

## What To Look For In A Forecasting Package

When evaluating a forecasting system, a user will first need to objectively consider the essential components of an integrated package. In order to prepare you for the presentation that follows, we at AFS have identified key aspects of forecasting that are ultimately necessary in providing the most powerful and flexible system technically available.

Presented on the following screens are questions that you should consider when selecting a forecasting package. It should be understood that Autobox answers each and every one of these needs. This article was published in the Journal of Business Forecasting Methods & Systems, Volume 15, Number 3, Fall 1996.

---

### Computational Considerations

To streamline the forecast process it is important for the user to have maximum flexibility in platform selection, the analytical techniques employed and output control. External user control combined with an automated approach to model simulation empowers the user with the capability to identify the most efficient process for generating product forecasts.

### Platform Compatability

Q. How does the forecasting engine operate and perform on different platforms? Does the desktop (PC) version give the same answer as the workstation or mainframe version? Are there desktop or workstation versions?

A. The engine must be a component in the PC interactive environment and deliver exactly the same analysis as the work station version. The engine version must be capable of running on a large number of platforms.

### Transparency

Q. Does the forecast engine deliver forecasts and statistical results like R-squared, model equation, audit trail, table of forecasts, etc.? Does the engine allow these to go to user defined files?

A. The engine must allow a user to specify file names where forecasts are sorted. These kinds of files are designed for post processing purposes and are machine readable. Statistical information must also be directed to specific files for subsequent reporting and managerial summary purposes.

### Targeted output

Q. Is there an approach which permits the user to speed up the process?

A. To gain computational speed the engine should allow the user to disable/enable various tests or simply increase/decrease the maximum number of iterations allowed. The user must also be able to control various print options or the level of output verbosity. Together, these options provide significant flexibility in optimizing simulations for computational speed.

### Computational

Q. How does the forecasting engine identify and select the optimum model? Does it try a fixed set of models or does it tailor the model to the data?

A. If the user specifies a starting model the engine should perform multiple passes of necessity and sufficiency checking, possibly augmenting with intervention variables and then testing for constancy of variance. If no model is

suggested by the user, the engine should evaluate a large number of possible combinations of ARIMA/TF structure and select the one that minimizes the AIC criterion with respect to the observed cross-correlation function. This is considered the automatically identified initial model. The path is then similar as above culminating in a statistically tight model that is both parsimonious and powerful.

Attempts to determine the best model from a list is doomed from the outset as the number of possible models is infinite unless the engine is truly flexible in identifying a model. The concept of automatic model building embodies an attempt to optimize the combination of a number of model components. Thus the approach is to iteratively optimize as more avenues of model construction are evaluated.

---

### **Ease of Use**

With a user friendly forecast engine the process of model identification and product forecasting needs is automated. This alone is not sufficient in providing an engine which is easy for the user to utilize. The forecast engine should be considered a tool box which contains all that is needed to support forecasting and the planning market strategy and provides all the information which is necessary to explain the results.

To provide a complete forecasting package, the developers must be available to provide expert assistance in overcoming hurdles which surface during day to day use.

### **A planning tool**

Q. How does the forecast engine allow the user to specify alternative scenarios for the future values of cause variables? Will the engine forecast series if the user does not want to specify these values or is simply unable to identify appropriate future cause variables?

A. The user should be able to specify future values of cause variables, such as price or length of a promotion, and simply plug these values into the equation previously determined by the engine to get an assessment of future impacts. By varying these inputs an optimal combination might be identified thus suggesting a best market strategy.

### **Is the model explicable?**

Q. How does the engine deliver the equation or model parameters? Does it report the statistical tests that proved that the obtained parameters are significant? Can you explain the model to your boss? Is saving, recovery and archiving of such models permitted?

A. An equation represents the best way to predict a specified series from either its past or other series. The equation predicts the expected value for a set of specific input series. This can be very useful in assessing impact and the corresponding anticipated effect.

### **Help desk**

Q. How does the forecast engine developer provide technical support? What is the quality of the support?

A. The engine designers and developers should man the phones providing timely and competent support. To investigate their capability call the vendor and ask the help desk to key in a series that you wish to read them .... Now see what their response is. If the engine is so automatic why can't they key in some 24 or 36 numbers and get a forecast which can be read back to you while you are on the phone.

---

## **Model Robustness**

The ability of a forecast engine to deal with the wide range of unique circumstances that occur when modeling a forecast is of the utmost importance. An engine can be considered robust if it is able to appropriately identify inherent data trends and structure. Proper modeling of cause variables requires an engine to have a certain degree of artful creativity which can accurately deal with holiday demand, promotions, and price changes.

## **Temporal Structure**

Q. How does the engine deal with the issue of incorporating the optimal lead/lag structure for each input series (exogenous or cause) in the model? Is the user permitted to specify a tentative structure or does it require the user to specify this structure? How does the forecast engine deal with unnecessary structure that might have been tentatively identified or specified by the user? How does the engine deal with under-identified structure? What process is used to identify and incorporate additional lags in the input series?

A. The engine should identify a model based upon the correlation between the dependent and cause variables. This model may be over-specified and necessity testing will eliminate the superfluous structure. If the model is under-specified the omitted structure will show up in error diagnostics leading to model augmentation.

If a user has some prior knowledge of the model and/or parameters these can be used but are not required. In either case necessity and sufficiency checking deals with incorrect structure eliminating model specification bias. Be wary of engines that require you to specify the temporal structure as that indicates a weakness and an inability to be empirical in the initial model form. This flaw is referred to in econometrics as model specification bias.

## **Required Dummy Variables**

Q. How does the forecasting engine identify the following kinds of dummy variables that might be required and useful in the model?

- One time only unusual values (pulse)
- One time per period unusual values (seasonal pulse)
- Locally changing mean of the residuals (level shift)
- Dynamically changing mean of the residuals (level shift)

A. The engine should identify intervention variables that may be simply a one time pulse or may be systematic (11th week of the year for example). Consider the case where an important variable like the occurrence of St. Patrick's day in predicting beer sales has been omitted from the model. The errors from that model would contain an unusual spike every March 17th (11th week of the year) and would help identify the omitted variable. The series may have changed level and to some statistically deficient procedures this might appear like a trend change but not to a superior engine. In some cases there is a gradual or asymptotic change to a new level. This process is identifiable as dynamic change.

The intervention series should be identified by a maximum likelihood procedure which augments the list of input series. This procedure is not simply taking the model residuals and standardizing them to determine the outliers. A number of software developers report this as outlier detection but this approach requires the errors to be independent of each other.

The concept of data mining should be incorporated where the engine detects deviation. Detecting deviation, which is the exact opposite of database segmentation, identifies outlying points in a data set (records that do not belong to any other cluster) and determines whether they can be represented by statistically significant indicator variables. Deviation detection is often the source of true discovery since outliers express deviation from some known

expectation and norm. It is very important that intervention detection be done automatically in both a univariate (non-causal) model and a multivariate (causal) model.

### **Local Time Trends**

Q. Does the forecast engine identify the need for local time trends? How are changes in trend detected? Does the user have control over the minimum number of points to determine a new trend?

A. In many applications there exists local trends which sometimes change abruptly. ARIMA models are deficient in dealing with this phenomena as it uses level shifts or differencing factors to mimic the process. In some cases dummy variables using the counting numbers is a more appropriate and visually correct structure keeping in mind that the trend may not have started at the first data point. In general the trend may have a dead period in the beginning or at the end. A number of trends may be necessary in conjunction with pulses etc. and an error process involving some arbitrary ARIMA structure.

### **Seasonal Structure**

Q. How does the forecast engine develop seasonal model structure? Does it encompass both stochastic seasonality and deterministic seasonality? What tests can be performed to determine which one is best and can the system handle both types simultaneously?

A. Seasonal stochastic structure like adjusting for the sales say 52 periods ago is a form of an ARIMA model. If seasonal dummies were more correct then the statistical model might suggest seasonal differencing. Franses test can be used to test the hypothesis of unit roots leading to either the acceptance of seasonal differencing or its replacement with seasonal dummies. Occasionally there are cases where both seasonal difference and a seasonal pulse might be required to describe a process. The engine should comfortably handle both kinds. Additionally, a seasonal ARIMA structure might also be appropriate.

### **Sound Model Construction**

Q. How does the forecast engine allow the user permitted to specify holiday, price, weather, or other possible cause variables? How does the forecast engine allow the user to distribute sales across consecutive weeks that would reflect historical demand patterns caused by what day in the week the holiday falls? How does the forecast engine allow the user to measure the effect of the length of a promotion and particular effects like the last week of a promotion or the week after a promotion?

How does a given system allow the user to adjust for price changes that occur mid week? Is the user permitted to deal with and measure the effect of cross-cannibalization?

A. The engine should allow the user to specify a large (up to 30) number of series in a model. These input series may be either stochastic (probabilistic) or deterministic (fixed) series.

If sales depend on the day of the week that a holiday occurs then the user should use a (0.2,0.8) or a (0.4,0.6) etc indicator rather than an (0,1) to reflect the historical and anticipated sales breakout.

The engine should suggest building indicator variables for a promotion week and if large sales are expected the last week of a promotion then a 0/1 variable could be constructed for the last week. Similarly if reduced sales have historically occurred the week after a promotion a 0/1 variable should be constructed for that variable.

Mid-week price changes can be accommodated by simply computing the weighted sum and using it. Sometimes the weighted prices of alternative packages for the same brand might be useful. Alternatively, or perhaps additionally, the price of a specific alternative product might be used to account for cross-cannibalization.

## Omitted Variables

Q. How does the forecasting engine deal with omitted series and if so how? What surrogate model forms can be developed to account for these missing series?

A. If the omitted variable is stochastic and has no internal time dependency (white noise) then its effect is simply to increase the background variance resulting in a downward bias of the tests of necessity and sufficiency. If however the omitted series is stochastic and has some internal autocorrelation then this structure evidences itself in the error process and can be identified as a regular phenomenon and appears as ARIMA structure. For example, if degree is needed but omitted a seasonal ARIMA structure will be identified and becomes a surrogate for the omitted variable.

If the omitted variable is deterministic and without recurring pattern it may be identified via surrogate (intervention) series.

## Constancy of Variance

Q. How does the forecast engine verify that the variance of the model errors is constant? Are individual tests conducted for:

- Deterministic points of change in the variance?
- Level dependent changes in the variance?

A. If a series is more volatile at different levels or is more volatile at different points in time this may significantly effect model identification and estimation/forecasting. The engine should test for these issues and either suggest transformations of the form log, square root, etc or weighting transformations akin to weighted regression where values are *believed* inversely proportional to their degree of disbelief.

---

## Empower the User

Control of the forecast process by way of parameter enabling along with engine self diagnostics empowers a user with the ability to assess, compensate, and tailor the model simulation according to requirements which are unique for each application. Specific control over confidence level thresholds, model constraints, and starting model definition provide a user with the capability to direct the forecast process.

## Seasonal products

Q. How does the forecast engine deal with products that are only sold in specific weeks of the year?

A. By simply redefining the year to include actual sales the model can be constructed. If the product is only sold the first 10 weeks of the year simply don't use the 42 non-sales weeks.

## Intermittent Demand

Q. How does the forecast engine deal with long periods of no demand that follow no specific pattern?

A. If a series has time points where sales arise and long periods of time where no sales arise it is possible to convert sales to sales per period by dividing the observed sales by the number of periods of no sales thus obtaining a rate. It is then possible to identify a model between rate and the interval between sales culminating in a forecasted rate and a forecasted interval.

### **Pooled Cross-Sectional data**

Q. How does the forecast engine allow the user to test the hypothesis of a common model across different groups?

A. The engine should have an option that implements an F-test to test the hypothesis that one model/set of parameters is useful overall and a common set of coefficients (parameters) suffices.

### **Too Much Data**

Q. How does the forecast engine deal with detecting changes in a model or parameters over time? What steps does it take to ignore historical data that is no longer needed or should not be used to construct a model?

A. The engine should perform a test for constancy of parameters. This test can reveal the presence of too much data and suggest the optimal number of data points that should be used leading to a purging of older and non-essential data.

### **Too few Observations/New Product Forecasting**

Q. How does the forecast engine deal with new products? Can you utilize models of products with a significant history as tentative models for an emerging product?

A. One can use a model for a mature product and estimate parameters for an emerging product using a few readings. As new data becomes available, new product models can then be automatically identified.

---