Statistical Forecasting with Microsoft Excel Forecast Models

OBJECTIVE INSIGHTS

Business Analysis for the Health Care Industry

Abstract

Current forecasting software can be problematic in its level of integration with Microsoft Excel forecasting models. Standalone forecasting software operates in a batch mode with Excel and is impossible to run interactively without custom programming. Excel forecasting add-ins are designed to work within Excel and can be automated to use with forecasting models, but lack support for damped trends, out-of-sample testing, and cumulative prediction intervals.

Objective Insights has developed the Objective Insights Prediction Engine (OIPE), a general-purpose set of trending and statistics routines designed to be incorporated into Microsoft Excel forecast models. The OIPE features a comprehensive set of trending methods and in-sample trend accuracy statistics. In addition, the OIPE calculates point and cumulative prediction intervals for exponential smoothing models and uses an innovative approach for out-of-sample accuracy testing. Finally, the OIPE is designed to be incorporated into custom forecasting solutions in Microsoft Excel, improving forecast productivity and aiding insights into market conditions.

The OIPE supports several varieties of trending methods, including moving average; linear, exponential, and polynomial regression; exponential smoothing; and Box-Jenkins (ARIMA). Exponential smoothing methods include simple, linear, and multiplicative trends, seasonality, and damping. Box-Jenkins includes both seasonal and non-seasonal methods.

The OIPE provides a number of in-sample statistics for trends calculated in an Excel forecasting model; in-sample statistics are those calculated from the same dataset used to determine the fitted statistical model. The statistics included are measures of forecast accuracy or are used to compare the quality of fit among different statistical models.

Another valuable capability in forecasting software is out-of-sample (holdout) testing, particularly when comparing the results of different forecasting methods. The OIPE implements out-of-sample testing with the additional enhancement of using rolling windows for the holdout samples. Rolling windows are essentially a sequential series of holdout samples and better take into account how variations in the data favor or disfavor particular trending methods compared to a single holdout sample.

Prediction intervals are another useful feature of the OIPE and are calculated for the most commonly used exponential smoothing methods. Prediction intervals represent a range of possible outcomes expected to contain an unknown future value of product demand, for instance, with some specified probability (e.g., 95% or 80%). In addition to point prediction intervals, the OIPE also calculates cumulative (lead-time) prediction intervals. When forecasts for a number of time points are aggregated into a cumulative total (e.g., summing monthly or quarterly forecasts into a calendar year), one cannot simply use a weighted average of the individual prediction intervals, as there is covariance between those error terms. Failing to take the covariance into account will tend to underestimate the aggregate prediction interval.

Objective Insights compared the forecast results of the OIPE against leading forecasting software using real-world customer data. Tests across eight such sets of data found that the OIPE performed as well or better than a leading software package.

The OIPE is designed to be incorporated into Microsoft Excel forecasting models and was originally developed by Objective Insights to support our own forecasting models without having to resort to third-party software. The OIPE enables the forecasting model to easily support your forecasting process and

workflow, rather than having to work around the way in which third-party software imports data and exports results.

Current State of Trend-Based Forecasting in Microsoft Excel

Trend-based forecasting in Microsoft Excel currently ranges from simple (linear regression available in the Excel Analysis Toolpak) to sophisticated (Excel add-ins such as ForecastX and ezForecaster). In addition to interactive use, some of the add-ins can be programmed with Excel's built-in Visual Basic for Applications, making them suitable for automatic use with forecasting models (e.g., updating forecast trends at the click of a button).

Current add-ins suffer from several shortcomings, however. First, these add-ins appear to lack support for calculating damped trends when using exponential smoothing. This is important because it makes it more difficult for the forecaster to damp (roll off) trends with longer time horizons, where the expectation is that a trend is not likely to continue growing linearly over time but instead should start to moderate.

Second, the add-ins do not have a provision for calculating forecast accuracy using out-of-sample testing. Out-of-sample testing withholds a portion of the dataset used for forecasting and uses the remaining data to determine the trend, which is then compared to the held-out sample to measure how well the trend predicted what really happened. The accuracy statistics provided by these add-ins are instead calculated in-sample (from the difference between the forecast model fitted to the entire dataset and the data itself), which is a less robust means of determining which forecast method yields the most accurate results.

Current Excel add-ins also lack cumulative (lead-time) prediction intervals. Prediction intervals (level of certainty about the range into which a forecast may fall) are normally calculated at the level of the individual unit of time used in the forecast (i.e., a month for monthly forecasts). When aggregating monthly forecasts into a calendar year, however, one cannot simply use a weighted average of the monthly prediction intervals, as there is covariance between the monthly error terms. Failing to take the covariance into account will tend to underestimate the aggregate prediction interval.

Standalone Forecasting Software

Beyond Excel, there are several forecasting software packages that vary from specialized to general purpose. The Autobox software focuses on Box-Jenkins methods, but excludes other approaches to statistical forecasting. Forecast Pro is a sophisticated, full-featured forecasting package, but it shares the common flaw with all standalone forecasting software that it is difficult at best to integrate with Excel forecasting models in an automated fashion. For example, both Autobox and Forecast Pro have programming interfaces utilizing Windows dynamic link libraries (DLL's), yet in practice, these interfaces cannot be controlled from Excel using Excel's Visual Basic for Applications (VBA). Instead, customers wishing to integrate this software is an Excel model must write custom wrappers in computer languages like C++ in order to call the trending routines from Excel.

Other alternatives to standalone forecasting software includes statistical software such as SAS and R, as well as trending functionality built into enterprise resource planning (ERP) systems. Statistical software is very comprehensive and flexible, offering a huge range of trending methods and accuracy statistics and the ability to add custom trending routines if needed. However, this comes at the cost of a steep learning curve and there is still the inability to integrate well with an Excel forecast model.

ERP software has the advantage of integration with corporate financial and manufacturing planning systems, but can be quite limited in terms of forecast method selection, customization of business rules, and fine-grained control over how individual products are trended. For example, the ability to calculate the model "backwards," such as with Excel's Goal Seek function, is missing.

As with Excel add-ins, standalone trending software that we have reviewed lacks the capability to calculate cumulative prediction intervals.

Objective Insights Prediction Engine

The Objective Insights Prediction Engine (OIPE) is a general-purpose set of trending and statistics routines. The OIPE runs in Microsoft Excel 97 or later (including all Macintosh versions except for Excel 2008, which lacks Visual Basic) and is written in Visual Basic for Applications. As such, it needs no other supporting components or software, such as DLL's.

The OIPE was originally written in order to improve the forecast process with Excel models and circumvent other shortcomings of third-party software. By enabling closer integration of the trending software with the Excel forecast model, the OIPE facilitates a better forecasting workflow and faster turnaround of forecast results.

The OIPE features a comprehensive set of trending methods and in-sample trend accuracy statistics. In addition, the OIPE calculates point and cumulative prediction intervals for exponential smoothing models and uses an innovative approach for out-of-sample accuracy testing. Finally, the OIPE is designed to be incorporated into custom forecasting solutions in Microsoft Excel, improving forecast productivity and aiding insights into market conditions.

The OIPE is not a "black box," so Objective Insights can review the calculation flow with clients to show what is being calculated and how. The formulas and calculation sequences can be laid bare to help explain the trending methods and the results.

Trending Methods

Trending a dataset with the OIPE assumes that the dataset has already been adjusted for extraneous factors that may distort the trend. For example, monthly prescription data is typically adjusted to account for number of days in the month and holiday effects; as an aside, Objective Insights has developed an innovative approach for such adjustments. As such, the trending methods below are intended to be applied to a single dataset at a time (i.e., there is no multivariate analysis).

The OIPE supports the following well-accepted trending methods:

- Moving average
- Linear, exponential, and polynomial regression
- Exponential smoothing
- Box-Jenkins (ARIMA)

Moving Average

The OIPE calculates a simple trailing average of the last *n* data points; the projected trend is the value of the moving average at the last *n* data points in the historical data series. The period used for the moving average is easily changed in forecast models implementing this method.

Regression Types

The linear regression routine in the OIPE uses the COBYLA (Constrained Optimization BY Linear Approximation, Powell, 1994) optimization algorithm to minimize the sum of squared errors. COBYLA is also used in the other regression routines, as well as in exponential smoothing and ARIMA, where minimization of the objective function is needed.

Exponential regression applies a log transform to the data before using univariate linear regression, thereby allowing more accurate appraisal of data growing on a percentage basis.

Polynomial regression adds quadratic and square root terms to univariate linear regression; the form of the regression equation is thus $y = m_1 x + m_2 x^2 + m_3 x^{1/2}$. Determination of the regression coefficients is conducted by minimizing the sum of squared errors using COBYLA as described above.

The length of the historical data window used in these different types of regression analysis can be easily changed in the forecast model.

Exponential Smoothing

Exponential smoothing (ES) is a time series analysis method related to a moving average. Whereas all data points in a moving average are equally weighted, exponential smoothing applies a series of exponentially decreasing weights to the data points (hence the name). Exponential smoothing is a robust, relatively simple forecasting technique that is frequently one of the top-scoring methods in forecasting competitions, such as the M series of contests organized by Makridakis (1982, 1993, 1998).

The degree to which recent data are more weighted than older data is determined by the smoothing parameters. These factors can be calculated for the level, slope (or trend), and seasonality of the dataset and are determined through an optimization process where an ES model is fitted to the historical dataset. Parameter values are selected to minimize the step-ahead squared errors between the ES model and the historical data and vary between zero and one, with values closer to one more heavily weighting recent data points.

The selection of exponential smoothing method depends on the nature of the data. A time series that has little or no discernible trend is best addressed with simple (one-parameter) exponential smoothing, where only the smoothing factor for the level is calculated. Trended data are best addressed with two-parameter ES, where both the level and trend smoothing factors are calculated. This method is referred to as Holt exponential smoothing (after the originator of the technique) when the ES model uses a linear trend. Exponential (multiplicative) trends can also be used in ES models.
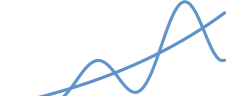
Another extension of exponential smoothing is where the time series is seasonal, such as with allergy products or influenza vaccines. This variation (known as Holt-Winters or simply Winters when the trend is linear) adds a third smoothing parameter for the cyclical component. Seasonality can be either additive (where the amplitude is constant over time) or multiplicative (where the amplitude varies with the level of the time series).

Trends may also incorporate a damping parameter, such that a trend's slope will approach zero at some point in the future. The more damping applied, the faster the trend will level out. Damping is useful where the forecaster does not expect a trend to linearly continue in the long term, such as during a product's initial growth phase. Damping can also be used to prevent the trend for a declining product from going below zero.

Another trending approach incorporated in the OIPE is the Theta method (Assimakopoulos & Nikolopoulos 2000), a forecasting model related to exponential smoothing. This method is useful for data where there is a slowly-evolving trend. The Theta method, also called simple exponential smoothing with drift, combines the level component found through simple exponential smoothing with a trend component computed as one-half the slope identified through linear regression on the entire time series.

The OIPE implements the state-space approach to exponential smoothing as described in Hyndman *et al* (2008). The results of the state-space approach match those of the classical methods of exponential smoothing (e.g., Gardner 1985), such as Holt and Holt-Winters; however, the state-space approach provides the underlying statistical models generally lacking in the classical methods. As such, prediction intervals may then be calculated for many of the exponential smoothing types

Exponential Smoothing Methods Implemented in OIPE

Trend Type	Seasonality		
	None	Additive	Multiplicative
None	Simple ES		
Additive	Holt ES		Holt-Winters ES
Multiplicative			

Box-Jenkins

The Box-Jenkins forecasting approach is a methodology combining autoregressive and moving average models with model identification and parameter estimation. Box-Jenkins models are also referred to as ARIMA (Auto Regressive Integrated Moving Average) models. The characteristics of an ARIMA model are embodied as ARIMA (p,d,q), where p is the order of the autoregressive component, d is the order of the integrated component, and q is the order of the moving average component. These components are integers equal to or greater than zero.

The autoregressive, or AR, part of the model is a straightforward multiple linear regression of a current value of the time series against one or more lagged (historical) values; this is rolled back through the entire time series to form the sets of dependent and independent variables in the regression. The number of past values considered is p, the order of the AR component.

The integrated component of the model is a differencing step intended to render the data series stationary in time, such that the mean, variance, and other statistical properties do not vary over the length of the series. For example, a linear trend is second order, as two differencing steps are necessary to remove the level and trend components to make such a time series stationary. Otherwise, with an increasing trend, one would expect the mean and variance to increase over time. In this case, the value of *d* is two.

The moving average, or MA, component is conceptually similar to the autoregressive component, but the current value of the time series is instead regressed against one or more lagged errors of the fitted model vs. historical values. In addition, the regression is done as an iterative process, since the errors are derived from the fitted model. The number of past error terms included is *q*, the order of the MA component.

Seasonality may also be taken into account in an ARIMA model in a similar fashion as described above, with seasonality segmented into autoregressive, differenced, and moving average components. A fully seasonal ARIMA model is denoted as ARIMA $(p,d,q) \times (P,D,Q)$, where the capitalized letters represent the order of the seasonal components.

One disadvantage of Box-Jenkins models is that they can require a fairly long time series (on the order of four or more years of monthly data) in order to accurately fit an appropriate model. Another disadvantage is that finding the appropriate Box-Jenkins model is computationally intensive, so this step can take several times longer than optimizing exponential smoothing models.

Objective Insights has licensed ARIMA routines from the [Numerical Algorithms Group](#) (NAG) to carry out Box-Jenkins forecasting in the OIPE. NAG is a software development firm well-known for its robust numerical methods routines.

Trend Statistics

The OIPE provides a number of in-sample statistics for trends calculated in an Excel forecasting model. In-sample statistics are those calculated from the same dataset used to determine the fitted statistical model.

The statistics in the following table are measures of forecast accuracy or are used to compare the quality of fit among different statistical models. Scaled accuracy measures (MAPE, sMAPE, MASE) can be compared across different datasets, while non-scaled measures (MAE, MSE) may only be used to compare model results within the same dataset. Fit measures such as R^2, AIC, and BIC inform the forecaster as to how well a model represents the data while adjusting for the complexity of the model.

OIPE Trend Statistics	
Fitted Model Error Variance	Variance of the one-period-ahead errors in the fitted model compared to the historical data
Mean Absolute Error (MAE)	Mean of the absolute value of the errors of the fitted trending model. Can only be used to compare forecasts computed from the same data series due to scale issues.
Mean Absolute Percentage Error (MAPE)	Mean of the absolute value of the errors of the fitted trending model divided by observed data. Can be used to compare forecasts computed from different data series.

OIPE Trend Statistics	

Symmetric Mean Absolute Percentage Error (sMAPE)	• More robust version of MAPE, especially when observed values are close to zero. Can be used to compare forecasts computed from different data series.
	    • sMAPE = mean(200*(Observed - Fitted)/(Observed + Fitted))
	    • Source: Makridakis (1993)

Mean Absolute Scaled Error (MASE)	• Mean of the errors of the fitted model scaled by the Mean Absolute Error of the in-sample flat-line forecast. MASE is less than one if the selected forecast method is better than the flat-line forecast and greater than one if worse than the flat-line forecast.
	    • Can be used to compare forecasts computed from different data series
	    • Source: Hyndman & Koehler (2006)

Mean Squared Error (MSE)	Mean of the squared errors of the fitted trending model. Can only be used to compare forecasts computed from the same data series due to scale issues.
Thiel's U	• Square root of the following: sum of the one-period-ahead forecast errors for the selected method divided by the sum of the one-period-ahead forecast errors for the flat-line method. • Somewhat similar to MASE, in that a Thiel's U value less than one indicates a better forecast than flat-line, while a value greater than one signifies a worse forecast than flat-line.

R^2 & Adjusted R^2	• R^2 is the coefficient of determination, the proportion of variability in a data set that is accounted for by a statistical model. The closer R^2 is to 1, the better the fit of the model to the observed data.
	    • R^2 = 1 - (sum of squared residuals)/(total sum of squares)
	    • Adjusted R^2 penalizes R^2 for the number of variables in the model. This statistic therefore rewards the model that best explains the data with the fewest variables

AIC (Akaike's Information Criterion)	• AIC = n*ln(RSS/n) + 2*k, where:
	    • n = number of periods in the historical data set
	    • k = number of parameters in the forecast model (e.g., Holt exponential smoothing has two parameters, level and trend)
	    • RSS = sum of the squared residuals

BIC (Schwarz/Bayesian Information Criterion)	• Similar to AIC, but the parameter term now becomes dependent on the log of the number of periods in the historical data set
	• BIC = n*ln(RSS/n) + k*ln(n), where:
	    • n = number of periods in the historical data set
	    • k = number of parameters in the forecast model (e.g., Holt exponential smoothing has two parameters, level and trend).
	    • RSS = sum of the squared residuals

Out-of-Sample Testing

Another valuable capability in forecasting software is out-of-sample (holdout) testing, particularly when comparing the results of different forecasting methods. An out-of-sample test is conducted by reserving the latter portion of a historical dataset. A trend is then calculated using only the remaining portion of the historical data and the trend for the equivalent part of the period is compared to the held-out sample for accuracy. For example, if you have a monthly ten year dataset, you might reserve two years of data for out-of-sample testing. You would then trend the first eight years of data and compare the two year trend to the equivalent two years of held-out data to evaluate how well the forecast matched what actually happened.

In-sample testing, in which the model is fitted and compared to the entire dataset, can be less accurate than holdout testing, partly because it is always possible to build a complicated model which exactly matches the dataset but might have poor predictive value.. In essence, in-sample testing allows the model to "cheat" by looking at the data it will be compared to and taking variations in that data into account. There is no such opportunity with out-of-sample testing, as the trend model is blind to the held-out data. Out-of-sample testing thus provides the truest test of the most accurate trending method when there are unexpected deviations in the data.

The OIPE implements out-of-sample testing with the additional enhancement of using rolling windows for the holdout samples. Rolling windows are essentially a sequential series of holdout samples and better take into account how variations in the data favor or disfavor particular trending methods compared to a single holdout sample. For instance, six rolling windows might be tested, each with a six-month holdout period. The first window would trend points 1 through *n*-6 and use as the holdout period points *n*-5 through *n*. The second window would trend points 1 through *n*-7 with the holdout of *n*-6 through *n*-1. The third window would trend points 1 through *n*-8 with the holdout of *n*-7 through *n*-2. This process continues until the sixth window has been tested.

Prediction Intervals

Prediction intervals are another useful feature of the OIPE. Prediction intervals represent a range of possible outcomes expected to contain an unknown future value of product demand, for instance, with some specified probability (e.g., 95% or 80%).

Prediction intervals should not be confused with confidence intervals. Confidence intervals reflect uncertainty due to measurement errors (i.e., "noise"), in that all sales demand data are samples used to estimate true product demand. Confidence intervals also only apply inside the historical data range, as opposed to prediction intervals, which apply outside the historical data range.

Prediction intervals are only calculated for the simple and linear exponential smoothing methods in the OIPE. Prediction intervals cannot be calculated for the other trending methods either because there is no underlying statistical process "producing" the data (e.g., moving average); the statistical model used to derive prediction intervals is too complicated, as is the case with most Box-Jenkins models and non-linear exponential smoothing; or the assumptions underlying the statistical model are often violated by real-world data. For example, linear regression assumes that forecast errors are independent of each other, yet such errors in time series data are frequently subject to autocorrelation as well as non-constant variance (heteroscedasticity).

Prediction intervals will tend to grow over time. This makes sense, as one would expect that the farther out in time, the greater the uncertainty. To illustrate with Holt exponential smoothing, the prediction interval half-width is governed by the following equation:

$$Z\frac{\alpha}{2}\sqrt{\left[1+(k\text{-}1)\alpha^2\left(1+k\gamma+\frac{k(2k\text{-}1)\gamma^2}{6}\right)\right]\sigma^2}$$

where:

- $Z_{\alpha/2}$ = certainty factor (e.g., 1.96 for 95% certainty)
- k = time index of forecast (e.g., for the second month of a forecast, k=2)
- α = optimized exponential smoothing level parameter
- γ = optimized exponential smoothing trend parameter

- σ^2 = variance of the one-step-ahead errors in the fitted model vs. observed data

If we expand and deconstruct the half-width equation, with

- $L_1 = k - 1$
- $L_2 = k(k - 1)$
- $L_3 = k(2k - 1)(k - 1)$

then the Holt exponential smoothing half-width equation becomes:

$$Z\frac{\alpha}{2}\sqrt{\left[1 + L_1\alpha^2 + L_2\alpha^2\gamma + \left(\frac{L_3\alpha^2\gamma^2}{6}\right)\right]\sigma^2}$$

The half-width of the prediction interval is therefore seen to vary with:

- $Z_{\alpha/2}$: The higher the certainty level (i.e., the more certain that the forecast will lie within the prediction interval), the wider the prediction interval
- k: The farther out in time that the forecast is, the wider the prediction interval
- α & γ: The higher the values for the smoothing parameters (i.e., the more the trend is weighted toward recent history), the wider the prediction interval
- σ^2: The noisier the observed data, the wider the prediction interval

In addition to computing prediction intervals for each forecasted time period, the OIPE also calculates cumulative (lead-time) prediction intervals. When forecasts for a number of time points are aggregated into a cumulative total (e.g., summing monthly or quarterly forecasts into a calendar year), one cannot simply use a weighted average of the individual prediction interval percentage half-widths, as there is covariance between those error terms. Failing to take the covariance into account will tend to underestimate the aggregate prediction interval. As with the point prediction intervals, cumulative prediction intervals are available only for simple and linear exponential smoothing models, as there are no exact analytical solutions for non-linear exponential smoothing.

Prediction intervals should not be interpreted as insurance that all future values will be encompassed by the forecast. Some actual values will still lie outside of the prediction interval, since the intervals are governed by the certainty level (e.g., a 95% certainty level implies that one in twenty cases will fall outside the prediction interval on average). In addition, any unadjusted trend-based forecast makes the assumption that the future will be like the past. If there is an unexpected upside or downside event, the prediction intervals will no longer be valid until that event has been incorporated into the trend.

To test the validity of prediction intervals produced by the OIPE, Objective Insights analyzed 334 series of "Monthly/Industry" data from the M3 forecasting competition (Makridakis & Hibon, 2000) using Holt exponential smoothing. The prediction intervals from the resulting trends were found to be somewhat too narrow, as is often observed in practice (Chatfield 1993). However, alternative methods of computing prediction intervals can yield ranges that are excessively wide (Yar & Chatfield 1990). Consequently, Objective Insights thinks that the prediction intervals produced by the OIPE are quite reasonable in light of test results and the weight of opinion expressed in the academic literature.

Forecast Accuracy

The true test of forecasting software is the accuracy with which it performs. Objective Insights recently tested the OIPE with eight sets of real-world customer data against Forecast Pro XE (version 5) using Holt

and Holt-Winters exponential smoothing. The results are summarized in the following table; bold type indicates where the value is both better (lower) and statistically significantly different (at the 95% confidence level).

MAPE Dataset	Holt		Winters (Additive)		Winters (Multiplicative)	
	FP	OIPE	FP	OIPE	FP	OIPE
1	4.4%	4.3%	4.7%	4.7%	4.7%	4.9%
2	3.2%	2.7%	3.7%	**3.5%**	4.1%	**3.8%**
3	3.1%	2.9%	**3.6%**	3.8%	3.9%	4.0%
4	3.1%	3.0%	3.6%	**2.8%**	3.9%	**3.0%**
5	3.1%	3.3%	3.8%	**3.1%**	4.1%	**3.2%**
6	3.0%	2.8%	3.1%	3.4%	3.4%	3.7%
7	3.3%	3.1%	3.2%	3.4%	3.5%	3.2%
8	2.8%	3.2%	3.3%	3.2%	3.7%	3.3%

FP = Forecast Pro XE (version 5)
MAPE = Mean Absolute Percentage Error (lower is better)

As shown above, the OIPE compares favorably to Forecast Pro. In general, the exponential smoothing results are statistically equivalent to those produced by Forecast Pro, though in several cases with the Holt-Winters methods, the OIPE actually outperformed Forecast Pro.

Integration with Custom Solutions

The OIPE is designed to be incorporated into Microsoft Excel forecasting models and was originally developed by Objective Insights to support our own forecasting models without having to resort to third-party software. The OIPE enables the forecasting model to easily support your forecasting process and workflow, rather than having to work around the way in which third-party software imports data and exports results.

The OIPE also allows fine-grained control of trending within the forecast model. The forecaster can have the option of running trends automatically or being able to pre-set methods and parameters down to the individual trend level.

Where complete transparency is needed, Objective Insights can open the OIPE to show what is being calculated and how. The formulas and calculation sequences can be laid bare to help explain the trending methods and the results.

Finally, the forecast model can be set up to structure trend results from the OIPE into tables and charts exactly the way the forecaster wants them. Conclusions can then be drawn and results more quickly conveyed to company decision makers, as opposed to separately constructing those tables and charts through intermediate steps.

Objective Insights has built the OIPE into three separate forecasting solutions to date:

- Short-Term Forecaster® (STF)
- Trend Explorer
- Trend Explorer Plus

Short-Term Forecaster™ (STF)

The STF™ is a tool designed to forecast near-term revenues based on historical market data. The model is easy-to-use, offering one-click trend generation and simple importation of product and market data.

The STF™ starts with customer demand, which is the cleanest, truest measure of demand, and builds an ex-factory revenue forecast. The STF™ projects customer demand trends using the OIPE. The model automatically creates as many trends as are needed for every key product in the market, including those from your company and the competitors. The STF™ easily handles products that have many subproducts (strengths, vial sizes, formulations, etc.) and indications.

When calculating trends, the STF™ modifies the raw demand data to take into account factors such as days per month and holiday effects. In addition, on the revenue side, the STF™ adjusts for factors such as price changes, discounts, product amount per prescription, and gross to net revenue conversion.

The STF™ allows entry of future events (assumptions) that are not reflected in the historical data. Such events encompass competitive product launches (including cannibalization by a company's successor products) and events affecting either specific products or the entire market. All assumptions can be clearly documented in the model through note fields. If there are no assumption changes from month to month, an updated forecast can be produced in minutes. Version control is handled through time and name stamps, and since it is an Excel model, a copy can be archived for future reference.

Another feature of the STF™ is the use of Monte Carlo simulation as an additional risk analysis measure. The STF™ calculates prediction intervals for the raw demand trends, but Monte Carlo simulation is used to gauge the effect of uncertainty in the assumptions for future (non-trended) events. The STF™ currently produces uncertainty deciles for net revenues, units, and prescriptions, though other measures could also be easily included.

Trend Explorer™

The Trend Explorer™ is a tool intended for trend analysis and comparison with a single dataset and uses the OIPE for all trend calculations. Trend Explorer™ has two primary functions: *ad hoc* analysis of a single dataset and comparing different trending methods with that dataset.

Ad hoc trend analysis is similar to the function found in other forecasting software. The trending method is user-selectable and parameters associated with the selected method may be adjusted prior to running the trend. For example, the number of data points used for a moving average or regression may be specified. Exponential smoothing and Box-Jenkins can either be run automatically (i.e., optimized for best results) or individual parameters may be specified to examine how they affect the trend. The accuracy of the trend may be evaluated using the trend statistics described in a previous section.

Trend Explorer™ also offers a trend comparison option where the results of all specified trending methods are automatically tested against a series of held-out samples from the end of the input dataset. Calculated trends are compared against rolling windows over the length of the held-out sample in order to minimize the effects from situations where different segments of data may favor one method over another. Both the number and length of the rolling windows are user-specifiable (subject to the length of the overall dataset). Results of the trend comparison may be ranked against any of the statistics described in the previous section.

Trend Explorer Plus™

The Trend Explorer Plus™ is a combination of the STF™ and Trend Explorer™. As with the Trend Explorer™, the Trend Explorer Plus™ evaluates which trending method works best through out-of-sample testing using rolling holdout windows. However, Trend Explorer Plus™ is designed to produce a complete product forecast (as opposed to the Trend Explorer's focus on trend analysis).

The Trend Explorer Plus™ easily imports product demand data, runs and tests demand trends, and then shows the results of the best trend and selected comparator trends in terms of units or prescriptions and gross and net revenues. As with the STF™, the Trend Explorer Plus™ can take into account factors such as days per month, holiday effects, price changes, and gross to net revenue conversion.

Comments & Questions

We welcome your comments or questions on this white paper. Please e-mail your comments or questions to info@objectiveinsights.com and we will happily respond.

Bibliography

Armstrong JS (ed.) (2001) Principles of Forecasting: A Handbook for Researchers and Practitioners. Kluwer Academic 2001. ISBN 0-7923-7930-6.

Assimakopoulos V, Nikolopoulos K (2000) The theta model: a decomposition approach to forecasting. *International Journal of Forecasting*; 16, 521-530.

Chatfield C (1993) Calculating Interval Forecasts. *Journal of Business & Economic Statistics*; 11(2), 121-135.

Chatfield C, Yar M (1991) Prediction intervals for multiplicative Holt-Winters. *International Journal of Forecasting*; 7, 31-37.

Gardner ES (1985) Exponential smoothing: The state of the art. *Journal of Forecasting*; 4 , 1–28.

Gardner ES (2006) Exponential smoothing: The state of the art—Part II. *International Journal of Forecasting*; 22(4), 637-666.

Gardner ES (1988) A Simple Method of Computing Prediction Intervals for the Time Series Forecasts. *Management Science*; 34(4), 541-546.

Hyndman RJ, Koehler AB (2006) Another look at measures of forecast accuracy. *International Journal of Forecasting*; 22, 679-688.

Hyndman RJ, Koehler AB, Ord JK, Snyder RD (2008) Forecasting with Exponential Smoothing: The State Space Approach. Springer-Verlag Berlin ISBN 978-3-540-71916-8.

Hyndman RJ, Koehler AB, Ord JK, Snyder RD (2005) Prediction Intervals for Exponential Smoothing Using Two New Classes of State Space Models. Journal of Forecasting; 24, 17-37.

Hyndman RJ, Koehler AB, Snyder RD. A State Space Framework For Automatic Forecasting Using Exponential Smoothing Methods. Working Paper 9/2000. Department of Econometrics and Business Statistics, Monash University. ISBN 0-7326-1078-8.

Hyndman RJ, Billah B. Unmasking the Theta Method. Working Paper 5/2001. Department of Econometrics and Business Statistics, Monash University. ISBN 0-7326-1085-0.

Makridakis S (1993) Accuracy measures: theoretical and practical concerns. *International Journal of Forecasting*; 9, 527-529.

Makridakis S, Hibon M (2000) The M3-competition: results, conclusions and implications. *International Journal of Forecasting*; 16, 451-476.

NAG (Numerical Algorithms Group) Fortran Library, http://www.nag.co.uk/numeric/FL/FLdescription.asp

Ord JK, Koehler AB, Snyder RD (1997) Estimation and Prediction for a Class of Dynamic Nonlinear Statistical Models. *Journal of the American Statistical Association*; 92(440), 1621-1629.

Powell MJD. "A direct search optimization method that models the objective and constraint functions by linear interpolation," in *Advances in Optimization and Numerical Analysis*, eds. Gomez S and Hennart JP (Kluwer Academic: Dordrecht, 1994), p. 51-67.

Powell MJD (1998) "Direct search algorithms for optimization calculations," *Acta Numerica* 7, 287-336 .

Yar M, Chatfield C (1990) Prediction intervals for the Holt-Winters forecasting procedure. *International Journal of Forecasting*; 6, 127-137.