WHITE PAPER

THE POWER TO KNOW®

Fundamental Issues in Business Forecasting

Content for this white paper was provided by Michael Gilliland and first appeared in his article "Fundamental Issues in Business Forecasting" in the Summer 2003 issue of *Journal of Business Forecasting*

Table of Contents

Introduction

Forecasting is a difficult and thankless endeavor. When accuracy is not quite where everyone wants it to be, we react by making significant new investments in technology, process and people. Unfortunately, investment in the forecasting function is no guarantee of better forecasts. There are often fundamental issues that affect an organization's ability to forecast accurately. Until those issues are recognized and addressed, further investment in the function may be wasted. This white paper identifies several fundamental issues that should be of concern to organizations creating a new forecasting function and to those struggling to improve the one they already have.

What is demand?

Perhaps the most fundamental question of all is, "What are we trying to forecast?" The usual answer is that we are trying to forecast customer "demand," with demand defined as "what the customers want and when they want it." A good forecast of demand, far enough into the future, allows the organization to invest only in the facilities, equipment, materials and staffing that it needs. This usual definition is not problematic until we try to operationalize it – that is, when we start to describe the specific, systematic way to measure it.

If customers place orders to express their "demand," and if the organization serves its customers perfectly by filling all orders in full and on time, then we have our operational definition. In this case, demand = orders = shipments. If both order and shipment data are readily available in the company's system, then we have the historical demand data, which we can use to feed our statistical forecasting models.

Unfortunately, few organizations serve their customers perfectly. As such, orders are not a perfect reflection of true demand. This is because when customer service is less than perfect, orders are subject to all kinds of gamesmanship. Here are a few examples:

1. An unfilled order may be rolled ahead to a future time bucket.
2. If shortages are anticipated, customers may artificially inflate their orders to capture a larger share of an allocation.
3. If shortages are anticipated, customers may withhold orders or direct their demand to alternative products or suppliers.

In the first example, demand (the order) appears in a time bucket later than when it was really wanted by the customer. Rolling unfilled orders causes demand to be overstated – the orders appear in the original time bucket, and again in future buckets, until the demand is filled or the order is cancelled.

In the second example, the savvy customer (or sales rep) has advanced knowledge that product is scarce and will be allocated. If the allocation is based on some criterion such as "fill all orders at X percent," the customer simply over-orders and ultimately may receive what he or she really wanted in the first place.

The third example not only contaminates the use of orders to reflect true demand, but it can also cause significant financial harm to your business. If you are in a situation of chronic supply shortages (due to either supply problems or much higher than anticipated demand), customers may simply go elsewhere. Customers may truly want your product (so there is real demand), but it won't be reflected in your historical data because no orders were placed. While orders are often perceived as "equal to or greater than" true demand, this third example shows that what is ordered may also be less than true demand.

As with orders, the use of shipments to represent demand has a number of potential problems. Shipments are often perceived as "equal to or less than" true demand. Thus, shipments and orders are thought to represent true demand's lower and upper bounds. But, as we see in example 3, orders can be lower than the true demand. Furthermore, by example 1, shipments can actually be greater than true demand in a particular time bucket. (This would occur when an unfilled order is rolled ahead into a future time bucket and then filled. In this situation the shipment occurs later than the true demand and inflates demand in the time bucket in which it is finally shipped.)

Shipments have an advantage over orders in that far fewer games can be played to manipulate the numbers. In an organization with generally good customer service (say, an order fill rate greater than 98 percent), then shipments are probably *good enough* to represent true demand.

A more complicated (but not necessarily better) operational definition of true demand can be constructed by some hybrid of orders and shipments. Examples include:

1) Demand = (Shipments + Orders) / 2

2) Demand = Shipments + Incremental Shortages

3) Demand = Shipments + Latest Shortages

The first case simply defines demand as halfway between orders and shipments; it assumes half of the shortages represent legitimate demand. If the order is 120 and the shipment is 100, then demand = 110.

The second case avoids over-counting repeat shortage rollovers by only adding increases in shortages to shipments. Thus, if the shortage in time period t is 20, and the shortage in period t+1 is again 20, then demand = shipment for period t+1 (the shortage amount, 20, did not increase from prior time period). If the shortage in period t+2 is 25, the demand in period t+2 is shipment + 5 (because there were an incremental five units of shortages from 20 to 25).

The third case also avoids over-counting repeat shortages by including in demand only those shortages still showing at the end of the time bucket. In this case, the demand for a month will include all shipments of that month + unfilled orders of the last week only. If, for example, shortages in a four-week month were 10, 20, 40 and 30, the total demand for the month would be shipments + 30 (the last week's shortages). Table 1 illustrates various demand definitions over a one-month period.

Table1

Illustration of Different Definitions of Demand

Week	**1**	**2**	**3**	**4**	**Month Total**
Orders	50	50	60	60	220
Shipments	50	40	55	40	185
Shortages		10	5	20	35
Incremental shortage		10		15	25
Latest shortage				20	20
Demand = (Shipments + Orders) / 2=(185+220) / 2=202.5					
Demand = Shipments + Incremental Shortages = (185+25) = 210					
Demand = Shipments + Latest Shortages = (185 +20) = 205					

To summarize, developing an operational definition of demand that fits your organization is a serious problem. Fortunately, for purposes of forecasting, it is probably sufficient to capture something *close enough* to this nebulous "demand" concept. Given that typical forecast error is 25 percent, 50 percent or more, a demand proxy that is within a few percentage points of true demand is probably adequate. There is no need to waste resources seeking perfection.

What to forecast?

An operational definition of demand leads directly to the topic of "what should we be forecasting?" Clearly we want to drive the supply chain with a forecast of future customer demand – but we have to be careful. It is important to distinguish "unconstrained" demand from the demand we actually expect to fulfill subject to supply-side constraints.

The supply side of an organization needs visibility to unconstrained customer demand – production and inventory planners need to be aware of what the customer really wants and when. However, once constraints are identified, it is proper to issue a shipment forecast – a best guess at what really is going to happen. Known constraints must be communicated to the sales organization. When a shortage is anticipated, the customer should be contacted and the demand redirected to a future date (i.e., when the demand can be fulfilled) or to alternative products. It is a failure of customer management to solicit orders that are *known in advance* to be unfillable.

Financial projections should be made from the (constrained) shipment forecast. Also, performance metrics (such as Mean Absolute Percent Error and bias) should be based on the shipment forecast. A forecast of shipments can be compared to actual shipments. Forecasting performance metrics should not be based on orders (or any version of "demand" that includes orders in its operational definition). This is because the proper organizational response to supply constraints is to redirect customer demand to alternative products or time frames – not seek and process orders that we know in advance cannot be filled. Orders are not a reliable component of performance metric calculations.

Performance measurement

Does your organization have a clue about how effectively it forecasts? Perhaps not. Commonly used metrics such as MAPE show the ultimate result of the forecasting process but give no indication of how efficient the organization was in achieving that level of forecast accuracy. Also MAPE, by itself, does not tell the organization whether other methods would have been equally or more accurate with less management effort.

Due to personal biases, company politics, lack of training and tools, or sheer incompetence, many (if not most) management efforts fail to improve the forecast – and may even make it worse! Traditional application of process performance metrics such as MAPE does not address this issue. By failing to consider the "forecast value added" (FVA) each step of the way by each participant in the forecasting process, the traditional approach to performance measurement misses a potential source of significant process improvement.

The most basic exercise is to compare the results of your newly structured forecasting process (using MAPE, MAD or other metric) to the results you *would have* achieved with a *naïve* forecasting method such as a random walk or moving average. The difference between your process results and the results of a naïve method is the value added by your efforts.

FVA analysis should be applied to each step and each participant in the process by comparing the forecast output of each step and each participant. For example, the naïve forecast (e.g., a random walk) should be compared to the statistical forecast (perhaps an exponential smoothing or ARIMA model). The statistical forecast can be compared to the sales force rollup, the marketing override, the forecaster analyst's override, the consensus forecast or the *evangelical* forecast (management's wish number). Similarly, a forecast based on point-of-sale data could be compared to the forecast based on orders or shipments. By evaluating the ultimate performance (MAPE, bias, etc.) at each stage, we can identify what helps, what hurts and what is simply a waste of effort.

Eliminating the non-value-adding steps and participants will help you to make the forecasting process more efficient, achieving results with fewer organizational resources. Identifying and eliminating those steps and participants that actually make the forecast *worse* will also improve forecast accuracy, again with fewer organizational resources.

Looking purely at MAPE as the end result of the forecasting process can create the wrong impression. It is a serious misconception to think, “If only we hired more analysts, if only we engaged more participants in the process, if only we had a bigger computer and more sophisticated software, then our forecasting problem would be solved.” In fact, it is more likely that results can be improved by doing less! In forecasting practice as in medical practice: “First, do no harm.”

Benchmarking performance

FVA analysis highlights a potential unfairness in forecasting performance benchmarking and comparison. One must be cautious in interpreting benchmarking results, as there are no standardized benchmarks of forecast errors, and they are usually based on self-reported rather than independently audited performance. This means we have no assurance that the reporting companies are using the same methods to gather data and measure error. We also have no assurance that companies have equally “forecastable” demand patterns (e.g., mature and stable product lines versus volatile, promoted items or lots of new products). This issue applies more generally to any comparisons of forecasting performance between companies, product lines, geographic regions, individual forecasters, etc.

A simple example, shown in Table 2, illustrates this point. If Company A has the best forecasting performance based on MAPE, does it mean that it truly has the most effective forecasting process? Not necessarily. One way to get the answer is to use FVA analysis. Suppose we computed the MAPE each company *would have* achieved by using the naïve forecasting model:

Table2

Comparison of Forecast Value Added

Company	MAPE	Naïve MAPE	FVA
A	20%	10%	-10%
B	30%	30%	0%
C	40%	50%	10%

FVA analysis reveals that Company A would have had more accurate forecasts had it just used the naïve model. All of the costs and efforts that went into the Company A's process only made the forecast worse! Company B also had a non-value-adding process, because it achieved the same MAPE it would have achieved by simply using the naïve. Only Company C, which actually had the most inaccurate forecasts, employed a forecasting process that added value. Only C's investment in the forecasting function is actually providing any benefit.

This example reveals the danger of uncritical acceptance of benchmark data. An organization may have "best-in-class" forecast accuracy simply because it has easy-to-forecast demand, not because its forecasting systems or processes merit distinction. Although FVA information is not readily available from public sources, it provides a better comparison than a simple error metric such as MAPE.

Evangelical forecasting

Forecasting is a highly visible and highly politicized element of the business environment. While an organization's forecast should represent an "unbiased best guess" at what really is going to happen, the forecast is more often an expression of the organization's targets or wishes. The most significant forecasting mistake an organization can make is to build plans around what it *wants* to see happen rather than with what it really *believes* will happen. If what your organization calls "forecasting" is simply an exercise to make the numbers match some predetermined financial objective, then stop wasting your time. You don't need professional forecasters to do this; just hire clerical support that knows how to do the arithmetic.

The primary purpose of a forecast is to drive the supply chain. The financial version of the forecast – for financial planning, analysis and reporting – should be derived from the supply chain forecast. Meeting customer demand with responsible management of company resources is a very serious objective. This objective should not be encumbered or contaminated by financial objectives that ignore conditions of the marketplace. Whether the business is consumer goods, industrial products or professional services, an honest and unpoliticized forecast gives the organization its best chance to meet customer service requirements with the appropriate investment in capacity, staffing and inventory.

Evangelical forecasting is an approach where the forecast is given "from above." In this kind of environment a charismatic owner, CEO, general manager or other executive determines the forecast, which may be nothing more than an expression of the revenue target. This forecast is established in dollars (rather than units), at some high level of aggregation (such as at total corporate, by brand, or by sales territory) and for a broad time frame (such as year or quarter). In this case, the job of a forecaster is merely to adjust product volume and mix to achieve the evangelical dollar target and then send the product unit forecast to the supply chain.

Evangelical forecasting has an advantage in that you don't need a cumbersome consensus or Sales and Operations Planning process to get everyone to agree upon a number. However, this approach can be very demoralizing for the forecasting department, when all their thoughtful modeling and analyses are brusquely overwritten by executive decree.

Forecasting is hard enough to begin with. At the very least, the supply chain should be driven by the organization's honest guess of future demand. The worst impact of evangelical forecasting, or any forecast derived from financial directives rather than by realities of the marketplace, is that it drives the supply chain with the wrong signal.

Organizational practices and demand volatility

Demand volatility has a significant impact on our ability to forecast accurately. The good news is that reducing the variability of a demand pattern is an almost sure way to improve the accuracy of your forecasts. Even better news is that reduced demand variability will get you better forecasts for free, without any change whatsoever in your process, systems or people.

Many organizational practices (such as sales contests or the quarter-end "push") increase volatility and make demand more difficult to forecast. Fortunately, management has control over these organizational practices and can change them.

Product consumption (e.g., consumers buying from retail stores) can be much less volatile than the shipment of product into the retailers. It is easy to compare the "inherent" volatility of consumption to the "artificial" volatility of shipments we ourselves encourage by misguided organizational practices. The inherent volatility in consumption can be measured by the coefficient of variation (CV) of point-of-sale data. This is the volatility of consumer pull from the stores. When compared to the CV of shipments (which is often two or three times higher), the difference is the "artificial" volatility caused by organizational practices.

Perhaps the surest way to get better forecasts is to make the demand forecastable. This can be achieved by re-engineering or eliminating those organizational practices that encourage customers to order in spikes or erratic patterns. Encouraging smooth and stable order patterns lowers supply chain costs and virtually guarantees more accurate forecasts with less management effort.

Improving the process

The above discussion covered many fundamental, yet frequently overlooked, areas of the forecasting function. Addressing these items can lead to immediate improvement in forecasting performance and should be done in conjunction with any new investment in technology or people. Before getting started, a few simple tests can help determine both opportunities for improvement and the organization's readiness to make those improvements:

- **Data/systems infrastructure**. Forecasting requires systems and data. Ask if your IT department can prepare a master file of items (with their attributes), a master file of customers (with their attributes), and a clean file of historical orders, shipments, forecasts, inventory, production and POS (for consumer products companies). If so, you have the minimum necessary data infrastructure in place. If not, a new forecasting system has little chance to succeed without the data to drive it.
- **Demand volatility**. Utilizing the historical shipment and POS data, determine the inherent volatility of consumption and the artificial volatility caused by your organizational practices. (For example, measure the CV of POS and the CV of shipments over the past 52 weeks.) Significant artificial volatility indicates opportunities to smooth shipments and thereby get better forecasts. This may call for re-engineering those practices that cause increased volatility in shipments.
- **Forecast Value Added**. For a quick and dirty analysis, compare your forecast accuracy over the past year to the accuracy you *would have* achieved by just using a simple method such as a random walk or moving average. Many organizations find that a simple method would have done better than all their elaborate systems and processes. Use FVA analysis to identify and eliminate the bad practices and wasted efforts. Focus on process efficiency and not just forecast accuracy, because efficiency breeds accuracy.

In short, when making investments in forecasting technology and people, make sure you've solved the basic stuff first. Big investments in the forecasting function – without first addressing fundamental issues – may yield little of the improvement you are seeking.

THE POWER TO KNOW.®

SAS Institute Inc. World Headquarters +1 919 677 8000
To contact your local SAS office, please visit: **www.sas.com/offices**