WHITE PAPER

What Management Must Know About Forecasting

Contents

The Forecasting Dilemma

Forecasts never seem to be as accurate as we would like, or need, them to be. As a result, there's a temptation to throw money at the problem in hopes of making it go away. There are plenty of consultants and software vendors willing to take that money in exchange for lots of promises, but are these promises ever fulfilled? How many organizations are you aware of – perhaps even your own – that have thrown thousands or even millions of dollars at the forecasting problem, only to end up with the same lousy forecasts?

The questions boil down to:

- Why do forecasts always seem to be wrong and sometimes terribly wrong?
- Is there anything you can do about it?

This white paper explores why forecasting is often poorly done and offers suggestions for improving it. Because forecasting is a core capability of AI, it also examines the emergence of AI and machine learning to enhance traditional time-series forecasting methods.

Why Are Forecasts Wrong?

There are at least four reasons why your forecasts are not as accurate as you would like them to be.

The first reason is unsuitable software – software that doesn't have the necessary capabilities, has mathematical errors or uses inappropriate methods. It is also possible that the software is perfectly sound but due to untrained or inexperienced forecasters, it is misused.

The second reason is untrained, unskilled or inexperienced forecasters who exhibit behaviors that affect forecast accuracy. One example is overadjustment, or as W. Edwards Deming put it, "fiddling" with the process. This happens when a forecaster constantly adjusts the forecast based on new information. Research has shown that much of this fiddling makes no improvement in forecast accuracy and is simply wasted effort.[1] Forecasting should be objective and scientific.

The third reason for forecasting inaccuracy is contamination by the biases and personal agendas of the process participants. Instead of presenting an unbiased best guess at what is going to happen, the forecast comes to represent what management wants to have happen – no matter what the marketplace is saying. Forecast value added (FVA) analysis, described later, can identify these sorts of biases and help streamline the forecasting process.

Finally, bad forecasting can occur because the desired level of accuracy is unachievable for the behavior you are trying to forecast. Consider the example of calling heads or tails in the tossing of a fair coin. It doesn't matter that you may want to achieve 60, 70 or 90 percent accuracy. The reality is that over a large number of tosses, you will only be right half of the time and nothing can change that. The nature of the behavior determines how well you can forecast it – and this applies to demand for products and services just as it does to tossing coins.

Inadequate, Unsound or Misused Software

A common mistake in bad or misused software is choosing a forecasting model based solely on the model's "fit to history" (often referred to as "best fit" or "pick best" functionality). The software provides (or the forecaster builds) several models so you can evaluate them against recent history. The model that has the best fit to history is selected to create forecasts of the future.

In Figure 1, the history consists of four weeks of actual sales: 5, 6, 4 and 7 units. You can see these as the four dots in each graph. Let's consider four models for forecasting future sales.

Model 1 is simply the average of the four points of history, and it forecasts 5.5 units for week 7. Model fit over the four points of history has a mean absolute percent error (or MAPE) of 18 percent.

Model 2 is a least squares regression line that shows an upward trend, and it forecasts 7.2 units for week 7. It has a fit error of 15 percent over the four points of history.

Model 3 is a quadratic equation with a fit error of only 8 percent, and it forecasts 16.5 units in week 7.

Finally, Model 4 is a cubic equation that fits the history perfectly with a fit error of 0 percent. It forecasts about 125 units in week 7.

Remember, the objective is not to fit a model to history – it is to find an appropriate model for forecasting future weekly sales. Fitting a model to history is easy. Anyone can do it, and it is always possible to find a model that has a perfect fit. But having perfect fit to history is no guarantee that the model will generate accurate forecasts. In this example, bad software (or misguided forecasters) using fit to history as the sole criterion for selecting the forecasting model would have chosen Model 4. Unless you have good reason to believe that sales of this product are about to explode, Model 4 actually appears to be the worst choice of the forecasting models. Models 1 or 2 (the ones with the worst fit) are probably the most appropriate models for forecasting the future, given the limited historical information.

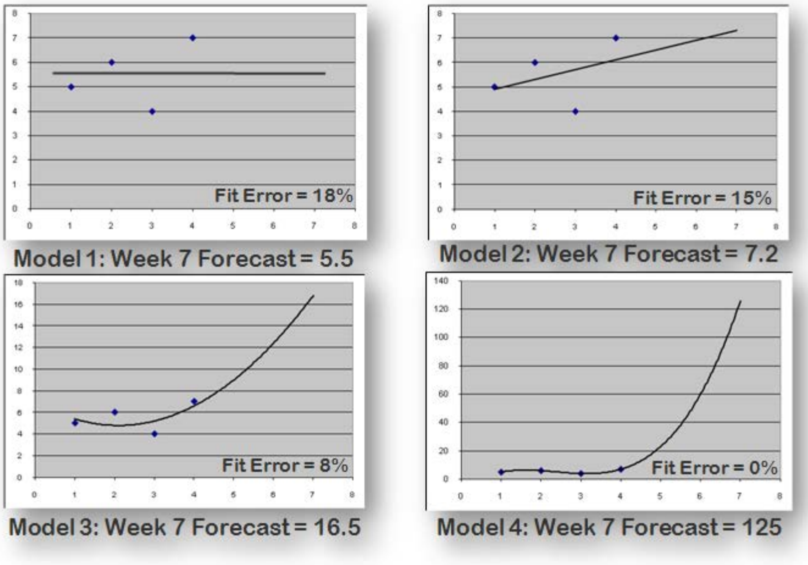

Figure 1: Confusing fit to history with appropriateness for forecasting.

Untrained or Inexperienced Forecasters

Few people enjoy the difficulties that go with being a forecaster, so few are willing to do it by choice. Do you draw forecasters from top-tier statisticians and analysts in your organization – or do you assign only your lowest ranking and least experienced employees to this task?

While master's- or PhD-level statistical modeling skills are not required to be a competent forecaster, it is necessary to have some understanding of statistical concepts like randomness and variation, as well as a good understanding of your business. Statistics alone will probably not solve the forecasting problem, but having no training in variation or process control methods can be debilitating.

Here are some questions to ask about your forecasting staff:

- Do the forecasters have any authority to set forecasts, or is all their work subject to final review by a management group that will end up forecasting whatever they please?
- If the forecasters cannot be trusted, then why are they in that role?
- If the forecaster can be trusted, why aren't they given the authority to create the forecasts?

While untrained, unskilled and inexperienced forecasters can play a role in inaccurate forecasting, forecasters (and management) can also undermine the process in other ways.

Contaminated and Politicized Forecasting Process

A significant problem is that there are points in the forecasting process that can be contaminated by the biases and agendas of process participants. In fact, the more elaborate the process with more human touch points, the more opportunity exists for these biases to contaminate what should be an objective and scientific process.

Those who have some input or authority in the forecasting process can use this to their own benefit. A sales rep at quota-setting time may try to lower expectations to get easier-to-beat quotas.[2] A product manager with a new product idea will forecast sales high enough to meet the minimum hurdles for new product development approval. Just about every participant can have a special interest of some sort, and these interests must be taken into account.

Elaborate and overly complex forecasting processes are also a poor use of organizational resources. Does each participant in your process actually make the forecast better, or can these participants be reassigned to more worthwhile activities in the organization? Most people do not like having to forecast, and have no training or skills in forecasting. Many do not add value to the process. It may make more sense to have your salespeople selling, your service people providing services and your management overseeing their parts of the business if their efforts are not improving the forecasts.

Unforecastable Demand

Finally, the problem may be that you have unforecastable demand. This doesn't mean you can't create a forecast. You can always create a forecast. Being unforecastable means that there is so much instability or randomness in your demand patterns that sophisticated methods don't help. You must either manage your operations to account for the inaccurate forecast or figure out ways to shape demand into patterns that can be forecast accurately.

Unrealistic accuracy expectations can lead to overconfidence in your forecasts and bad business decisions. A forecast is just an estimate, a best guess of what is going to happen in the future. In better software, some expression of confidence or an error range accompanies the estimate. The organization should have an appreciation of the uncertainties involved in every forecast. Knowing that demand is likely to be 100 +/- 10 units instead of 100 +/- 100 units can lead to very different plans of action.

Unrealistic accuracy expectations can also lead to inappropriate forecasting performance targets. Unachievable targets, such as correctly calling a coin toss 60 percent of the time, can discourage forecasters, who may wonder why they bother trying when the goal is unreachable. Even worse, an unreachable goal may encourage them to manipulate the metrics to achieve the necessary result.

Worst Practices in Business Forecasting

A review of some of the worst practices in business forecasting is important for a fuller understanding of the issues forecasters face. There are many things an organization can do that only make the forecast worse. There is no need to waste time and company resources repeating them.

Unrealistic Accuracy Expectations

Contrary to the claims of some consultants and vendors, there is no magic formula to ensure accurate forecasts in all situations. Sometimes nothing seems to work, but that doesn't stop you from wanting to believe unsubstantiated claims about solving your forecasting and business problems.

Ultimately, forecast accuracy is limited by the nature of the behavior you are trying to forecast. If the behavior exhibits smooth, stable and repeating patterns, then you should be able to forecast it accurately using simple methods. If the behavior is wild and erratic with no structure or stability, then there is little hope of forecasting it well, no matter how much time, money and resources you invest. The most sophisticated methods in the world aren't going to forecast unforecastable behavior, and your organization may have to adjust to that reality.

A worst practice is having unrealistic expectations and wasting resources trying to pursue unachievable levels of accuracy. So what can you do about it?

A better practice in this situation is to use a naïve forecasting model – something simple and easy to compute. The classic naïve forecast (known as the random walk or no-change model) uses the last known value as the future forecast. It forecasts no change from the latest observation. Thus, if you sold 12 last week, your forecast for this week is 12. If you sell 10 this week, your forecast for next week becomes 10, and so on.

An alternative naïve model, when you have seasonal data, is to use a seasonal random walk. With this model, your forecast is based on the observed value from the corresponding period of the prior year. For example, the forecast for April 2019 would be your actual results from April 2018.

The naïve model is a no-cost alternative: You don't need sophisticated systems, an elaborate forecasting process or an expensive staff. A naïve model will achieve some level of forecast accuracy; for example, let's say 60 percent. This 60 percent accuracy level becomes the baseline for evaluating your forecasting efforts. If your process cannot improve on a naïve model, then why bother?

Assuming Model Fit Equals Forecast Accuracy

Assume you have chosen an appropriate model for forecasting. How accurate can you expect that model to forecast?

Inexperienced forecasters, and those outside of forecasting, may assume that a model's fit to history indicates how accurately the model will forecast the future. So if the error of the historical fit is 20 percent, then the error of the future forecasts will also be 20 percent. This is an incorrect assumption. One of the dirty tricks of selling is for forecasting software vendors to only show how well their models fit your history, but never show you how well they really forecast the future.

There are a couple of reasons why forecast accuracy will almost always be worse, and often much worse, than the fit of the model to history. You may have chosen an inappropriate model; one that happens to fit the history, but does not capture the underlying mechanisms that guide the behavior. Or, you may have specified a model that correctly expresses the behavior, but then the behavior changes. Whenever you are evaluating software to purchase, or even reviewing the performance of the software you already have, make sure to focus on accuracy of the forecasts and not on the fit to history.

Inappropriate Performance Objectives

Failing to identify an accurate, reasonable forecast for your demand patterns can lead to setting incorrect performance objectives. As mentioned above, you cannot consistently guess the tossing of a fair coin correctly other than 50 percent of the time, so it makes no sense to give you a goal of achieving 60 percent accuracy. The same applies to forecasting demand. While management may want to achieve 90 percent accuracy, the nature of the demand patterns may be such that 90 percent is unachievable.

Goals are often assigned based on the level of accuracy management believes it needs. It isn't uncommon for there to be goals such as "greater than 80 percent accuracy for all products" with no consideration for whether this is reasonable.

Goals are sometimes based on industry benchmarks that purport to identify best-in-class forecasting performance. However, industry benchmarks should never be used to set forecasting objectives for your organization.

Perils of Industry Benchmarks

Benchmarks of forecasting performance are available from several sources, including professional organizations and journals, academic research and private consulting/benchmarking organizations. But there are several reasons why industry benchmarks are irrelevant in setting your own forecasting performance objectives.

1) Can you trust the data? Are the numbers based on rigorous audits of company data or responses to a survey? If they are based on unaudited survey responses, do the respondents actually know the answers or are they just guessing?

2) Is measurement consistent across the respondents? Are all organizations forecasting at the same level of granularity, such as by product, customer or region? Are they forecasting in the same time interval, such as weekly or monthly? Are they forecasting by the same lead time offset, such as three weeks or three months in advance? Are they using the same metric? It is important to note that even metrics as similar sounding as MAPE, weighted MAPE and symmetric MAPE can deliver very different values from the same data.

3) Finally, and most important, is the comparison relevant? Does the benchmark company have equally forecastable data?

Consider this worst-case example. Suppose a benchmark study shows that Company X has the lowest forecast error. Consultants and academics then converge on Company X to study its forecasting process and publish reports touting Company X's best practices.

Suppose you read these reports and begin to copy Company X's purported best practices. However, upon further consideration, the method of FVA analysis (described below) is applied and you discover that Company X had very easy demand to forecast, and it would have had even lower forecast error if it had used the naïve no-change model. In other words, Company X's so-called best practices just made the forecast worse than doing nothing and just using the naïve model!

This example is not far-fetched. Organizations at the top of the benchmark lists are probably there because they have the easiest-to-forecast demand. Many organizational practices, even purported best practices, may only make your forecast worse. In a series of stunning articles published in the journal Foresight,[3] Steve Morlidge of the consulting firm CatchBull found that 30-50 percent of the forecasts used by organizations to run their businesses were less accurate than the no-change model. Benchmarks tell you the accuracy that best-in-class companies are able to achieve, but they do not tell you whether the forecastability of their demand patterns was similar to yours. Without that information, industry benchmarks are irrelevant and should not be used to set performance objectives.

Also, objectives should not be set arbitrarily. It is unreasonable to set aspirational, blanket objectives (such as a "forecast accuracy must be greater than 80 percent") without any consideration of forecastability. If the objective is set too high, it will discourage forecasters and encourage cheating to hit the goal. If the objective is set too low, so a naïve forecast can beat it, then forecasters can sit idle and still meet the goal.

The better practice is to tie the forecasting performance objective to the underlying demand patterns, and the way to do this is to use a naïve forecast as the baseline. The only reasonable objective is to beat the naïve model (or at least do no worse) and to continuously improve the forecasting process. You improve the process not only by making the forecasts more accurate and less biased, but by making the process more efficient – using fewer and fewer resources and automating whenever possible. This is where automated forecasting software, such as SAS® Forecast Server, SAS Forecasting for Desktop or SAS Visual Forecasting, can be very effective.

Adding Variation to Demand

The forecastability of demand is largely dependent on the demand volatility. When your demand is smooth and stable, you should be able to forecast it accurately. When demand is erratic and random, you should not expect accurate forecasts.

The scatter plot in Figure 2 compares forecast accuracy (from 0 to 100 percent on the vertical axis), to the volatility of the sales pattern as measured by the coefficient of variation (CV) along the horizontal axis. It is based on one year of weekly forecasts for 5,000 SKUs at a consumer goods company. The dots show the volatility of the sales during the 52 weeks and the forecast accuracy achieved for each of the 5,000 SKUs. For SKUs with greater volatility (moving to the right in the plot), forecast accuracy tended to decrease. This is aptly referred to as a "comet chart."

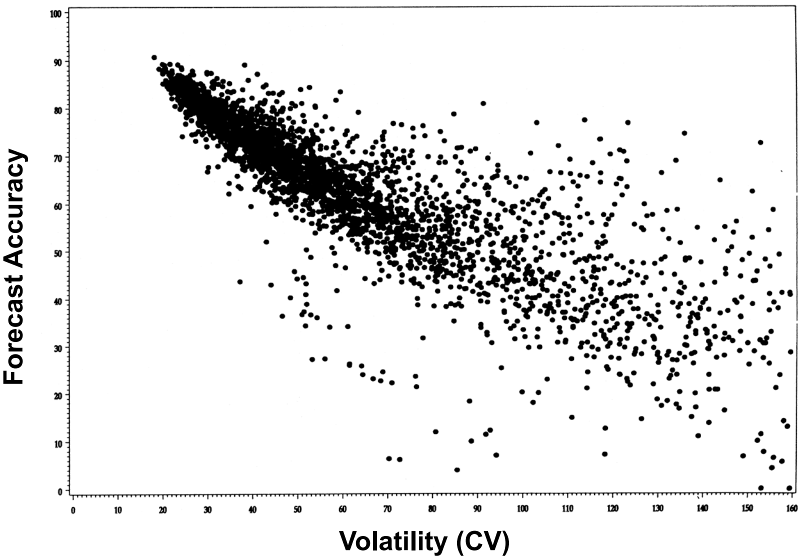

Figure 2: Forecast accuracy versus volatility.

Here is another version of the same kind of analysis, created by Rob Miller of Covidien.[4] The scatter plot in Figure 3 is similar to the one in Figure 2, but uses forecast error (MAPE) as the vertical axis instead of forecast accuracy. This again illustrates the consequences of unbridled volatility in demand patterns. Using three years of monthly data for more than 6,000 items, Miller found that 87 percent of items fell in the upper right-hand quadrant (both MAPE and CV above 50 percent), and only 5 percent of items had a CV of less than 50 percent. The chart highlights ample opportunity for volatility reduction efforts and suggests taking a critical look at organizational practices that encourage increased volatility.

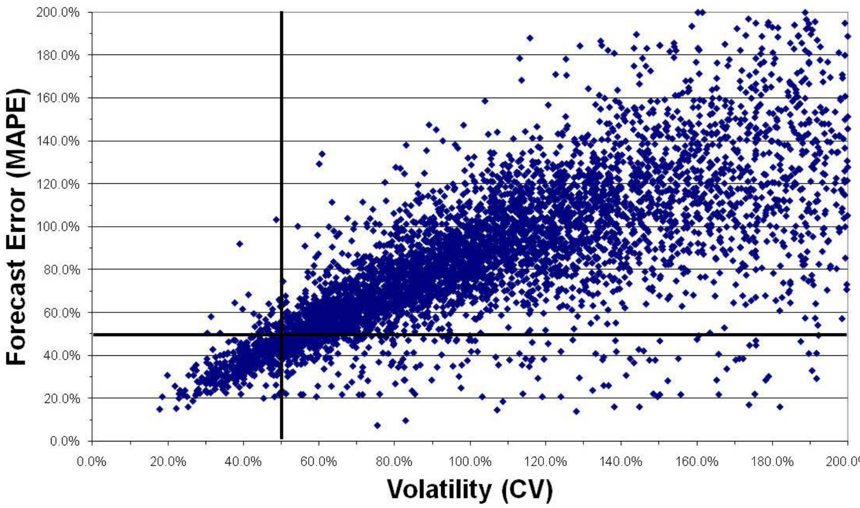

Figure 3: MAPE versus volatility.

This kind of analysis suggests that whatever you can do to reduce volatility in the demand for your products, the easier they should be to forecast. Unfortunately, most organizational policies and practices add volatility to demand rather than make it more stable – a worst practice.

Everyone is familiar with the quarter-end push, or “hockey stick” 0000 when companies do everything possible at the end of the quarter to make the sales target. Figure 4 shows shipments from a consumer goods manufacturer to retail stores. The thin line shows shipments, and you can see the big spikes at quarter end and the big drop at the start of every new quarter.

The thicker line shows consumer purchases from the retail store, and you can see it is fairly stable. You could have just forecast the mean (of about 2.4 million pounds) and the forecast would have been quite accurate.

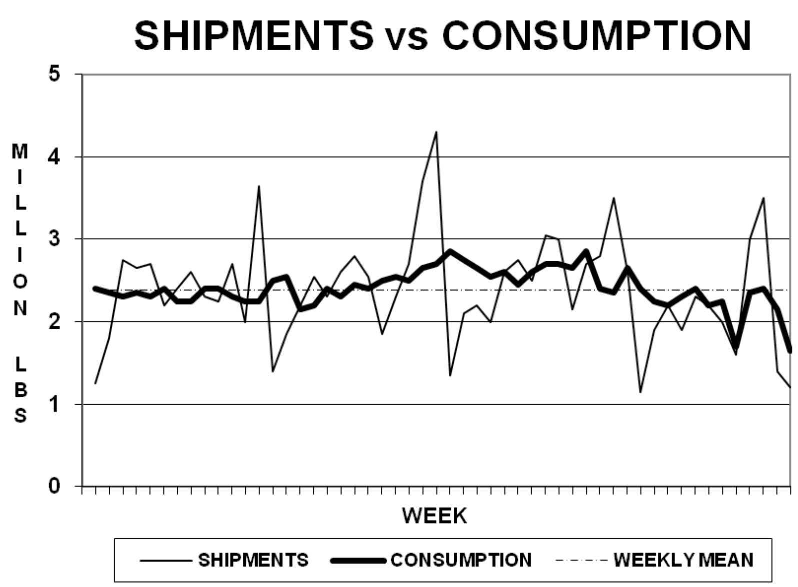

Figure 4: Shipments versus consumption.

The variation of the shipment pattern is three times the variation of the retail sales pattern. These highly erratic and hockey stick patterns are encouraged by financial practices, such as hitting the quarter-end revenue targets, and by sales and promotional practices, such as cutting prices or offering other incentives that spike the demand. In many markets, customers have been trained to wait for the end of the quarter to get the best deals.

Instead of policies that encourage volatile demand from your customers, a better practice is to remove those kinds of incentives, or create incentives that encourage smooth and stable demand. One example might be to reward the percent of weeks at or above plan, since this creates a disincentive for falling behind and then cramming all sales into the last week of the quarter.

It is important to note that, in addition to being able to forecast smooth demand more accurately, smooth demand should be easier and cheaper to service, so you can reduce costs.

It is difficult to change the way that public companies do business; however, organizations should apply these same sorts of analyses to their own data to better understand the depth of issues that are created by their policies and practices.

Almost every company has weekly or monthly sales records for the past year, and many have also kept track of their historical forecasting performance. (All of them should!) With access to this data, it is easy to use SAS to create visualizations of the relationship between volatility and the ability to forecast accurately. This can draw management attention to the difficulties and costs of doing business as usual. You may be able to re-engineer incentives to customers, rewarding them for buying in smooth patterns rather than all at once in large quantities. You may be able to re-engineer incentives to your sales force, for example, rewarding those with smooth and steady growth, rather than rewarding record weeks.

Unfortunately, few organizations have tried these changes, but there are some examples. One consumer goods manufacturer shifted the fiscal calendar by sales region, so that each region's quarter end fell on a different week. This had the effect of smoothing demand at the company level, but each region still had the same pattern as before with a big spike at their quarter end and a drop at the start of their new quarter. Because each sales region was tied to a regional distribution center, each distribution center still had to deal with the same erratic highs and lows. This approach did not address the fundamental problem.

The same company tried another approach, which was to alter incentives to the sales force. Salespeople were given extra credit for sales at the start of each new quarter and reduced credit for sales at the end of the quarter. This was clever, because it eliminated the incentive for cramming sales into the end of a quarter and encouraged the sales force to get ahead of plan early. This approach was successful in reducing overall volatility.

A final example involves a major medical supply manufacturer that sells to distributors, who then sell to hospitals, clinics and doctor's offices. The manufacturer promised three-day deliveries upon receipt of orders and made no attempts to shape demand. Volatile order patterns made for challenging forecasting, inventory management and production scheduling. Its order-fill rate was only 85 percent. This company took the proactive step of asking its customers (the distributors) what was most important to them. Receiving the shipment the day it was promised was the most important. Otherwise, temporary warehouse employees might need to be hired if the shipment came on the wrong day. The distributors agreed to, and the company met, a five-day lead time rather than a three-day lead time that was sometimes late (or early). The company was then able to schedule weekly deliveries to the distributors on a designated day. This spread shipments evenly across the week, achieving a better order-fill rate (greater than 97 percent) with less inventory and more steady production. The savings were significant and the company shared in the cost benefits by offering 1 percent discounts to the distributors when placing orders on their designated day.

Customers are so well trained to expect better deals at quarter end that they may not be happy with these sorts of changes. One way to get them on board is by letting them share in the benefits of overall reduced costs.

New Product Forecasting (NPF)

Good forecasting helps drive better decision making, and many critical decisions are needed in the release of a new product. The worst practice here is making business decisions based on the assumption that your new product forecasts are going to be highly accurate – because they probably won't be.

Since there is no historical demand data for a new product, forecasting is largely based on judgment. Often, the product advocate (e.g., product manager) provides the forecast. Almost assuredly, the forecast will be high enough to exceed internal hurdles for getting new products approved for development. When justification for the forecast is required, a common method is to refer to past products, sometimes called "like items," that are similar to the new product. This is forecasting by analogy. While this approach is legitimate, it is vulnerable to the advocate only choosing prior products that were successful. Because most new products fail in the marketplace, basing a forecast only on successful product introductions creates an unjustifiably optimistic view.

While there are dozens of methods purporting to improve new product forecasting accuracy, the important thing is being aware of the uncertainties and the likely range of outcomes. Too much confidence in the accuracy of your new product forecast can lead to risky business decisions.

Expectations of accuracy for new product forecasts should be modest and acknowledged upfront. The structured analogy approach[5] allows the organization to both statistically and visually assess the likely range of new product demand, so that it can be managed accordingly. Rather than lock in elaborate sales and supply plans based on a point forecast that is likely to be wrong, an organization can use the structured analogy process to assess alternative demand scenarios and mitigate risk.

Judgment is always going to be a big part of new product forecasting – a computer will never be able to tell you whether lime green or bright orange is going to be the hot new fashion color. But judgment needs help to keep it on track and as objective as possible. While the structured analogy approach can be used to generate new product forecasts, it is mostly of value in assessing the reasonableness of forecasts that are provided from elsewhere in the organization. The role of structured analogy software is to do the heavy computational work and provide guidance – making the NPF process as automated, efficient and objective as possible.

Forecast Value Added (FVA) Analysis

There is scant evidence that complex forecasting processes result in more accurate forecasts. In fact, just the opposite may occur. More elaborate processes, involving more human touch points, provide more opportunity for the forecast to be contaminated by the biases and personal agendas of process participants. The best way to determine failings in both process and modeling is through FVA analysis.

FVA is defined as:

> *The change in a forecasting performance metric (such as MAPE, accuracy or bias) that can be attributed to a particular step or participant in the forecasting process.*

FVA is concerned about whether process steps and participants are adding value by making the forecast better. In conducting the analysis, many organizations find process activities where FVA is negative – in other words, these efforts are making the forecast worse. (See the Morlidge articles in Foresight for more discussion of this.)

The objective of FVA analysis is to identify these non- or negative-value-adding activities and either correct them or remove them from the forecasting process. FVA analysis is about rooting out the waste and inefficiencies from forecasting efforts.

FVA analysis can be illustrated with a simple example. Consider the rudimentary forecasting process in Figure 5, consisting of demand history being fed into a statistical forecast model. The model generates a forecast that is reviewed (and possibly adjusted) by the forecast analyst.

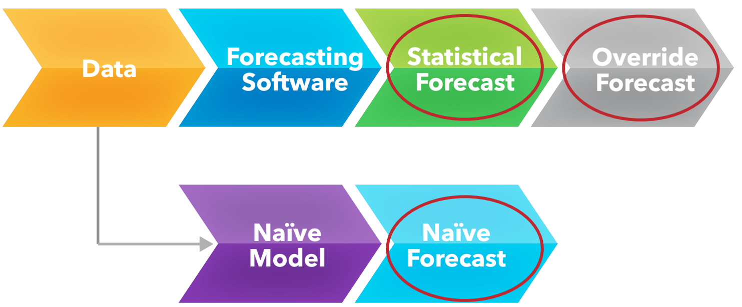

Figure 5: Simple forecasting process.

Is this process adding value by making the forecast better? To find the answer, you need to record the forecasts generated by the statistical model and the final forecast after any adjustments from the forecast analyst. You then compare these forecasts, as well as the naïve forecast, to the future demand as it becomes available.

For more information on forecast value added analysis, including case studies and step-by-step instructions for conducting the analysis, see the SAS white paper, **Forecast Value Added Analysis: Step by Step.**

Once you've gathered this data, an FVA analysis "stairstep" report looks something like Figure 6:

Process Step	Forecast Accuracy	FVA vs. Naïve	FVA vs. Statistical
Naïve Fore-cast	60%	–	–
Statistical Forecast	65%	5%	–
Analyst Override	62%	2%	-3%

Figure 6: FVA stairstep report for a simple forecasting process.

Down the left column are the process steps, which in this simple example include the statistical model and analyst override, as well as the naïve model. A naïve model is always part of FVA analysis because it shows what forecasting performance you could have achieved by essentially doing nothing and just using a naïve model for forecasting.

The second column shows the forecasting performance achieved by each step in the process. In this example, you are measuring forecast accuracy (although you could be measuring forecast error, or bias, or other metric of your choice). The naïve model achieved 60 percent accuracy, but the statistical model achieved 65 percent accuracy. After the analyst had a chance to review and adjust the forecast, accuracy fell to 62 percent.

The two right columns show the forecast value added. The statistical model added five percentage points of value compared to the naïve model, but the analyst override step had negative FVA, reducing accuracy by three percentage points versus the statistical model.

Again, forecast accuracy is largely determined by the nature of the behavior you are trying to forecast. The objective should be to efficiently develop forecasts that are as accurate as anyone can reasonably expect them to be.

FVA analysis identifies the non-value-adding activities in the forecasting process, so they can be eliminated. This allows organizations to streamline their process and redirect the non-value-adding efforts into more productive activities, such as having salespeople selling rather than forecasting.

When organizations eliminate activities that are making the forecast worse, they can actually achieve better forecasts with less effort.

The Role of AI and ML in Forecasting

Technology continues to advance, with more data, computational power and statistical sophistication available now than ever before. Artificial intelligence and machine learning are providing new approaches to longtime business problems, delivering automation and learning capabilities to emulate human tasks. There is considerable research going on – by both academics and industry data scientists – to determine the value added by AI/ML to the task of business forecasting. While these approaches have delivered positive results in many areas, it's wise to proceed with a degree of caution in forecasting.

There's a long trail of evidence showing that when it comes to forecasting, simple methods can perform better than complex methods.[6] While complex models generally provide a better fit to history – in fact, a perfect fit to history is always possible – fit to history is not a reliable indicator of the accuracy of future forecasts.

As technology advances and new forecasting methods evolve, it's essential to put new methods to the test to make sure they fulfill their promise. This is what FVA analysis allows you to do.

So far, there's limited evidence about the relative performance of machine learning methods compared to traditional statistical methods for time series forecasting,[7] but the M4 Forecasting Competition held in 2018 yielded telling results.

Among the 60 competing methods and benchmarks forecasting 100,000 time series, the six pure machine learning methods did very poorly – all failing to match a benchmark of simple time series methods. Yet the two top-performing methods combined machine learning with traditional time series models.

You can find a thorough recap of the results, analysis and commentary on the M4 in a special issue of the International Journal of Forecasting.[8] The M5 Forecasting Competition, starting in 2020, should yield even more important insights on the value of AI/ML in forecast modeling.

Another promising role for machine learning is in the area of assisted demand planning. Independent efforts at Kellogg's and in SAS Research & Development have demonstrated favorable results using ML to guide manual overrides to statistical forecasts.

This work has shown considerable promise for both reducing the time spent reviewing and adjusting statistical forecasts (by 47% in the research sample), and improving forecast accuracy (by guiding against adjustments that just made the forecast worse). While this research is still in its early stages, it illustrates another potentially valuable use of ML in forecasting.

Summary

This white paper has provided information about forecasting that is important for organizational management to understand. It described several reasons why forecasts are often wrong, and some of the organizational practices that contribute to poor forecasting. It also provided a brief introduction to the popular method of forecast value added (FVA) analysis that many organizations now use to remove waste and inefficiency from their forecasting process.[9]

Notes

1. Robert Fildes and Paul Goodwin, "Good and Bad Judgment in Forecasting: Lessons From Four Companies." Foresight: The International Journal of Applied Forecasting, Issue 8, Fall 2007, 5-10. For a more detailed account of the same research, see Robert Fildes, Paul Goodwin, Michael Lawrence and Konstantinos Nikolopoulos, "Effective Forecasting and Judgmental Adjustments: An Empirical Evaluation and Strategies for Improvement in Supply-Chain Planning," International Journal of Forecasting 25 (2009), 3–23.
2. Michael Gilliland, "Role of the Sales Force in Forecasting." Foresight, Issue 35, Fall 2014, 8-13.
3. Steve Morlidge, "Forecast Quality in the Supply Chain." Foresight, Issue 33, Spring 2014, 26-31. Several subsequent articles by Morlidge also published in Foresight.
4. Rob Miller, "Green Forecasting: A Collaborative Approach to Cleaning Polluted Processes," IBF Supply Chain Forecasting and Planning Conference in Phoenix, February 2010.
5. Michael Gilliland and Sam Guseman, "Forecasting New Products by Structured Analogy." Journal of Business Forecasting, Winter 2009, 12-15.
6. Kesten Green and J. Scott Armstrong, "Simple Versus Complex Forecasting: The Evidence." Journal of Business Research, 68, 1678-1685 (2015).
7. Spyros Makridakis, Evangelos Spiliotis and Vassilios Assimakopoulos, "Statistical and Machine Learning Forecasting Methods: Concerns and Ways Forward. PLoS ONE 13(3): e0194889. https://doi.org/10.1371/journal.pone.0194889.
8. "International Journal of Forecasting" 35:4 is a special issue containing results, analysis and commentary on the M4 Forecasting Competition.
9. For more thorough discussion on all of the topics in this white paper, see Michael Gilliland's book, *The Business Forecasting Deal* (Wiley, 2010) and also the compilation *Business Forecasting: Practical Problems and Solutions* edited by Michael Gilliland, Len Tashman and Udo Sglavo (Wiley, 2015). A concise handling of many of these same topics can be found in *The Little (Illustrated) Book of Operational Forecasting* by Steve Morlidge (Troubador, 2018), a pocket-sized compendium of essential facts and guidance for the forecasting practitioner and organizational management.

To contact your local SAS office, please visit: sas.com/offices