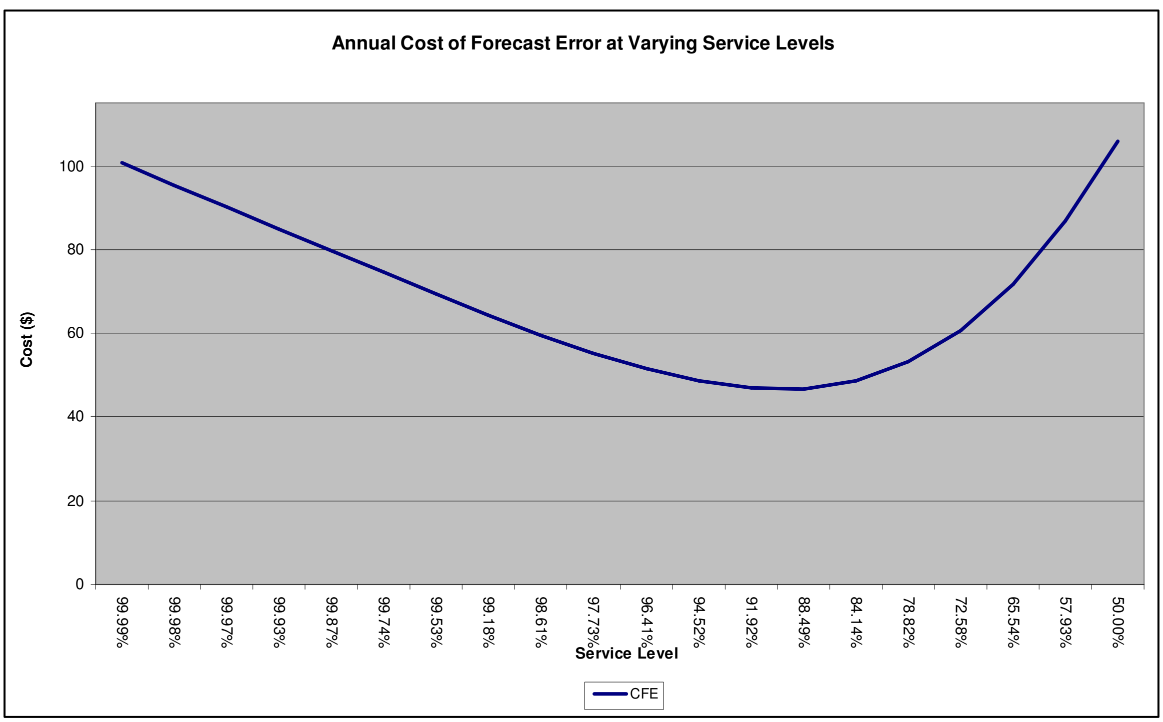

Contact Details

Dr Peter Catt
Supply Chain Solution Manager

cell: 027 542 2053
ddi: 09 571 7105
peter.catt@soltius.co.nz

www.soltius.co.nz

Conclusions

The ability to quantify the costs associated with forecast error yields an appreciation of the costs and benefits of a desired service level. It also allows an objective comparison between competing forecast methods.

Statistical measures of forecast accuracy do not make cost trade-offs explicit and hence need to be augmented with financial figures. Without knowing the costs associated with a forecast error, one cannot determine the acceptability of the available methods.

Consider downloading the necessary data from your ERP system to an Excel® spreadsheet for calculation. Better yet, consider developing a template within your business intelligence or ERP application.

References

Corsten, D. & Gruen, T. (2004). Stock-outs cause walkouts, *Harvard Business Review* (May, 2004), 26-28.

Flores, B., Olson, D. & Pearce, S. (1993). Use of cost and ac¬curacy measures in forecasting method selection: A physical distribution example, *International Journal of Production Research*, 31(1), 139-160.

Roberts, S. & Whybark, D. (1974). Adaptive forecasting tech¬niques, *International Journal of Production Research*, 12(6), 635-645.

Schwartz, B. (1968). A new approach to stockout penalties, *Management Science*, 12 (12), 538-544.

Silver, E., Pyke, D. & Peterson, R. (1998). *Inventory Management and Production Planning and Scheduling* (3rd ed.), New York: John Wiley & Sons.

Timme, S. & Williams-Timme, C. (2003). The real cost of holding inventory, *Supply Chain Management Review*, 7(4), 30-37.

Contact Details

Dr Peter Catt
Supply Chain Solution Manager

cell: 027 542 2053
ddi: 09 571 7105
peter.catt@soltius.co.nz

www.soltius.co.nz