

Retek[®] Demand Forecasting[™] 10.0

User Guide



The software described in this documentation is furnished under a license agreement, is the confidential information of Retek Inc., and may be used only in accordance with the terms of the agreement.

No part of this documentation may be reproduced or transmitted in any form or by any means without the express written permission of Retek Inc., Retek on the Mall, 950 Nicollet Mall, Minneapolis, MN 55403, and the copyright notice may not be removed without the consent of Retek Inc.

Information in this documentation is subject to change without notice.

Retek provides product documentation in a read-only-format to ensure content integrity. Retek Customer Support cannot support documentation that has been changed without Retek authorization.

Corporate Headquarters:

Retek Inc.
Retek on the Mall
950 Nicollet Mall
Minneapolis, MN 55403

888.61.RETEK (toll free US)
+1 612 587 5000

Retek[®] Demand Forecasting[™] is a trademark of Retek Inc.

Retek and the Retek logo are registered trademarks of Retek Inc.

European Headquarters:

Retek
110 Wigmore Street
London
W1U 3RW
United Kingdom

Switchboard:
+44 (0)20 7563 4600

Sales Enquiries:
+44 (0)20 7563 46 46
Fax: +44 (0)20 7563 46 10

This unpublished work is protected by confidentiality agreement, and by trade secret, copyright, and other laws. In the event of publication, the following notice shall apply:

©2004 Retek Inc. All rights reserved.

All other product names mentioned are trademarks or registered trademarks of their respective owners and should be treated as such.

Printed in the United States of America.



Customer Support

Customer Support hours

Customer Support is available 7x24x365 via e-mail, phone, and Web access.

Depending on the Support option chosen by a particular client (Standard, Plus, or Premium), the times that certain services are delivered may be restricted. Severity 1 (Critical) issues are addressed on a 7x24 basis and receive continuous attention until resolved, for all clients on active maintenance. Retek customers on active maintenance agreements may contact a global Customer Support representative in accordance with contract terms in one of the following ways.

Contact Method	Contact Information
E-mail	support@retек.com
Internet (ROCS)	rocs.retek.com Retek's secure client Web site to update and view issues
Phone	1 612 587 5800

Toll free alternatives are also available in various regions of the world:

Australia	1 800 555 923 (AU-Telstra) or 1 800 000 562 (AU-Optus)
France	0800 90 91 66
United Kingdom	0800 917 2863
United States	1 800 61 RETEK or 800 617 3835

Mail	Retek Customer Support Retek on the Mall 950 Nicollet Mall Minneapolis, MN 55403
-------------	---

When contacting Customer Support, please provide:

- Product version and program/module name.
- Functional and technical description of the problem (include business impact).
- Detailed step by step instructions to recreate.
- Exact error message received.
- Screen shots of each step you take.

Contents

Chapter 1 – Retek Demand Forecasting Overview 1

What is Retek Demand Forecasting?	1
Forecasting challenges and solutions	2
Selecting the best forecasting method	2
Handling items with limited demand histories	2
Handling lost sales and unusually high demand	3
Forecasting demand for new products and locations	3
Incorporating the effects of promotions and other event-based challenges on demand	4
Providing detailed sales predictions based on an assortment plan	4
Retek Demand Forecasting Features	5
Automated alert system for user-maintained exception reporting	5
Generation of both final level and source level forecasts	6
Retek Demand Forecasting Modules	7
Predict	7
Promote	7
Curve	7
Preprocessing	8

Chapter 2 – Predict (Statistical Forecasting)..... 9

Overview	9
Workbooks and Wizards	9
Forecasting Methods	9
Forecast Administration Workbook	12
Forecast Administration Wizard	13
Basic Settings workflow tab	13
Advanced Settings workflow tab	20
Forecast Maintenance Workbook	26
Forecast Maintenance Wizard	26
Basic Settings workflow tab	27
Advanced Settings workflow tab	29
Run Batch Forecast	37
Batch Forecast Wizard	37

Delete Forecasts	38
Delete Forecasts Wizard	38
Forecast Approval Workbook	39
Forecast Approval Wizard	40
Final Forecast Worksheet	41
Source Level Worksheet	43
Approval Worksheet	44
Final or Source Parameters Worksheet	45
Confidence Worksheet	47
Interactive Forecasting Workbook	47
Forecasting Parameter Worksheet	49
Interactive Forecasting Worksheet	50
Forecast Scorecard Workbook	51
Forecast Scorecard wizard	52
Final Error Measure Worksheet	53
Final Actuals vs. Forecasts Worksheet	55
Export Forecasts	56
Export Forecasts Wizard	57
Forecast Export Administration Workbook	58
Access the Forecast Export Administration Workbook	58
Export Parameters Worksheet	59
Procedures in Statistical Forecasting	63
Set or change default parameters for a forecast	63
Adjust forecast parameters for specific products or locations	64
Manually request an unscheduled forecast	65
Review, modify, and approve forecast values	66
Manually export a forecast to another program	68
Set or change forecast export parameters	69
Monitor the accuracy of Retek Demand Forecasting forecasts	69
Delete an unnecessary forecast	70

Chapter 3 – Promote (Promotional Forecasting) 71

Overview	71
What is Promotional Forecasting?	71
Comparison between promotional and statistical forecasting	72
Developing promotional forecast methods	72
Retek’s approach to promotional forecasting	74
Promotional forecasting terminology and workflow	74
Promote Workbooks and wizards	75

Promotion Planner Workbook Template.....	76
Promotion Planner wizard	77
Promotion Planner Workbook and Worksheet	77
Promotion Selector Template.....	78
Promotion Selector Wizard	78
Causal Maintenance Workbook Template	79
Causal Maintenance Wizard.....	79
Causal Maintenance Workbook and Worksheets	80
Promotion Analysis and Planning Workbook Template.....	82
Promotion Analysis and Planning Wizard.....	82
Promotion Analysis and Planning Workbook	83
Procedures in Promotional Forecasting.....	84
Set up the system to run a promotional forecast.....	84
Chapter 4 – Curve (Profile-based Forecasting)	85
Overview	85
Profile Management Terminology and Workflow	86
Profiles (spreading ratios).....	86
Profile generation techniques	87
Building system profiles based on the historic level	88
Overriding profiles by adjusting the source measure	88
Defining the training window for historical data.....	89
Aggregating historical data.....	89
Unreliable data, thresholds, and threshold exceptions.....	91
Loading default ratios	92
Approving the generated profiles	92
Profile Administration Workbook.....	95
Profile Administration Wizard	95
Core Profile Worksheet	96
Creating new profile templates.....	98
Profile Maintenance Workbook	98
Profile Maintenance Wizard.....	98
Profile Maintenance Workbook.....	99

Run Batch Profile	100
Profile Approval Workbook.....	101
Profile Approval Wizard	101
Profile Approval Worksheet.....	102
Source Level Worksheet.....	103
Final Profile Worksheet.....	105
Procedures in Profile-based Forecasting	106
Set default parameters for a profile	106
Modify the training window for a subset of items in a profile	107
Manually request a generated profile	107
Approve a profile.....	108
Set Retek Demand Forecasting to perform plan explosion	109
Create a new profile template.....	110
Chapter 5 – Preprocessing.....	113
Overview	113
Preprocessing Workbook templates and wizards.....	113
Administration Workbook.....	114
Administration Wizard	114
Administration Worksheet.....	114
Run Batch Workbook.....	116
Run Selection Update.....	116
Preprocessing methods	118
Available preprocessing methods.....	118
Mathematical formulation of preprocessing methods	122
Chart-form examples of preprocessing effects.....	125
Chapter 6 – Retek Demand Forecasting methods	129
Forecasting techniques used in RDF	129
Exponential smoothing.....	129
Regression analysis	130
Bayesian analysis.....	130
Prediction intervals.....	130
Automatic method selection	130
Source level forecasting.....	131
Promotional forecasting.....	131

Time series (statistical) forecasting methods	131
Why use statistical forecasting?	132
Confidence limits.....	133
Auto Exponential Smoothing (AutoES) forecasting	133
Definitions of equation notation used in this section.....	134
Components of exponential smoothing	135
Simple moving average	135
Simple exponential smoothing	136
Croston's method.....	137
Simple/Intermittent Exponential Smoothing	138
Holt exponential smoothing.....	138
Multiplicative Winters exponential smoothing	139
Additive Winters exponential smoothing	140
Seasonal Exponential Smoothing (SeasonalES).....	141
Seasonal Regression	142
Bayesian Information Criterion (BIC).....	144
Automatic forecast level selection (Auto-Select)	147
Profile-based forecasting.....	150
Forecast method.....	150
Profile based method and new SKUs	151
Example.....	151
Bayesian forecasting	152
Sales plans vs. historic data	153
Forecasting algorithm.....	153
Guidelines.....	155
Causal (promotional) forecasting methods	155
The causal forecasting algorithm.....	157
Causal forecasting algorithm process	159
Causal forecasting array interface description.....	161
Causal forecasting at the daily level	161
Final considerations about causal forecasting	163
Glossary	165
Index.....	173

Chapter 1 – Retek Demand Forecasting Overview

What is Retek Demand Forecasting?

Retek® Demand Forecasting™ is a Windows-based statistical and causal forecasting solution. It uses state-of-the-art modeling techniques to produce high quality forecasts – with minimal human intervention. Forecasts produced by the Demand Forecasting system enhance the retailer's supply-chain planning, allocation, and replenishment processes, enabling a profitable and customer-oriented approach to predicting and meeting product demand.

Today's progressive retail organizations know that store-level demand drives the supply chain. The ability to forecast consumer demand productively and accurately is vital to a retailer's success. The business requirements for consumer responsiveness mandate a forecasting system that more accurately forecasts at the point of sale, handles difficult demand patterns, forecasts promotions and other causal events, processes large numbers of forecasts, and minimizes the cost of human and computer resources.

Forecasting drives the business tasks of planning, replenishment, purchasing, and allocation. As forecasts become more accurate, businesses run more efficiently by buying the right inventory at the right time. This ultimately lowers inventory levels, improves safety stock requirements, improves customer service, and increases the company's profitability.

The competitive nature of business requires that retailers find ways to cut costs and improve profit margins. The accurate forecasting methodologies provided with Retek Demand Forecasting can provide tremendous benefits to businesses.

A connection from Retek Demand Forecasting to Retek Predictive Planning products is built directly into the business process by way of the automatic approvals of forecasts, which are then fed directly to Retek Predictive Planning. This process allows you to accept all or part of a generated sales forecast. Once that decision is made, the remaining business measures are planned within Retek TopPlan.

Forecasting challenges and solutions

A number of challenges affect the ability of organizations to forecast product demand accurately. These challenges include selecting the best forecasting method to account for level, trending, seasonal, and spiky demand; generating forecasts for items with limited demand histories; forecasting demand for new products and locations; incorporating the effects of promotions and other event-based challenges on demand; and accommodating the need of operational systems to have sales predictions at more detailed levels than planning programs provide.

Selecting the best forecasting method

One challenge to accurate forecasting is the selection of the best model to account for level, trending, seasonal, and spiky demand. Retek's AutoES (Automatic Exponential Smoothing) forecasting method eliminates this complexity.

The AutoES method evaluates multiple forecast models, such as Simple Exponential Smoothing, Holt Exponential Smoothing, Additive and Multiplicative Winters Exponential Smoothing, Crostons Intermittent Demand Model, and Seasonal Regression forecasting to determine the optimal forecast method to use for a given set of data. The accuracy of each forecast is evaluated, as is the complexity of the forecast model, in order to determine the most predictive method. You simply select the AutoES forecast generation method, and the system finds the best model.

Handling items with limited demand histories

Another challenge to creating accurate forecasts is predicting demand for items with a limited sales history. The accuracy of time series forecasts is diminished when short or sparse historical data is fed into the forecasting model. The Demand Forecasting approach to this challenge is to generate forecasts at multiple levels of data aggregation.

Forecasts are typically difficult to generate directly at the SKU/Store/Week level due to the low sales volumes of the items in question. Although forecast information is often required at a low level, the final forecast level, data is often too sparse and noisy to identify clear demand patterns. For this reason, it is sometimes necessary to aggregate sales data from this base level to a higher forecasting level in order to generate a reasonable forecast. At the higher level, trends and seasonality in the data can be more readily identified. Once the forecast is created at the higher level, the source level, the results can be spread back down to the final level based on that lower level's relationship to the total. Demand Forecasting generates forecasts directly at the lower level to determine demand proportions across SKUs.

Handling lost sales and unusually high demand

Future forecasts should be based on past demand. Unfortunately, retailers cannot record demand, but instead record sales. These two figures can differ when inventory drops to zero and there is demand, but no sales. To deal with this situation, RDF has a module called Preprocessing that recognizes when sales might be lower than actual demand, and adjusts sales values up to a level to a predicted level of where demand might have been.

Another situation when a retailer might want to adjust sales before using them as proxies for demand in forecasting is when past sales were unusually high due to an external event that is not expected to repeat. For example, if bad weather causes a power outage, then sales of batteries and flashlights for the week are likely to soar. These high sales values, however, are not good predictors of expected sales for the *next* week, when a power outage is less likely to occur. In this case, one would want to adjust sales downward, not so much to reflect true demand as to reflect “normal” demand. The same interpolation algorithms used to adjust sales up to correct for lost sales can be used to adjust sales down to correct for unusually high demand.

For more details on the Preprocessing module, see Chapter 5.

Forecasting demand for new products and locations

Retek Demand Forecasting can also forecast demand for new products and locations for which no sales history exists. You can model a new product’s demand behavior based on that of an existing similar product for which you do have a history. Forecasts can thus be generated for the new product based on the history and demand behavior of the existing one. Likewise, the sales histories of existing store locations can be used as the forecast foundation for new locations in the chain. For more details, see the Forecast Maintenance – Advanced Tab section in Chapter 2.

Incorporating the effects of promotions and other event-based challenges on demand

Promotions, non-regular holidays, and other causal events create another significant challenge to accurate forecasting. Promotions such as advertised sales and free gifts with purchase can have a significant impact on a product's sales history, as can irregularly occurring holidays such as Easter.

Using Promote (an optional add-on to Retek Demand Forecasting) promotional models of forecasting can be developed to take these and other factors into account when forecasts are generated. Promote attempts to identify the causes of deviations from the established seasonal profile, quantify these effects, and use the results to predict future sales when conditions in the selling environment will be similar. This type of advanced forecasting identifies the behavioral relationship of the variable you want to forecast (sales) not only to its own past, but also to explanatory variables such as promotion and advertising.

Suppose that your company has a large promotional event during the Easter season each year. The exact date of the Easter holiday varies from year to year; as a result, the standard time series forecasting model often has difficulty representing this effect in the seasonal profile. Promote tool allows you to identify the Easter season in all years of your sales history, and then define the upcoming Easter date. By doing so, you can causally forecast the Easter-related demand pattern shift.

Providing detailed sales predictions based on an assortment plan

The planning process attempts to establish the correct balance between different products in order to maximize sales opportunities within the available selling space. To facilitate this process, an assortment plan is often created. The assortment plan provides details of anticipated sales volumes and stock needs at aggregated levels. However, many operational systems require base data at much lower hierarchical levels, because these systems are responsible for ensuring that proper quantities of individual products are present in the right stores at the right time.

To address this need, Demand Forecasting contains an optional predictive solution (Curve) that transforms organization-level assortment plans into base-level weekly sales forecasts. Curve generates these lower level sales predictions by applying sets of profiles, or spreading ratios, to the assortment plan. The plan is thus exploded, or allocated, across the product, location, and time hierarchies.

Retek Demand Forecasting Features

In the forecast generation process, the Demand Forecasting application is supported and characterized by features such as:

- Multidimensional databases and database components (dimensions, positions, hierarchies)
- Product, location, and calendar hierarchies
- Aggregation and spreading of sales and forecast data
- Client-server architecture and master database
- Workbooks and worksheets for displaying and manipulating forecast data
- Wizards for creating and formatting workbooks and worksheets
- Menus, quick menus, and toolbars for working with sales and forecast data
- An automated alert system that provides user-defined and user-maintained exception reporting
- Generation of both final level and source level forecasts

More details about most of these features are in the Retek Predictive Application Server User Guide and online help. Following are descriptions of the last two features.

Automated alert system for user-maintained exception reporting

Alert management is a feature that provides user-defined and user-maintained exception reporting. Through the process of alert management, you define measures that are checked daily to see if any values fall outside of an acceptable range or do not match a given value. When this happens, an alert is generated to let you know that a measure may need to be examined and possibly amended in a workbook.

The Alert Manager is a dialog box that is displayed automatically when you log on to the system. This dialog provides a list of all identified instances in which a given measure's values fall outside of the defined limits. You may pick an alert from this list and have the system automatically build a workbook containing that alert's measure. In the workbook, you can examine the actual measure values that triggered the alert and make decisions about what needs to be done next.

Generation of both final level and source level forecasts

Often, forecast information is required for items at a low level in a hierarchy. Problems can arise when historic sales data for these items is too sparse and noisy to identify clear selling patterns. In such cases, generating a reliable forecast requires aggregating sales data from a low level up to a higher level in the hierarchy. After a forecast is generated at the higher level, the resulting data can be allocated (spread) back down to the lower level, based on the lower level's relationship to the total. Before you can spread forecast data back down to a lower level, you should have an understanding of the relationship between the lower level and the higher level dimensions.

Frequently, an additional forecast will be generated at the low level, to help determine this relationship. This low level is called the final forecast level. Forecast data at this level might be sufficient to generate reliable percentage-to-whole information, but the actual forecast numbers will be more robust when they are generated at an aggregate level. This aggregate level from which forecast data is spread is referred to as the source forecast level.

Some high-volume items may possess sufficient sales data for robust forecast calculations directly at the final forecast level. In these cases, forecast data generated at an aggregate level and then spread down to lower levels can be compared to forecasts run directly at the low level. Comparing the two forecasts, each generated at a different hierarchy level, can be an invaluable forecast performance evaluation tool.

Your Retek Demand Forecasting system may include multiple final forecast levels. Forecast data must appear at some final level for the data to be approved and exported to some other system.

Retek Demand Forecasting Modules

In addition to the standard Administration and Analysis modules there are four components, or modules, associated with the Retek Demand Forecasting application: Predict, Promote, Curve, and Preprocessing,

Predict

The Predict module consists of the workbook templates, measures, and forecasting algorithms that are needed to perform time series forecasting. This includes the Forecast Administration, Forecast Maintenance, Run Batch Forecast, Forecast Approval, Forecast Scorecard, Interactive Forecasting, Forecast Export Administration, Export Forecasts, and Delete Forecast Workbook templates. The Predict module also includes the AutoES batch forecasting routine and all of its component algorithms.

For more information on the Predict Workbooks and Worksheets, see Chapter 2. A detailed discussion of statistical forecasting methods is in Chapter 6.

Promote

The Promote module consists of the templates and algorithms required to perform promotional forecasting, which uses both past sales data and promotional information (for example, advertisements, holidays) to forecast future demand. This module includes the Causal Maintenance, Promotion Planner, Promotion Selector, and Promotion Planning and Analysis templates, as well as all promotional forecasting algorithms.

For more information on the Promote Workbooks and Worksheets, see Chapter 3. A detailed discussion of promotional forecasting methods is in Chapter 6.

Curve

The Curve module consists of the workbook templates and batch algorithms necessary for the creation, approval, and application of curves, or profiles, that are used to explode an assortment plan across a forecasting dimension. This module includes the Profile Administration, Profile Maintenance, Profile Approval, Spread Maintenance, and Run Batch Profile Workbook templates, as well as the profile generation algorithm and all profile-based forecasting algorithms.

For more information on the Curve Workbooks and Worksheets, see Chapter 4. A detailed discussion of profile-based forecasting methods is in Chapter 6.

Preprocessing

The Preprocessing module is intended to be run before running a batch forecast. This module consists of workbook templates and forecasting methods for correcting past data points that represent unusual values not representative of the general demand pattern. For example, if an item is out of stock, it cannot be sold. In that case, the actual sales volume is likely to be lower than the actual demand for that item. If that situation is not corrected in the data, the forecast algorithm will likely replicate the “dip” caused by the past lost sales. The Preprocessing module is provided for you to automatically make adjustments to the raw POS (point of sales) data, so the subsequent demand forecasts do not replicate undesired patterns.

For more information on the Preprocessing Workbooks and Worksheets, see Chapter 5.

Chapter 2 – Predict (Statistical Forecasting)

Overview

The Predict Workbook template group allows you to perform functions related to statistical time series forecasting. This includes defining and maintaining parameters for forecasts, running, reviewing, and approving forecasts, exporting forecasts to other processes, and related tasks.

Workbooks and Wizards

The forecasting workbooks included in the Predict module include:

- Forecast Administration Workbook
- Forecast Maintenance Workbook
- Run Batch Forecast Workbook
- Delete Forecasts Workbook
- Forecast Approval Workbook
- Interactive Forecast Workbook
- Forecast Scorecard Workbook
- Export Forecasts Workbook
- Forecast Export Administration Workbook

Forecasting Methods

A forecasting system's main goal is to produce accurate predictions of future demand. Retek's Demand Forecasting solution utilizes the most advanced forecasting algorithms to address many different data requirements across all retail verticals. Furthermore, the system can be configured to automatically select the best algorithm and forecasting level to yield the most accurate results.

The following section summarizes the use of the various forecasting methods employed in the system. This section is referenced throughout this document when the selection of a forecasting method is required in a workflow process. (See also Chapter 6 for more details on forecasting algorithms):

- Average – Retek Demand Forecasting uses a simple moving average model to generate forecasts.

- AutoES - Retek Demand Forecasting fits the sales data to a variety of exponential smoothing models of forecasting, and the best model is chosen for the final forecast. The candidate methods considered by AutoES are: Simple ES, Intermittent ES, Trend ES, Multiplicative Seasonal, Additive Seasonal, and Seasonal ES. The final selection between the models is made according to a performance criterion (Bayesian Information Criterion) that involves a tradeoff between the model's fit over the historic data and its complexity.
- Simple ES - Retek Demand Forecasting uses a simple exponential smoothing model to generate forecasts. Simple ES ignores seasonality and trend features in the demand data and is the simplest model of the exponential smoothing family. This method can be used when less than one year of historic demand data is available.
- Intermittent ES - Retek Demand Forecasting fits the data to the Croston's model of exponential smoothing. This method should be used when the input series contains a large number of zero data points (that is, intermittent demand data). The original time series is split into a Magnitude and Frequency series, then the Simple ES model is applied to determine level of both series. The ratio of the magnitude estimate over the frequency estimate is the forecast level reported for the original series.
- Simple/Intermittent ES – A combination of the Simple ES and Intermittent ES methods. This method applies the Simple ES model unless a large number of zero data points are present, in which case the Croston's model is applied.
- Trend ES – Retek Demand Forecasting fits the data to the Holt model of exponential smoothing. The Holt model is useful when data exhibits a definite trend. This method separates base demand from trend, then provides forecast point estimates by combining an estimated trend and the smoothed level at the end of the series.
- Multiplicative Seasonal – Also referred to as Multiplicative Winters Model, this model extracts seasonal indices that are assumed to have multiplicative effects on the un-seasonalized series.
- Additive Seasonal – Also referred to as Additive Winters Model, this model is similar to the Multiplicative Winters model, but is used when zeros are present in the data. This model adjusts the un-seasonalized values by adding the seasonal index (for the forecast horizon).
- Seasonal ES – This method, a combination of several Seasonal methods, is generally used for known seasonal items or forecasting for long horizons. This method applies the Multiplicative Seasonal model unless zeros are present in the data, in which case the Additive Winters model of exponential smoothing is used. If less than 2 years of data is available, then a Seasonal Regression model is used. If there is too little data to create a seasonal forecast (in general, less than 52 weeks), then the system will select from the Simple ES, Trend ES and Intermittent ES methods.

- Seasonal Regression – Seasonal Regression cannot be selected as a forecasting method, but is a candidate model used only when the Seasonal ES method is selected. This model requires a minimum of 52 weeks of history to determine seasonality. Simple Linear Regression is used to estimate the future values of the series based on a past series. The independent variable is the series history one-year or one cycle length prior to the desired forecast period, and the dependent variable is the forecast. This model assumes that the future is a linear combination of itself one period before plus a scalar constant.
- Causal – Causal is used for promotional forecasting and can only be selected if Promote is implemented. (See Chapter 3 for more information on Promote.) Typically, the Causal method is used at the Final Levels (for example, item/str/week). Causal uses a Stepwise Regression sub-routine to determine the promotional variables that are relevant to the time series and their lift effect on the series. AutoES utilizes the time series data and the future promotional calendar to generate future baseline forecasts. By combining the future baseline forecast and each promotion's effect on sales (lift), a final promotional forecast is computed.
- No Forecast - No forecast will be generated for the product/location combination.
- Bayesian – Useful for short life cycle forecasting and for new products with little or no historic sales data, the Bayesian method requires a product's known sales plan (created externally to RDF) and considers a plan's shape (the selling profile or life cycle) and scale (magnitude of sales based on Actuals). The initial forecast is equal to the sales plan but as sales information comes in, the model generates a forecast by merging the sales plan with the sales data. The forecast is adjusted so that the sales magnitude is a weighted average between the original plan's scale and the scale reflected by known history.
- Profile Based – Retek Demand Forecasting generates a forecast based on a seasonal profile that can be created in Curve or a legacy system. Profiles can also be copied from another profile and adjusted. Using historic data and the profile, the data is de-seasonalized and then fed to the Simple ES method. The Simple forecast is then re-seasonalized using the profiles.
- Profile Spread – Used at the final level to utilize a profile (either generated externally or with Curve) to determine the spreading ratios from the Source level forecast down to the Final level.

Forecast Administration Workbook

Forecast Administration is the first step in setting up Demand Forecasting to generate forecasts. In Forecast Administration, you establish a final level forecast horizon and a frequency for reviewing information; you also set the forecast parameter default values for both final level and source level forecasts.

Forecast horizon information includes values such as the horizon start date and the number of periods (typically weeks or months) for which sales information is to be forecasted. The forecast review time (number of days between each forecast generation) is also established in the Forecast Administration Workbook. Each of these fields is explained in detail in the Final Level Worksheet section of this chapter.

The Forecast Administration Workbook allows you to set default values for parameters affecting the final level and source level forecasts, thus eliminating the need to define these parameters individually for each product and location in the database. If certain products or locations require parameter values other than the defaults, these fields can be amended on a case-by-case basis in the Forecast Maintenance Workbook.

The Forecast Maintenance Workbook automatically extracts default information from the Forecast Administration Workbook so you can change these parameters as needed. For example, suppose that the majority of your company's products can be forecasted using an AutoES forecast generation method. In the Forecast Administration Workbook, you would set the default generation method to be AutoES. But perhaps the performance of a few products is best predicted using a simple moving average model of forecasting. The Forecast Maintenance Workbook allows you to change the forecast generation method to Average for those few products, overriding the defaults specified in Forecast Administration. For more information on the Forecast Maintenance Workbook functionality, see the Forecast Maintenance Workbook section in this chapter.

The Forecast Administration Workbook consists of a wizard and a workbook with two workflow tabs:

- Basic Settings
- Advanced Settings

Often, forecast information is required for items at a very low level in the hierarchy. Problems can arise, however, in that data is often too sparse and noisy to identify clear patterns at these lower levels. For this reason, it sometimes becomes necessary to aggregate sales data from a low level to a higher level in the hierarchy in order to generate a reasonable forecast. Once this forecast is created at the higher or *source level*, the results can be allocated to the lower or *final level* dimension based on the lower level's relationship to the total.

In order to spread this forecasted information back down to the lower level, it is necessary to have some idea about the relationship between the final level and the source level dimensions. Often, an additional forecast is run at the low level in order to determine this relationship. Forecast data at this low level might be sufficient to generate reliable percentage-to-whole information, but the actual forecast numbers are more robust when generated at the aggregate level.

The Final Level Worksheet represents forecast parameters for the lower (final) level, the level to which source forecast values are ultimately spread. Forecast data must appear at some final level in order for the data to be approved or exported to some other system. The Source Level Worksheet represents the default values for forecast parameters at the more robust aggregate (source) level.

Forecast Administration Wizard

To create a Forecast Administration Workbook, or perform the following steps.

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Administration and click OK.
- 3 The Forecast Administration wizard opens and prompts you to select the level of the final forecast. The final forecast level is a level at which approvals and data exports can be performed. Depending on your organization's setup, you may be offered a choice of several final forecast levels. Make the appropriate selection and click the Finish button to open the workbook.

Basic Settings workflow tab

The Basic Settings workflow tab contains forecast administration settings that may have to be reviewed and changed on a regular basis. On the Basic Settings workflow tab, there are two worksheets:

- Final Level Worksheet
- Source Level Worksheet

Final Level Worksheet – Basic Settings

The Forecast Administration Final Level Worksheet allows you to set the forecast horizon information, frequency of review, and all default parameters for the lower or final level forecast (the level to which aggregate forecast data will ultimately be spread). Forecast approvals and data exports can only be performed on forecasts at a final level.

The screenshot shows a window titled "Final Level Worksheet" with a "Forecast Level" dropdown set to "Item/Store/Week - Final (Simp)". The main area contains a table of settings:

Data Plan	
Data Source	Weekly Sales - Total
Default Approval Policy	Automatic
Default Confidence %	90.00
Default Forecast Method	SimpleES
Default Source Level	Item/Chain/Week - Source
Forecast Cycle (Days)	7
Forecast Length	13
Forecast Start Date	4/5/2002
History Start Date	10/16/1998
Keep Last Changes	None
Label	Item/Store/Week - Final
Next Run Date	9/30/2003
Profile	

At the bottom, there is a "Measure" dropdown and navigation arrows.

Final Level Worksheet – Basic Settings

The following is a description of the Basic Settings parameters contained in the Final Level Worksheet:

Data Plan

A drop-down list from which you can select the sales plan associated with the final level forecast. Sales plans, when available, provide details of the anticipated shape and scale of an item's selling pattern. The sales plan is required when Bayesian forecasting is selected as the forecast generation method.

Data Source

A drop-down list from which you can select the default measure to be used as the data source (for example, POS) for the generation of forecasts.

Default Approval Policy

A drop-down list from which you select the default automatic approval policy for forecast items. Valid values are:

- Manual – The system-generated forecast will not be automatically approved. Forecast values must be manually approved by accessing and amending the Forecast Approval Workbook.
- Automatic – The system-generated quantity will be automatically approved as-is.
- By alert “*name of the alert*” – This is not one alert, but actually a set of alerts, which could be an empty set. If you have an alert hit, the system will use the Manual setting, in other words it will not automatically approve that product/location. If you do not have an alert hit, the system will automatically approve that product/location.

Note: if you select a specific alert as your approval method and later on you delete the alert, the approval will work in the same manner as the *Manual* option. The same will happen if the alert is on a wrong intersection.

Default Confidence %

The percentage used to generate the upper confidence limit of a two-sided confidence interval. For example, actual sales values should fall into a forecast’s 90% confidence interval 90% of the time. Confidence limits are useful guides in analyzing expected forecast performance.

Default Forecast Method

A drop-down list from which you can select the default forecast generation method for items forecasted at the lower level. Typically, a lower-level forecast (interim forecast) is run in tandem with an upper-level forecast, and the information gained by generating the former is used to determine the manner in which forecast data is spread from the higher level. Valid options depend on your system setup. A summary of methods is provided in the overview of Chapter 2 and Chapter 6 covers each method in detail.

Default Source Level

The level at which the aggregate, more robust forecast is run. You can select the desired aggregate level from this measure’s drop-down list. Forecast data from this level is ultimately spread down to the lower (final) level based on the relationship between the two levels in the hierarchy.

Forecast Cycle

The amount of time (measured in days) that the system waits between each forecast generation. Once a scheduled forecast has been generated, this field automatically updates the Next Run Date field.

Forecast Length

This is also commonly referred to as the Horizon Length. It is the number of time periods (typically weeks or months) that the system forecasts beyond the horizon start date. Company preferences determine the length of these time periods; this information is built in Retek Demand Forecasting at the time of installation.

Forecast Start Date

The starting date of the forecast. The default is the system date unless otherwise specified.

History Start Date

The starting date for historical sales data. For example, if your system start date is January 1, 2003, but you only want to use historical sales data from the beginning of 1999, then you need to set your History Start Date to January 1, 1999. Only history after this date is used for generating the forecast. The default is the system start date unless otherwise specified. If sales data is collected weekly or monthly, Retek Demand Forecasting generates forecasts only using data from sales periods after the one containing the history start date.

It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 1999 and an item/location does not have sales history until February 1, 1999, the system will consider the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

Keep Last Changes

A drop-down list from which you select the default change policy for forecast items. Valid values are:

- Keep Last Changes (None) –There are no changes that are introduced into the adjusted forecast.
- Keep Last Changes (Total) –Considers both the Last System Forecast and the Last Approved Forecast in determining change policy. For each forecasted item/week-combination, Retek Demand Forecasting automatically introduces the same quantity that was approved in the Last Approved Forecast into the change, *if that quantity differed from that in the Last System Forecast*. If the quantities are the same, then Retek Demand Forecasting will introduce the current system-generated forecast into the adjusted forecast.

- **Keep Last Changes (Diff)** –Considers both the Last System Forecast and the Last Approved Forecast in determining approval policy. For each forecasted item/week-combination, Retek Demand Forecasting determines the difference between the Last System Forecast and the Last Approved Forecast. This difference (positive or negative) is then added to the current system forecast, and calculated as the adjusted forecast.
- **Keep Last Changes (Ratio)** –Considers both the Last System Forecast and the Last Approved Forecast in determining change. For each forecasted item/week-combination, Retek Demand Forecasting determines the difference between the Last System Forecast and the Last Approved Forecast; this difference is expressed as a percentage. This same percentage is used to calculate the adjusted forecast.

Label

The exact name given to a forecast that is displayed any time the forecast is referenced in the Retek Demand Forecasting application (for example, in a wizard or a workbook).

Next Run Date

The date on which the next batch forecast generation process will automatically be run. Retek Demand Forecasting automatically triggers a set of batch processes to be run at a pre-determined time period, usually at night. When a scheduled batch is run, the Next Run Date automatically updates based on the value in Forecast Cycle.

Profile

A drop-down list from which you can select the seasonal profile associated with the final level forecast. Seasonal profiles, when available, provide details of the anticipated seasonality of an item's selling pattern. The seasonal profile is required when Profile based forecasting is selected as the forecast generation method. The seasonal profile can be generated or loaded, depending on your configuration.

Source Level Worksheet – Basic Settings

The Source Level Worksheet allows you to set the default parameters for the source level forecasts (the aggregate levels from which forecast data is spread to the lower level). A position along the Forecast Level axis exists for every source level defined in the system.

The screenshot shows a window titled "Source Level Worksheet" with a "Forecast Level" tab. It contains a table with four columns representing different forecast levels. The rows define parameters for each level.

	SKU/Chain/Week (AutoES	Style-Color/Chain/Week (/	Style/Chain/Week (AutoE
Data Plan			
Data Source			
Default Confidence %	90.00	90.00	90.00
Default Forecast Method	Simp/Intr ES	Simp/Intr ES	Simp/Intr ES
Label	SKU/Chain/Week	Style-Color/Chain/Week	Style/Chain/Week
Profile			

At the bottom left, there is a "Measure" button and a scroll bar.

Source Level Worksheet – Basic Settings

The following is a description of the Basic Settings parameters contained in the Source Level Worksheet:

Data Plan

A drop-down list from which you can select the sales plan associated with the final level forecast. Sales plans, when available, provide details of the anticipated shape and scale of an item's selling pattern. The sales plan is required when Bayesian forecasting is selected as the forecast generation method.

Default Confidence %

The percentage used to generate the upper confidence limit of a two-sided confidence interval. For example, actual sales values should fall into a forecast's 90% confidence interval 90% of the time. Confidence limits are useful guides in analyzing expected forecast performance.

Default Forecast Method

- A drop-down list from which you can select the default forecast generation method for items forecasted at the lower level. Typically, a lower-level forecast is run in tandem with an upper-level forecast, and the information gained by generating the former is used to determine the manner in which forecast data is spread from the higher level. Valid options depend on your system setup, and may include the following. A summary of methods is provided in the overview of Chapter 2 and Chapter 6 covers each method in detail.

Profile

A drop-down list from which you can select the seasonal profile associated with the final level forecast. Seasonal profiles, when available, provide details of the anticipated seasonality of an item's selling pattern. The seasonal profile is required when Profile based forecasting is selected as the forecast generation method. The seasonal profile can be generated or loaded, depending on your configuration.

Label

Displays the exact name given to a forecast that is displayed any time the forecast is referenced in the Retek Demand Forecasting application (for example, in a wizard or a workbook).

Advanced Settings workflow tab

The Forecast Administration Advanced Settings workflow tab is used to set advanced forecasting parameters. The parameters on this workflow tab aren't as likely to be changed on a regular basis as the ones on the Basic Settings workflow tab.

Final Level Worksheet – Advanced Settings

The Final Level Worksheet allows you to set the advanced parameters for the final level forecasts.

Forecast Level: SKU/Store/Week - Final	
Bayesian Alpha	1.00
Croston's Min Gaps	5
DD Duration (Weeks)	0
Generate Baseline Forecasts	<input checked="" type="checkbox"/>
Generate Cumulative Intervals	<input type="checkbox"/>
Generate Intervals	<input type="checkbox"/>
Generate Methods	<input type="checkbox"/>
Generate Parameters	<input type="checkbox"/>
Holt Min Hist (Periods)	13
Like TS Duration	0
Max Alpha (Profile)	1.00
Max Alpha (Simple, Holt)	1.00
Max Alpha (Winters)	1.00
Max Number of Forecasts	10
Trend Damping Factor	0.50
Updating Last Week Forecast	No Change
Updating Last Week Forecast Number of Weeks	1
Winters Min Hist (Years)	2

Measure ◀ ▶

Final Level Worksheet – Advanced Settings

The Final Level Worksheet – Advanced Settings contains the following parameters:

Croston Min Gaps

The default minimum number of gaps between intermittent sales for Retek Demand Forecasting to consider Croston's as a potential AutoES forecasting method. If there are not enough gaps between sales in a given product's sales history, then the Croston's model is not considered a valid candidate model. If this field is left blank, the default is five minimum gaps between intermittent sales.

Holt Min Hist (Periods)

The default minimum number of periods of historical data necessary for Retek Demand Forecasting to consider Holt as a potential AutoES method of forecasting. Retek Demand Forecasting fits the given data to a variety of AutoES candidate models in an attempt to determine the best method; if not enough periods of data are available for a given item, then Holt will not be considered as a valid option. If this field is left blank, the default is 13 periods.

Generate Cumulative Interval

Specifies whether you want Retek Demand Forecasting to generate cumulative intervals. Select the check box to generate cumulative intervals. Cumulative Intervals are a running total of Intervals. The cumulative interval is necessary if forecast information is to be exported to the Retek Replenishment application. If you do not need cumulative intervals, you can eliminate excess processing time and save disk space by clearing the check box. The calculated cumulative intervals can be viewed within RDF (see Forecast Approval Workbook for more details on the calculation of Cumulative Intervals when the user adjusts forecasts.

Generate Interval

Specifies whether you want Retek Demand Forecasting to generate intervals. Select the check box to generate intervals. Intervals can be displayed in the Forecast Approval Workbook. If you do not need intervals, you can eliminate excess processing time and save disk space by clearing the check box. For many forecasting methods they are calculated as standard deviation but for Simple, Holt, and Winters, the calculation is more complex. (See Chapter 6 for more details.). Intervals are not exported. Also, Intervals are capped.

Generate Parameters

Specifies whether you want Retek Demand Forecasting to output the alpha, level, and trend parameters for each fitted time series. These parameters can be displayed in the Forecast Approval Workbook. Select the check box if you want these parameters generated, or clear it if you do not.

Generate Methods

Specifies whether you want Retek Demand Forecasting to output the chosen method for each fitted time series. The chosen method can be displayed in the Forecast Approval Workbook. Select the check box if you want these methods generated, or clear it if you do not.

Generate Baseline Forecasts

Specifies whether you want Retek Demand Forecasting to output the baseline forecast. This parameter must be selected if you are using Promote to support promotional forecasting. (See Chapter 3 for more information on Promote.)

Winters Min Hist (Years)

The default minimum number of years of historical data necessary for Retek Demand Forecasting to consider Winters as a potential AutoES method of forecasting. If not enough years of data are available for a given item, then Winters will not be considered as a valid forecasting model. If this field is left blank, the default is two years of required history.

Max Alpha (Simple, Holt) (range (0,1])

In the Simple or Holt model fitting procedure, alpha, which is a model parameter capturing the level, is determined by optimizing the fit over the time series. This parameter displays the maximum value (that is, cap value) of alpha allowed in the model fitting process. An alpha cap value closer to 1 allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to 0 only allows less reactive models. The default is 1.

Max Alpha (Winters) (range (0,1])

In the Winters model fitting procedure, alpha, which is a model parameter capturing the level is determined by optimizing the fit over the time series. This parameter displays the maximum value (that is, cap value) of alpha allowed in the model fitting process. An alpha cap value closer to 1 allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to 0 only allows less reactive models. The default is 1.

Max Alpha (Profile) (range (0,1])

In the Profile-based model-fitting procedure, alpha, which is a model parameter capturing the level, is determined by optimizing the fit over the deseasonalized time series. The time series is deseasonalized based on a seasonal profile created by the Curve tool. This parameter displays the maximum value (that is, cap value) of alpha allowed in the model fitting process. An alpha cap value closer to 1 allows more reactive models (alpha = 1, repeats the last data point) whereas alpha cap closer to 0 only allows less reactive models. The default is 1.

Bayesian Alpha (range (0,infinity))

In the Bayesian Demand Forecasting method, historic data is combined with a known sales plan in creating the forecast. As POS data comes in, a Bayesian forecast is adjusted so that the sales magnitude is a weighted average between the original plan's scale and the scale reflected by known history. This parameter displays the value of alpha (the weighted combination parameter). An alpha value closer to 0 weights the sales plan more in creating the forecast whereas alpha closer to infinity weights the known history more. The default is 1.

Max Number of Forecasts

The maximum number of generated forecasts that will be kept in the system. If the number of generated forecasts in the system is more than the max number of forecasts, the forecasts are deleted in FIFO order (that is, the oldest generated forecast is deleted first).

Updating Last Week Forecast

A drop-down list from which you can select the method for updating the last Updating Last Week Forecast Number of Weeks week(s) of the forecast horizon for Approved Forecast measure. This option is valid only if the approval method was selected to be Manual (or approve by alert and alert was rejected) in FMA.

- **No Change** – The last Updating Last Week Forecast Number of Weeks week(s) of the forecast horizon for Approved Forecast measure are not updated.
- **Replicate** – The last Updating Last Week Forecast Number of Weeks week(s) of the Approved Forecast will be filled in by replicating the value of the Forecast Length minus Updating Last Week Forecast Number of Weeks week of the Approved Forecast measure.
- **Use Forecast** – The last Updating Last Week Forecast Number of Weeks week(s) of the Approved Forecast measure will be filled in using the Adjusted Forecast value(s).

Updating Last Week Forecast Number of Weeks

The number of weeks that have to be updated in the Approved Forecast using the method specified from the updating method pick list.

Like TS Duration (weeks)

The number of weeks of history required after which Retek Demand Forecasting stops using the substitution method and starts using the system forecast generated by the forecast engine. A value must be entered in this parameter if using Like SKU/Sister Store functionality. (See Forecast Maintenance Workbook: Advanced Settings in this chapter for more information on Like SKU / Sister Store functions.)

DD Duration (weeks)

The number of weeks of history required after which Retek Demand Forecasting stops using the DD approach and defaults to the “normal” Profile-Based method for all time series, for which the Default Forecast Method is set to Profile Based.

Trend Damping Factor (range (0,1))

Displays the default value of the trend damping factor. A value close to 0 is a high damping while a value of 1 implies no damping. The default is 0.5.

Source Level Worksheet – Advanced Settings

The Source Level Worksheet allows you to set the advanced parameters for the source level forecasts.

	SKU/Chain/W/Week (AutoES)	Style-Color/Chain/W/Week (AutoES)	Style/Chain/W/Week (AutoES)
Bayesian Alpha	1.00	1.00	1.00
Crostons Min Gaps	5	5	5
Holt Min Hist (Periods)	13	13	13
Max Alpha (Profile)	1.00	1.00	1.00
Max Alpha (Simple, Holt)	1.00	1.00	1.00
Max Alpha (Winter)	1.00	1.00	1.00
Trend Damping Factor	0.50	0.50	0.50
Winters Min Hist (Years)	2	2	2

Source Level Worksheet – Advanced Settings

The Source Level Worksheet – Advanced Settings contains the following parameters:

Croston Min Gaps

The default minimum number of gaps between intermittent sales for Retek Demand Forecasting to consider Croston's as a potential AutoES forecasting method. If there are not enough gaps between sales in a given product's sales history, then the Croston's model is not considered a valid candidate model. If this field is left blank, the default is five minimum gaps between intermittent sales.

Holt Min Hist (Periods)

The default minimum number of periods of historical data necessary for Retek Demand Forecasting to consider Holt as a potential AutoES method of forecasting. Retek Demand Forecasting fits the given data to a variety of AutoES candidate models in an attempt to determine the best method; if not enough periods of data are available for a given item, then Holt will not be considered as a valid option. If this field is left blank, the default is 13 periods.

Winters Min Hist (Years)

The default minimum number of years of historical data necessary for Retek Demand Forecasting to consider Winters as a potential AutoES method of forecasting. If not enough years of data are available for a given item, then Winters will not be considered as a valid forecasting model. If this field is left blank, the default is two years of required history.

Max Alpha (Simple, Holt) (range (0,1])

In the Simple or Holt model fitting procedure, alpha, which is a model parameter capturing the level, is determined by optimizing the fit over the time series. This parameter displays the maximum value (that is, cap value) of alpha allowed in the model fitting process. An alpha cap value closer to 1 allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to 0 only allows less reactive models. The default is 1.

Max Alpha (Winters) (range (0,1])

In the Winters model fitting procedure, alpha, which is a model parameter capturing the level, is determined by optimizing the fit over the time series. This displays the maximum value (that is, cap value) of alpha allowed in the model fitting process. An alpha cap value closer to 1 allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to 0 only allows less reactive models. The default is 1.

Max Alpha (Profile) (range (0,1])

In the Profile based model fitting procedure, alpha, which is a model parameter capturing the level, is determined by optimizing the fit over the deseasonalized time series. The time series is deseasonalized based on a seasonal profile created by the Curve tool. This displays the maximum value (that is, cap value) of alpha allowed in the model fitting process. An alpha cap value closer to 1 allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to 0 only allows less reactive models. The default is 1.

Bayesian Alpha (range (0,infinity))

In the Bayesian Demand Forecasting procedure, historic data is combined with a known sales plan in creating the forecast. As POS data comes in, a Bayesian forecast is adjusted so that the sales magnitude is a weighted average between the original plan's scale and the scale reflected by known history. This parameter displays the value of alpha (the weighted combination parameter). An alpha value closer to one (or infinity) weights the sales plan more in creating the forecast, whereas alpha closer to zero weights the known history more. The default is 1.

Trend Damping Factor (range (0,1])

The default value of the trend damping factor. A value close to 0 is a high damping, while a value if 1 implies no damping. The default is 0.5.

Forecast Maintenance Workbook

The Forecast Maintenance Workbook allows you to select and modify forecasting parameters for product/location combinations when the values of these parameters differ from the default values that are assigned in the Forecast Administration Workbook.

Suppose, for example, that the default forecast method of all the products in the database was set in the Forecast Administration Workbook to be AutoES. For a particular product, however, you know that a seasonal model is a better predictor of sales performance. To make this change, you must access the Forecast Maintenance Workbook, select the product / location intersection to be reviewed, make the appropriate change to the forecast method, and commit the workbook to the Retek Demand Forecasting master database.

The Forecast Maintenance Workbook is split into two workflow tabs:

- Basic Settings – Includes all Final Level Worksheets and their respective Source Level Worksheets.
- Advanced Settings – Interfaces with the Substitute Item functionality. This functionality manages cases that arise when forecasting new product introductions and new store openings. This tab includes the measures Like SKU, Sister Store, Adjustment %, Substitute Method, Start Date and End Date measures.

Forecast Maintenance Wizard

To create a Forecast Maintenance Workbook, perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Maintenance and click OK.
- 3 The Forecast Maintenance Wizard opens and prompts you to select the level of the final forecast. The final forecast level is a level at which approvals and data exports can be performed. Depending on your organization's setup, you may be offered a choice of several final forecast levels. Make the appropriate selection and click Next.

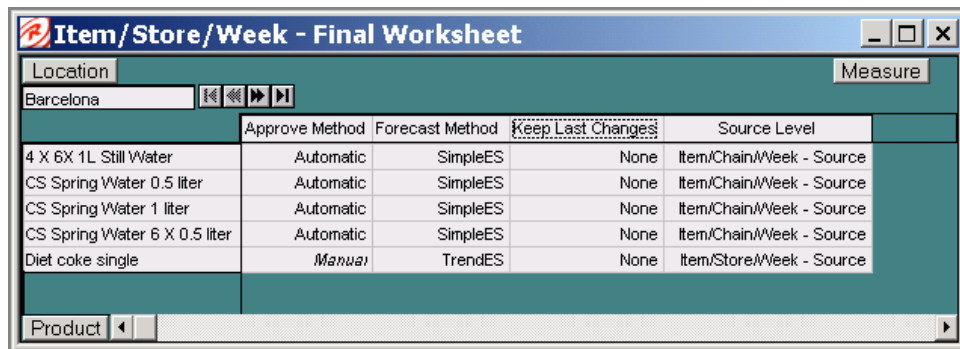
- 4 Select the products whose parameters you want to update. Highlight the appropriate products on the Available Products side of the screen and click the right arrow button to move them to the Selected Products side. To access a quick menu of additional functions, such as changing the product rollout, click the right mouse button. Click Next.
- 5 Now perform the same operation for locations in the location hierarchy. Highlight the appropriate locations on the Available Locations side of the screen and click the right arrow button to move them to the Selected Locations side. Click Next.
- 6 Additional wizard screens may prompt you to select any additional measures (that is, measures not standard in the Forecast Maintenance Workbook) that you would like included. Make the appropriate selections (if any) and click the Finish button to display the workbook.

Basic Settings workflow tab

The Basic Settings workflow tab of Forecast Maintenance includes all Final Level Worksheets and their respective Source Level Worksheets.

Final Level Worksheet – Basic Settings

The Final Level Worksheet allows you to view and set the basic parameters for the final level forecasts.



Item	Approve Method	Forecast Method	Keep Last Changes	Source Level
4 X 6X 1L Still Water	Automatic	SimpleES	None	Item/Chain/Week - Source
CS Spring Water 0.5 liter	Automatic	SimpleES	None	Item/Chain/Week - Source
CS Spring Water 1 liter	Automatic	SimpleES	None	Item/Chain/Week - Source
CS Spring Water 6 X 0.5 liter	Automatic	SimpleES	None	Item/Chain/Week - Source
Diet coke single	Manual	TrendES	None	Item/Store/Week - Source

Forecast Maintenance Final Level Worksheet – Basic Settings

The Final Level Worksheet – Basic Settings contains the following parameters:

Approval Method

A drop-down list from which you select the approval policy for individual product/location combinations. Valid values are:

- **Manual** – The System Forecast and Adjusted Forecast are not automatically approved. Forecast values must be manually approved by accessing and amending the Forecast Approval Workbook.
- **Automatic** – The Adjusted Forecast is automatically approved as-is.
- **By alert “*name of the alert*”** - This is not one alert, but actually a set of alerts, which could be an empty set. If you have an alert hit, the system will use the Manual setting, in other words it will not automatically approve that product/location. If you do not have an alert hit, the system will automatically approve that product/location.

Note: If you select a specific alert as your approval method and later on you delete the alert, the approval will work as manual. The same will happen if the alert is on a wrong intersection.

Forecast Method

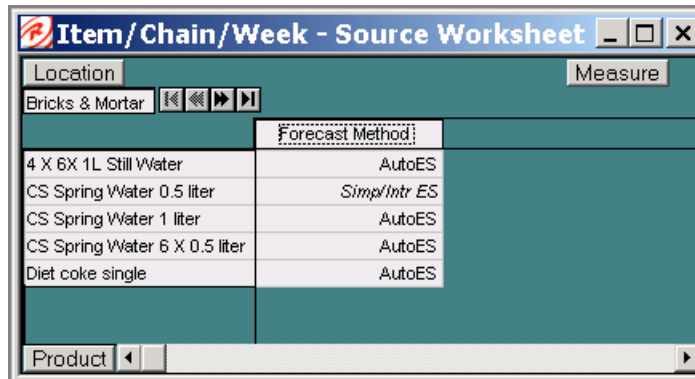
- A drop-down list from which you can select the forecast generation method for items forecasted at the lowest level. Valid options depend on your system setup, and may include the following. A summary of methods is provided in the overview of Chapter 2, and Chapter 6 covers each method in detail.

Source Level

The level at which the aggregate, more robust forecast is run. Forecast data from this level is spread down to the lowest level, based on the relationship between the two levels in the hierarchy. Choose “Auto” if the best level is to be chosen automatically. See Auto Source for more details.)

Source Level Worksheet – Basic Settings

The Source Level Worksheet allows you to modify the forecast method for intersections at source level forecast that vary from the default method. There should be a Source Worksheet for every Source Level associated to the Final Level that was selected in the wizard process.



Forecast Maintenance Source Level Worksheet – Basic Settings

The Source Level Worksheet – Basic Settings contains the following parameters:

Forecast Method

- Displays a drop-down list from which you can select the forecast generation method for items forecasted at the lowest level. A summary of methods is provided in the overview of Chapter 2, and Chapter 6 covers each method in detail.

Advanced Settings workflow tab

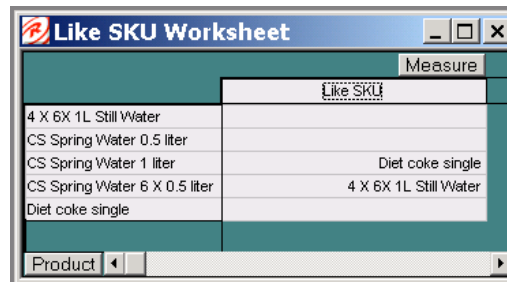
The Advanced Settings workflow tab is used to set substitute item, like SKU and like Store functionality and also allows you to override dates to generate and forecast and history start dates for any item/location that varies from the default settings in the Forecast Administration Workbook. This tab includes the measures such as Like SKU, Sister Store, Adjustment %, Substitute Method, Forecast Start Date, Forecast End Date, and History Start Date.

The Advanced Settings workflow tab displays the following worksheets:

- Like SKU Worksheet – Advanced Settings
- Sister Store Worksheet – Advanced Settings
- Advanced Parameter Worksheet – Advanced Settings

Like SKU Worksheet

The Like SKU Worksheet is used in combination with the Advanced Parameter Worksheet to forecast a new item by modeling it after an existing item.



Like SKU Worksheet

The Like SKU Worksheet contains the following parameter:

Like SKU

Displays the SKUs selected during the wizard process. In the example shown in the above worksheet, the Like Item (existing item) is selected on the right from a pick-list, across from the new item. The Like Item's forecast or sales will be used for the new item based on the parameters selecting in the Advance Parameter Worksheet. In the example above, CS Spring Water 1 liter is a new item and Diet Coke single is being selected as the existing item that will be used to model CS Spring Water 1 liter's forecast.

Sister Store Worksheet

The Sister Store Worksheet is used in combination with the Advanced Parameter Worksheet to forecast demand for a new store by modeling it after an existing store.



Sister Store Worksheet

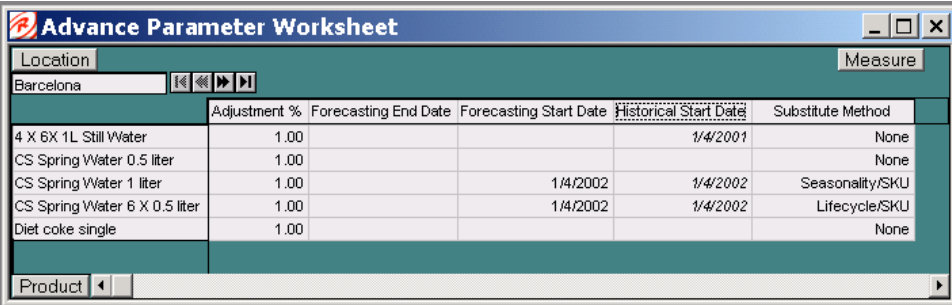
The Sister Store Worksheet contains the following parameter:

Sister Store

Displays the SKUs selected during the wizard process. In the example shown in the above worksheet, the Sister Store (existing Store) is selected on the right from a pick-list, across from the new Store. The Sister Store's forecast or sales will be used for the new Store based on the parameters selecting in the Advance Parameter Worksheet. In the screen shot above, the Seattle is the new store and Minneapolis is being selected as the store to be used for modeling Seattle's forecast.

Advanced Parameter Worksheet

The Advanced Parameter Worksheet is used to set additional parameters required to support Like SKU / Sister Store functionality and to manage item/locations whose Forecast Start Date, Forecast End Date or History Start Date varies from the default settings in the Forecast Administration Workbook. Since this workbook can be used to support many processes it is important to pay close attention to the various usage of each parameter. These are defined in detail below and examples are provided where appropriate.



	Adjustment %	Forecasting End Date	Forecasting Start Date	Historical Start Date	Substitute Method
4 X 6X 1L Still Water	1.00			1/4/2001	None
CS Spring Water 0.5 liter	1.00				None
CS Spring Water 1 liter	1.00		1/4/2002	1/4/2002	Seasonality/SKU
CS Spring Water 6 X 0.5 liter	1.00		1/4/2002	1/4/2002	Lifecycle/SKU
Diet coke single	1.00				None

Advanced Parameter Worksheet

The Advanced Parameter Worksheet contains the following parameters:

Adjustment %

This parameter is relevant to Like SKU / Sister Store functionality. The user may enter an adjustment % to apply to the forecast for the new product / location combination displayed. This is a real number between $[0, \infty)$. The default (NA) value is 100%; that is, 1, which translates to no adjustment.

Forecast Start Date

This parameter represents the date to start forecasting for an item/location combination. This parameter can be set in the future if using Like Sku or Sister Store functionality and upon reaching that time, the forecast will be generated. If this date is set to the past, it is ignored in favor of the Forecast Start Date from the Forecast Administration Workbook. This means that you do not need to go in and change the Forecast Start Date once it is no longer in the future. For Like SKU or Sister Store the Forecast Start Date and History Start Date should be set to the same date. It is important to also understand how Forecast Start Date should be used in conjunction with Forecast End Date (see below).

Forecast End Date

This parameter represents the last point in time for which the Forecasting Engine will forecast for an item/location combination. Should this parameter be set to a date less than the Forecast Start Date plus the Forecast Length (in Forecast Administration), the engine will forecast 0 past this date. If Forecast End Date is more than Forecasting Start Date plus Forecasting Length, you do NOT get a forecast outside Forecasting Start Date plus Forecasting Length. In other words, both Forecast Start Date and Forecasting End Date are relevant for time periods within the forecast horizon set at the global level.

Forecast End Date can be used for new item or location forecasting if the item or location needs to be forecasted for a period shorter than the Like TS Duration (set globally in Forecast Administration).

History Start Date

This parameter represents the first point in time from which the Forecasting Engine will begin training and modeling (i.e., if there are 2 years of history but you only want to use one year, you set the start date to a year ago). This parameter overrides at the sku/store level from the global settings in the Forecast Administration Workbook. This can be used to level out past sales. For example, if you have a large spike in the first 3 weeks of sales for an item was on sale, you can set the Historical Start Date to one week past that period, and those first few weeks will not be used when generating the forecast.

It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 1999 and an item/location does not have sales history until February 1, 1999, the system will consider the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

Unless you are using Like SKU or Sister Store, if this parameter is set into the future, there would be no modeling as the training window is read as zero (for example, there will be no forecast). If you are using Like SKU or Sister Store, the History Start Date for the new item or new store should be set with the same date as the Forecast Start Date.

Note: When using the ‘Lifecycle Method’ (see Substitute Methods below) the History Start Date for the substitute item or location must be set to the point in the sales history that the new item or location will begin using as its sales. In the screen shot above, ‘CS Spring Water 6X .05 Liter’ is a new item using the Seasonality_SKU method and ‘4 X 6X 1L Still Water’ is its substitute item. The History Start Date for the substitute item is set to one year prior to the Forecast Start Date. Therefore the System Forecast for the new item will be identical to the sales for the substitute item one year ago in its history.

Substitute Method

Displays a drop-down list from which you can select the substitute method. Valid options are:

- None – There is no substitution for this product/location combination. This is the default value.
- Seasonality/New Item - You provide a Like Item that has a similar seasonality pattern that sells at the same store. The new product’s forecast will be the like item’s demand forecast with the applied adjustment. The forecast will be set to 0 for all dates before the new product’s start date.
- Seasonality/New Store - You provide a Sister Store that has a similar seasonality pattern that sells the same product. The product’s forecast at the new store will be the demand forecast of the same product at the Sister Store with the applied adjustment. The forecast will be set to 0 for all dates before the new store’s open date.
- Seasonality/New Item /New Store – You provide a Like Item that sells at a Sister Store that has a similar seasonality pattern. The new product’s forecast at the new store will be the demand forecast of the Like Item at the Sister Store with the applied adjustment. The forecast will be set to 0 for all dates before the new product’s start date after the new store opens.
- Lifecycle/New Item - You provide a Like Item that had a similar lifecycle pattern that sells at the same store. The new product’s forecast will be the like items actual sales with the applied adjustment shifted such that the like item’s first sales matches the new product’s start date.
- Lifecycle/New Store - You provide a Sister Store that had a similar lifecycle pattern that sells the same product. The product’s forecast at the new store will be the products actual sales at the Sister Store with the applied adjustment shifted such that the Sister Store’s first sales matches the new store’s open date.
- Lifecycle/New Item /New Store – You provide a Like Item that sells at a Sister Store that has a similar lifecycle pattern. The new product’s forecast at the new store will be Like Item’s actual sales at a Sister Store with the applied adjustment shifted such that the Like Item’s first sales at the Sister Store matches the new product’s start date after the new store opens.

Steps Required for Forecasting using Like SKU / Sister Store Methods

The following outlines the steps required for using each of the above Substitution Methods:

To support Like SKU / Sister Store functionality the 'Like TS Duration' must be set in the Forecast Administration Workbook – Advance Tab. This parameter sets the number of weeks of history required after which Retek Demand Forecasting stops using the substitution method and starts using the system forecast generated by the forecast engine.

- Seasonality/SKU: Introduction of a new item at an existing store (like SKU with a similar forecast)
 - Like SKU Worksheet: Select a like SKU from the drop-down list, across from the new SKU.
 - Advanced Parameter Worksheet: Set the Forecast Start Date for the new SKU at an existing store.
 - Advanced Parameter Worksheet: Set the History Start Date for the new SKU at the existing store to the same date as the Forecast Start Date.
 - Advance Parameter Worksheet: Set the Adjustment % (optional) for the new SKU at the existing store.
- Seasonality/STR: Introduction of an existing SKU at a new store (sister store with similar forecast)
 - Sister Store Worksheet: Select a Sister Store from the drop-down list, across from the new store.
 - Advanced Parameter Worksheet: Set the Forecast Start Date for the existing SKU at the new store.
 - Advanced Parameter Worksheet: Set the History Start Date for the existing SKU at the new store to the same date as the Forecast Start Date.
 - Advance Parameter Worksheet: Set the Adjustment % (optional) for the existing SKU at the new store.

- Seasonality/SKU_STR: Introduction of a new SKU at a new store (like SKU and sister store with a similar forecast).
 - Like SKU Worksheet: Select a like SKU from the drop-down list, across from the new SKU.
 - Sister Store Worksheet: Select a Sister Store from the drop-down list, across from the new store.
 - Advanced Parameter Worksheet: Set the Forecast Start Date at the intersection of the new SKU and the new store.
 - Advanced Parameter Worksheet: Set the History Start at the intersection of the new SKU and new store.
 - Advance Parameter Worksheet: Set the Adjustment % (optional) at the intersection of the new SKU and new store.
- Lifecycle/SKU - Introduction of a new SKU at an existing store (Like SKU's sales history to be used as the forecast for the new item)
 - Like SKU Worksheet: Select a like SKU from the drop-down list, across from the new SKU.
 - Advanced Parameter Worksheet: Set the Forecast Start Date for the new SKU at the existing store.
 - Advanced Parameter Worksheet: Set the History Start Date for the new SKU at the existing store to the same date as the Forecast Start Date.
 - Advanced Parameter Worksheet: Set the History Start Date for the Like SKU at the existing store to the point in its sales history that will map to the new SKU's forecast.
 - Advance Parameter Worksheet: Set the Adjustment % (optional) for the new SKU at the existing store.

- Lifecycle/STR: Introduction of an existing SKU at new store (Sister Store's sales history to be used as the forecast for the new store)
 - Sister Store Worksheet: Select a Sister Store from the drop-down list, across from the new Store.
 - Advanced Parameter Worksheet: Set the Forecast Start Date for the existing SKU at the new store.
 - Advanced Parameter Worksheet: Set the History Start Date for the existing SKU at the new store to the same date as the Forecast Start Date.
 - Advanced Parameter Worksheet: Set the History Start Date at the intersection of the Sister Store and existing SKU to the date in its sales history that will map to the new store's forecast.
 - Advance Parameter Worksheet: Set the Adjustment % (optional) for the existing SKU at the new store.
- Lifecycle/SKU_STR: Introduction of a new item at a new store (Like SKU's and Sister Store's sales history to be used as the forecast for a new item at a new store)
 - Like SKU Worksheet: Select a like SKU from the drop-down list, across from the new SKU.
 - Sister Store Worksheet: Select a Sister Store from the drop-down list, across from the new store.
 - Advanced Parameter Worksheet: Set the Forecast Start Date at the intersection of the new SKU and new store.
 - Advanced Parameter Worksheet: Set the History Start Date at the intersection of the new SKU and new store to the same date as the Forecast Start Date.
 - Advanced Parameter Worksheet: Set the History Start Date at the intersection of the Like SKU and Sister Store to the date in its sales history that will map to the new SKU's and new store's forecast.
 - Advance Parameter Worksheet: Set the Adjustment % (optional) at the intersection of the new SKU and new store.

Run Batch Forecast

The forecast generation process creates sales forecasts for all product/location combinations whose forecast generation method is set to any method other than No Forecast. Forecasts are typically run automatically as scheduled batch jobs. Retek Demand Forecasting regularly triggers a set of processes to be run at a pre-determined time when system use is at a minimum, such as overnight.

Scheduling of the automatic batch forecasting process is performed in the Forecast Administration Workbook, where a default value is set for the forecast cycle (number of days between forecast runs). The Forecast Cycle measure and Next Run Date field in the Forecast Administration Workbook control the automatic scheduling of batch forecasting jobs in the Retek Demand Forecasting application. Refer to the Forecast Administration Workbook description for further information.

The Run Batch Forecast wizard allows you to manually start the forecast generation process. To kick off a forecast run at a time other than the regular overnight batch job, you must perform the following tasks:

- 1 Create or access a Forecast Administration Workbook. In the Forecast Start Date measure, enter the starting date of the forecast horizon.
- 2 Access the Run Batch Forecast Wizard to begin the batch forecast process for all product/location combinations set to some method of forecast generation other than No Forecast.

Note: Ensure that there are no other users logged on to the system. Other users cannot access the system until the forecast generation process has run to completion.

Batch Forecast Wizard

To run a batch forecast manually, perform the following steps.

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Run Batch Forecast and click OK.

The Run Batch Forecast Wizard opens and prompts you to select the level of the final (approval level) forecast. Click Next.

- 3 Select which batch type to run.

- **Scheduled Runs and Overrides.** Generates all regularly scheduled batch forecasts set to be generated on today's date. When a scheduled batch forecast is run, the Next Run Date measure in the Forecast Administration Workbook is updated with an appropriate value, based on the established Forecast Cycle. If you select this option (Scheduled Runs and Overrides) and no batch runs are scheduled for this date, the system reads the Forecast Start Date field of the Forecast Administration Workbook. If the Forecast Start Date contains a value, a forecast will be generated now (based on the established Forecast Start Date value) and the Next Run Date value will be updated based on the Forecast Cycle length. If the Next Run Date is set to some future date, then no forecast will be run.

Note: After a forecast is generated, the Forecast Start Date field is cleared. This ensures that the same forecast will not be generated again on the same date.

- **Overrides Only.** Causes the system to consider only the Forecast Start Date field in the Forecast Administration Workbook. If Forecast Start Date contains a value, a forecast will be generated now. The Next Run Date value is not updated; as a result, any regularly scheduled run for this date will be performed during the overnight batch process. If no value is set for the Forecast Start Date (that is, Forecast Start Date = NA), then no forecast will be generated.

- 4 Click Next to begin the forecast generation process.

Delete Forecasts

The Delete Forecasts Wizard guides you through the process of deleting unwanted forecasts from the system. Deletion of forecasts is permanent, so you should only delete those forecasts that are no longer needed.

Delete Forecasts Wizard

To access the Delete Forecasts Wizard to permanently remove a forecast from the Retek Demand Forecasting system, perform the following steps.

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Delete Forecasts and click OK.
- 3 Select the generation date of the forecast you want to delete. Click Next.
- 4 Select Yes to verify forecast deletion, or No to cancel deletion and exit the Forecast Deletion Wizard. Remember that deletion of forecasts from Retek Demand Forecasting is permanent!

- 5 Click Finish to process the request. If Yes was selected on the final wizard screen, the forecast is deleted.

Forecast Approval Workbook

After the forecast is generated, the next steps in the forecasting process are analysis and approval. Approval of forecasts is required before the forecasted data can be exported to other processes, such as replenishment programs. The Forecast Approval Workbook allows you to view, analyze, correct, and approve forecast results.

Some system forecasts may be set to be automatically approved by Retek Demand Forecasting. The default approval policy for items in a forecast is set in the Forecast Administration Workbook, and these policies can be amended and maintained for individual product/location combinations in the Forecast Maintenance Workbook. Any forecasts not automatically approved must be accessed, evaluated, possibly amended, and ultimately approved before subsequent processes are executed.

You can view and analyze forecast data at multiple forecast levels (source level and final level) simultaneously. Revisions to and approvals of final level forecast values can be made on the appropriate worksheets in the Forecast Approval Workbook. The Forecast Approval Workbook can contain up to five types of worksheets:

- Final Forecast Worksheet – Allows you to review final level system-forecasted quantities and make revisions to them if needed
- Source Level Worksheet – Displays the system-generated source level forecast, and allows you to compare this data with final level forecast values.
- Approval Worksheet – Allows you to specify the manual approval policy of forecasts by product and location.
- Final Parameters Worksheet – this option is only available if “Generate System Parameters” or “Generate Methods” is turned on in the Forecast Administration Workbook.
- Confidence Worksheet – Allows you to set the percent of confidence in both final level and source level forecasts.

When the Forecast Approval Workbook is displayed, you must approve the system-generated forecast for each product or overwrite it with an amount you feel is more reasonable. Forecast values are overwritten in the Final Forecast measure of the Final Forecast Worksheet. Approvals are made for each product/location combination in the Approval Method measure of the Forecast Approval Worksheet.

After you complete your work, you can save the workbook using the Save function on the File menu. To update the master database with the approved forecast values, you must commit the workbook using the Commit Now or Commit Later option on the File menu. Once the workbook is committed, the forecast values are stored in the master database and can be used by other processes.

Forecast Approval Wizard

To access Forecast Approval, select Open from the File menu to bypass the Forecast Approval Wizard and open an existing Forecast Approval Workbook, or perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Approval and click OK.

The Forecast Approval wizard opens and prompts you to select the final level at which to approve forecast values. Make your selection and click Next.

- 3 Select the forecast level(s) to include in the Forecast Approval Workbook. The final forecast level selected on the previous screen will appear in your workbook, plus any source-level forecasts chosen here. If the forecast was not run at a source level as well as a final level, then this screen will not appear. Select as many forecast levels as necessary for comparison, then click Next.
- 4 Now select the generation date of the forecast you wish to approve. Click 'Last Chronological Forecast' to build a workbook containing the most recent forecast values, or click 'List of all generated forecast dates' to select from a list of all previously generated forecasts. Click Next.
- 5 Select the specific products whose forecasts you want to view. Highlight the products on the Available Products side of the screen and click the right arrow button to move them to the Selected Products side. You can access a quick menu of additional functions, such as changing the view of the product rollup, by clicking the right mouse button. Click Next.
- 6 Perform the same operation for locations that you want to view. Highlight the locations on the Available Locations side of the screen and click the right arrow button to move them to the Selected Locations side. Click Next.
- 7 Select whether you want to view data for any time periods before the forecast horizon. If you select Yes on this wizard screen, you will be prompted to select the specific time periods you want to include. If you select No, a calendar selection wizard will not be displayed. Select the appropriate radio button and click Next.

- 8 From the list provided, check any additional forecasting measures specific to this forecast that you would like to view. Click Next.
- 9 Next, highlight any additional registered measures you would like to view in your workbook. The measures available on this wizard screen are not specific to this forecast, and can be chosen for any new workbooks. If 'Generate Intervals' or 'Generate Cumulative Intervals' was selected in the Forecast Administration Workbook, the associated measures can be selected for view.
- 10 Click Finish to build and open your workbook.

Final Forecast Worksheet

The Final Forecast Worksheet allows you to review the system-forecasted quantities and make revisions to those forecasted amounts if needed.

Your main objective in the Forecast Approval Workbook is to review and edit forecast values using the Adjusted Forecast field on the Final Level Worksheet and, ultimately, to approve forecasts.

Product	Location	01/04/02	01/11/02	01/18/02	01/25/02	02/01/02	02/08/02	02/15/02	02/22/02
CS Spring Water 0.5 liter	London-Kensington								
Weekly Sales - Total		117	120	117	104	116	105	117	115
Adjusted Forecast		110.83	111.29	118.50	105.51	109.81	109.29	114.83	109.32
Approved Forecast		110.83	111.29	118.50	105.51	109.81	109.29	114.83	109.32
Last Approved Forecast		110.83	111.29	118.50	105.51	109.81	109.29	114.83	109.32
System Forecast		110.83	111.29	118.50	105.51	109.81	109.29	114.83	109.32
Upper Confidence Limit		168.55	192.92	218.47	220.95	238.88	250.67	267.54	272.57

Forecast Approval – Final Level Worksheet

The following is a description of the standard measures that are contained in the Final Level Worksheet. Additional registered measures may appear here if you made any such selections in the Forecast Approval wizard.

Adjusted Forecast

The user-adjusted forecast values. To fill the Adjusted Forecast field, you can accept those values predicted by the Demand Forecasting system, or you can overwrite System Forecast values with values you feel are more reasonable. Entering data in this field does not change the System Forecast values.

System Forecast

The quantity that the system predicts will be required for the product, location, and calendar combination displayed.

Approved Forecast

The forecast quantity that was approved for the product/location combination during the most recent approval.

Last Approved Forecast

The forecast quantity that was approved for the product/location combination the last time an approval was done.

Weekly Sales (or the name of your sales measure)

This is the sales measure used to generate the forecast. This allows the user to compare Actuals to forecasted values.

Upper Confidence Limit (if selected)

The upper bound of a normal probability distribution that is centered about the forecast value. It is calculated based on the Default Confidence % parameter value entered in the Forecast Administration Workbook, or on the Confidence Level value entered in the Confidence Worksheet of the Forecast Approval Workbook (if such a value has been entered to override the default).

System Baseline

A forecast generated by the system based on past sales data that contains no promotions. To see a generated Baseline Forecast, Promote must be implemented and Generate Baseline must be selected in the Forecast Administration Workbook.

Intervals and Cumulative Intervals

When a user adjusts the forecast for week, the cumulative intervals for all the subsequent weeks will be recalculated according to the following formula:

$$CumInt_{Adj}(t) = CumInt_{Adj}(t-1) + CAP \left[\frac{F_{Adj}(t)}{F_{Sys}(t)} \times [CumInt_{Sys}(t) - CumInt_{Sys}(t-1)] \right]$$

Where:

t : $W, W+1, W+2, \dots$, End of forecast horizon

$CumInt_{Adj}(t)$: Adjusted Cumulative Interval for week t

$CumInt_{Sys}(t)$: System Cumulative Interval for week t

$F_{Sys}(t)$: System Forecast for week t

$F_{Adj}(t)$: Adjusted Forecast for week t

$CAP[x]$: A capping function that limits x to $0.3 \times F_{Adj}(t) \leq CAP[x] \leq F_{Adj}(t)$

For more information on interval calculation, see Chapter 6.

Source Level Worksheet

The Source Level Worksheet displays the system-generated source level forecast. Source level forecast values are useful for evaluating the accuracy of system forecasts. Final level forecast values in the Final Level Worksheet can be viewed alongside and compared with their corresponding source level forecasts.

	01/04/02	01/11/02	01/18/02	01/25/02	02/01/02	02/08/02	02/15/02	02/22/02
Weekly Sales - Total	2027	2091	2034	1816	2017	1827	2034	1995
System Forecast	1928.28	1936.38	2061.77	1835.85	1910.65	1901.48	1997.93	1902.03

Forecast Approval – Source Level Worksheet

The following is a description of the standard measures that are contained in the Source Level Worksheet. Additional registered measures may appear here if you made any such selections in the Forecast Approval wizard.

Weekly Sales (or the name of your sales measure)

This is the sales measure used to generate the forecast. This allows the user to compare Actuals to forecasted values.

System Forecast

The quantity that the system predicts will be required for the product, location, and calendar combination displayed.

Upper Confidence Limit (if selected in the wizard)

The upper bound of a probability distribution that is centered about the forecast value. It is calculated based on the Default Confidence % parameter value entered in the Forecast Administration Workbook, or on the Confidence Level value entered in the Confidence Worksheet of the Forecast Approval Workbook (if such a value has been entered to override the default).

Approval Worksheet

The Forecast Approval Worksheet allows you to specify the manual approval policy of forecasts by product and location. This worksheet can also be used to view the approval date of forecast values, as well as the name of the user who manually approved forecast values for a given product/location combination. The default Approval Method is set in Forecast Administration Workbook and for product/location combinations that vary from the default, the Forecast Maintenance Workbook can be used. This Workbook is used to approve forecasts that require manual approval or forecasts that have been modified using the Adjust Forecast measure in the Final Level Worksheet.

Location	Measure	Approval Comment	Approval Date	Approved	Approved By
Barcelona				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	
				<input type="checkbox"/>	

Approval Worksheet

The Approval Worksheet(s) contain the following standard measures:

Approved

A Boolean flag. Your main objective in the Forecast Approval Workbook is to fill the Approved Forecast field of the Final Level Worksheet with approved forecast values. The Approve (Boolean flag) allows the selection of the product/location combinations that can be approved.

- Do Not Approve (Not checked) – Does not approve the adjusted forecast value that is displayed in the Adjusted Forecast measure of the Final Level Worksheet.
- Approve (Checked) – Approves the forecast quantity entered in the Adjusted Forecast field of the Final Level Worksheet as it exists. The adjusted forecast quantity is approved and the forecast approval date measure is stamped with the current system date.

Approved By

A read-only field showing the name of the user who manually approved a forecast value for a given product/location combination. If the forecast value was automatically approved by Demand Forecasting, the string “System” is entered in this field. This measure is blank until a forecast quantity is approved.

Approval Comment

A comment field in which to enter notes regarding the approval of forecast values for specified product/location combinations.

Approved Date

A read-only field showing the date that the forecasted quantity was authorized. This information is necessary for Demand Forecasting to carry out any subsequent processes, such as replenishment procedures. Whenever a forecast is approved, this field is automatically updated with the current date.

Final or Source Parameters Worksheet

The Final or Source Parameters Worksheets allows you to view the alpha, level and trend parameters for each fitted time series. These parameters are only available if you checked “Generate Parameters” or “Generate Methods” in the Advanced Tab of the Forecast Administration Workbook.

Method Picked-Item/Chain/System	Alpha-Item/Chain/System	Level-Item/Chain/System	Trend-Item/Chain/System
SIMPLE	0.90	358.09	0.00
MMINTERS	0.06	1916.14	9.77
NO MODEL	0.00	0.00	0.00
NO MODEL	0.00	0.00	0.00
SIMPLE	0.05	544.15	0.00

Source Parameters Worksheet

The Final and Source Parameters Worksheet(s) contain the following standard measures:

Method Picked

The name of the method that was used to generate the forecast for the given product/location combination. This field is useful when combined methods are requested. Then this field displays the actual method the system picked from the combined methods. In case stand-alone methods are chosen then in general this field is the same as the method chosen in forecast administration or forecast maintenance. However if the requested method is unable to produce a good fit then the system will default to a “simpler method” and that method is displayed here.

System Alpha

The system-calculated alpha value (which is an internal optimization parameter which corresponds to the rate of decay of the weighting on the historical values) for the corresponding product/location combination if the chosen method is one of the following methods: Simple, Holt, Additive Winters, Multiplicative Winters, and Profile Based.

System Level

The system-calculated level (which is the constant baseline forecast) if the chosen method is one of the following methods: Simple, Holt, Additive Winters, Multiplicative Winters, and Seasonal Regression.

System Trend

The system-calculated trend (which is the rate of change of the baseline forecast with time) if the chosen method is one of the following methods: Holt, Additive Winters, Multiplicative Winters, and Seasonal Regression.

Confidence Worksheet

The Confidence Worksheet allows you to set the percent of confidence in the forecast. The value displayed in the Confidence Worksheet defaults to the percent value set in the Forecast Administration Workbook; the Confidence Worksheet allows you to override the default and change the confidence %. The confidence % is used to calculate the value of the upper confidence limit, which is useful in evaluating the accuracy of a forecast.

Statistical forecasting typically represents uncertainty in a forecast with confidence percentages and confidence limits. The forecast is centered between the upper and lower confidence limits; the confidence limits describe the spread of the probability distribution above and below the forecast. The confidence percentage is used to determine the values of the upper and lower confidence limits. A 90% confidence interval for a forecast contains the actual sales quantity 90% of the time. This means that 5% of the time actual sales will be above the upper confidence limit, and 5% of the time they will be below the lower confidence limit.

By accessing the Confidence Worksheet and changing the confidence percentage to a value either higher or lower than the system default confidence, the upper confidence limit changes. The table below illustrates how the upper confidence limits are changed in the Final Level and Source Level Worksheets when the confidence percentage is changed.

System Forecast	Confidence %	Upper Confidence Limit
3.96	10%	4.16
3.96	50%	5.06
3.96	90%	6.64

Note that the forecast does not change when the confidence limits change. Confidence limits are useful guides to analyze expected forecast performance and should not be used alone to measure the accuracy of a forecast.

Interactive Forecasting Workbook

The Interactive Forecasting Workbook allows you to make changes to the forecast parameters for a small group of SKU/Store combinations and see the results, without having to wait for the nightly batch run which generates forecasts for all the SKU/Store combinations set in the Forecast Maintenance Workbook.

In this workbook, you can edit various forecast parameters including sales history, forecast method, history start date, forecast length and forecast starting date. A new forecast is produced based on the changed parameters.

In addition to demand forecasts, the Interactive Forecasting Workbook can also generate fit in historical region, system picked model, confidence interval, and cumulative interval.

The forecasting methods used in the Interactive Forecasting Workbook include: Average, AutoES, Simple ES Trend ES, Seasonal ES, No Frst, Intermittent ES, Multiplicative Winters, Additive Winters, and Simple/Intermittent ES.

To access Interactive Forecasting , perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Interactive Forecasting and click OK.
- 3 Select the forecast level.
- 4 Select the forecast source measure.
- 5 Enter the default forecasting starting date.
- 6 Select the forecast output; that is, decide whether interval, cumulative interval, and system-picked model should be output.
- 7 Decide whether the forecast is produced at the time of workbook creation. If 'No' is selected a forecast will not be generated when the workbook build processes is completed. If 'Yes' is selected, the system will use the selections chosen during this wizard process in addition to the following parameters as the basis for the forecast:
 - Forecast Length (set in Forecast Administration Workbook)
 - History Start Date (either from Forecast Administration Workbook or the Forecast Maintenance Workbook if an override was set for an item/location intersection being evaluated)
 - Forecast Method (either from Forecast Administration Workbook or the Forecast Maintenance Workbook if an override was set for an item/location intersection being evaluated)
- 8 Range the Time dimension.
- 9 Range the Product dimension.
- 10 Range the Store dimension.
- 11 Select extra measures to be included in the forecasting workbook.

Once the wizard is completed, the Interactive Forecasting Workbook is displayed. The workbook contains two worksheets:

- Forecasting Parameter Worksheet
- Interactive Forecasting Worksheet

Forecasting Parameter Worksheet

The Forecasting Parameter Worksheet is based on the intersection of the Product and Location dimensions.

	Forecast Horizon	Forecast Method	Forecast Starting Date	History Start Date	System Chosen Method
Ban Roll On - Unscented	13.00	No Frst	10/1/1999	1/5/1997	No Frst
Ban Roll On - scented	13.00	No Frst	10/1/1999	1/5/1997	No Frst
Crest Tooth Paste (150ml)	13.00	AutoES	3/1/2002	1/1/1998	Simple
Crest Tooth Paste (75ml)	13.00	AutoES	3/1/2002	1/1/1998	Simple

Forecasting Parameter Worksheet

The Forecasting Parameter Worksheet contains the following measures. Of these measures, all measures but System picked model are editable. Any changes to the editable parameters, that is, Sales History, Forecast Method, History Start Date, Forecast Start Date, and Forecast Horizon, trigger a forecast of the time series.

Forecast Method

A drop-down list from which you can select the default forecast generation method for items forecasted at the lower level. Typically, a lower-level forecast is run in tandem with an upper-level forecast, and the information gained by generating the former is used to determine the manner in which forecast data is spread from the higher level. Valid options depend on your system setup. A summary of methods is provided in the overview of Chapter 2, and Chapter 6 covers each method in detail.

History Start Date

The starting date for historical sales data. For example, if your system start date is January 1, 2003, but you only want to use historical sales data from the beginning of 1999, then you need to set your History Start Date to January 1, 1999. Only history after this date is used for generating the forecast. The default is the system start date unless otherwise specified. If sales data is collected weekly or monthly, Retek Demand Forecasting generates forecasts only using data from sales periods after the one containing the history start date.

It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 1999 and an item/location does not have sales history until February 1, 1999, the system will consider the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

Forecast Start Date

The starting date of the forecast. The default is the system date unless otherwise specified.

Forecast Length

The number of time periods (typically weeks or months) that the system forecasts beyond the horizon start date. Company preferences determine the length of these time periods; this information is built in Retek Demand Forecasting at the time of installation.

System Picked Model

The name of the method that was used to generate the forecast for the given product/location combination. This field is useful when combined methods are requested. Then this field displays the actual method the system picked from the combined methods. In case stand-alone methods are chosen then in general this field is the same as the method chosen in forecast administration or forecast maintenance. However if the requested method is unable to produce a good fit then the system will default to a “simpler method” and that method is displayed here.

Interactive Forecasting Worksheet

The Interactive Forecasting Worksheet is based on the intersection of the Product, Location, and Time dimensions.

Item/Chain/Week - Source Interactive Forecasting Worksheet		Calendar					
Location							
Bricks & Mortar							
		03/01/02	03/08/02	03/15/02	03/22/02	03/29/02	04/05/02
Ban Roll On - Unscented	Sales History	541.00	578.00	547.00	576.00	564.00	559.00
	System Forecast Item/Chain/Week - Source	0.00	0.00	0.00	0.00	0.00	0.00
	System Cumulative Interval Item/Chain/Week - Source	0.00	0.00	0.00	0.00	0.00	0.00
	System Interval Item/Chain/Week - Source	0.00	0.00	0.00	0.00	0.00	0.00
Ban Roll On - scented	Sales History	1720.00	1725.00	1690.00	1648.00	1745.00	1607.00
	System Forecast Item/Chain/Week - Source	0.00	0.00	0.00	0.00	0.00	0.00
	System Cumulative Interval Item/Chain/Week - Source	0.00	0.00	0.00	0.00	0.00	0.00
	System Interval Item/Chain/Week - Source	0.00	0.00	0.00	0.00	0.00	0.00
Crest Tooth Paste (150ml)	Sales History	692.00	737.00	633.00	653.00	721.00	679.00
	System Forecast Item/Chain/Week - Source	696.35	696.35	696.35	696.35	696.35	696.35
	System Cumulative Interval Item/Chain/Week - Source	208.90	361.83	511.71	660.61	809.08	957.32
	System Interval Item/Chain/Week - Source	208.90	295.44	361.83	417.81	467.12	511.71
Crest Tooth Paste (75ml)	Sales History	690.00	691.00	696.00	705.00	695.00	703.00
	System Forecast Item/Chain/Week - Source	694.75	694.75	694.75	694.75	694.75	694.75
	System Cumulative Interval Item/Chain/Week - Source	208.43	361.00	510.54	659.10	807.23	955.13
	System Interval Item/Chain/Week - Source	208.43	294.76	361.00	416.85	466.06	510.54
Product	Measure						

Interactive Forecasting Worksheet

The Interactive Forecasting Worksheet contains the following measures. Of these measures, only Sales History is editable; all other measures are read-only.

Sales History

Your historical sales data. This field is editable so you can change out-of-character sales if needed. For example, if your battery sales went up during a major power outage, you can lower the sales data back to a more normal level so the unusually high sales will adversely affect your forecast.

System Cumulative Interval

The cumulative standard deviation of the fit error over the history. Cumulative Intervals are a running total of Intervals. The cumulative interval is necessary if forecast information is to be exported to the Retek Replenishment application. See Forecast Approval Workbook for more details on the calculation of Cumulative Intervals when the user adjusts forecasts.

System Forecast

The quantity that the system predicts will be required for the product, location, and calendar combination displayed.

System Interval

The system generated confidence interval, which requires the standard deviation of the fit error over the forecast history. For many forecasting methods they are calculated as standard deviation but for Simple, Holt, and Winters the calculation is more complex. (See Chapter 6 for more details.) Intervals are not exported. Also, Intervals are capped.

Forecast Scorecard Workbook

This section describes the purpose and content of the Forecast Scorecard and the steps required in order to create and access one. This section also discusses the two types of worksheets contained in the Forecast Scorecard Workbook, as well as the definitions of parameters that exist in each. Evaluating forecast accuracy through the use of error statistics is discussed, as is the process of comparing historical forecasts to actual sales data.

The purpose of the Forecast Scorecard is to monitor the accuracy of both system-generated forecasts and approved final forecasts. Once a forecast has been generated and actual point-of-sale data is received for the forecasted period, statistical information can be reviewed to help you analyze the accuracy of forecasting models and methods.

Forecast Scorecard wizard

The Forecast Scorecard template consists of a wizard and two worksheets.

Based on your selections in the wizard, the Forecast Scorecard provides statistical information and comparison data that allow you to monitor the accuracy of system-generated forecasts and final approved forecasts.

- The Final Error Measure Worksheet displays statistical information, such as mean error and root mean squared error, that reflects the accuracy of the forecast.
- The Final Actuals vs. Forecasts Worksheet displays forecast results and actual point-of-sale values for specified product/location/calendar combinations.

To access Forecast Scorecard, select Open from the File menu to bypass the Forecast Scorecard Wizard and open an existing Forecast Scorecard, or perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Scorecard and click OK.
- 3 The Forecast Scorecard wizard opens and prompts you to select the final forecast level that you want to evaluate. Make a selection and click Next.
- 4 Select the generation date of the forecast you want to evaluate. Click the radio button corresponding to your selection and click Next.
- 5 Select the forecasts you want to evaluate (such as System Forecast, Final Forecast, etc.). Click Next.
- 6 Check the boxes corresponding to the error statistics that you want calculated for the forecast data. Click Next.
- 7 Select the specific products whose forecasts and actual sales data you want to view. Highlight the products on the Available Products side of the screen and click the right arrow button to move them to the Selected Products side. You can access a quick menu of additional functions, such as changing the product rollup, by clicking the right mouse button. Click Next.
- 8 Perform the same operation for locations that you want to view. Highlight the locations on the Available Locations side of the screen and click the right arrow button to move them to the Selected Locations side. Click Next.
- 9 Select any additional measures (that is, measures not standard in the Forecast Scorecard Workbook) that you would like included. Click the Finish button to display the workbook.

Final Error Measure Worksheet

The Final Error Measure Worksheet displays statistical information that reflects the accuracy of the forecast. You may need to evaluate a variety of such statistics to verify and compare forecast accuracy.

Given the situation, different levels of forecast accuracy can be useful. For example, in a situation with noisy data and no forecast previously available, then a 200% error can be considered excellent and useful. In another situation with smooth data, if an old method of forecasting provided forecasts with a 10% error, then a new method of forecasting with a 20% error would not be considered useful.

Different levels of forecasting accuracy are obtained at different levels of product aggregation. SKU-level forecasts with a 200% error can roll up into division-level forecasts with a 10% error. Therefore, the error measures are most useful when comparing different methods of forecasting or when looking at a particular model's accuracy over time. Accurate forecasts should lead to a reduction in inventory holding costs, or to an increase in customer service levels; forecast errors should be converted into dollars saved.

Note: All error statistics are calculated over the first period in the forecast to the last period in the forecast for which forecast data is available.

	10000115 Crest Tooth Pa	10000117 Crest Tooth Pa	10000118 Glue Gun Stick
Averaged Actual Over Horizon	519.00	524.00	233.00
Averaged System Forecast Over Horizon	517.19	546.00	247.94
System Forecast Mean Absolute Error	0.36	2.77	2.88
System Forecast Mean Absolute Percentage Error	0.00	0.00	0.00
System Forecast Mean Error	0.14	-1.69	-1.15
System Forecast Percentage Absolute Error	0.00	0.00	0.00
System Forecast Root Mean Squared Error	0.47	3.06	3.74

Final Error Measure Worksheet

The following is a description of the measures which can be included in the Final Error Measure Worksheet: The error statistics will be shown for each forecast you chose to include in the workbook (such as System Forecast, Final Forecast, etc.).

Averaged System Forecast Over Horizon

The sum of all forecasted sales from the first period in the forecast to the last period in the forecast for which forecast data is available.

Averaged Actual Over Horizon

The sum of all actual sales from the first period in the forecast to the last period in the forecast for which point-of-sale data is available.

Mean Error

The error of a forecast observation is the difference between the forecasted value and the actual POS value. The Mean Error statistic is a measure of the average error over time, calculated by summing the errors for all observations, then by dividing by the number of observations to obtain the average. It measures forecast accuracy by calculating the error in units. Because a positive error in one period can cancel out a negative error in another period, this measure is useful when you are interested in how well the forecast predicts over the forecast horizon rather than on a period-to-period basis. Mean error is, however, useful as a measure of forecast bias; a negative mean error suggests that, overall, the forecasting model overstates the forecast, while a positive mean error indicates forecasts that are generally too low.

$$\frac{1}{\# \text{ of observations}} \sum_{i=1}^{\# \text{ of observations}} (Actual_i - Forecast_i)$$

Mean Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecasted value and the actual POS value. The Mean Absolute Error statistic is a measure of the average absolute error, calculated by summing the absolute errors for all observations, then by dividing by the number of observations to obtain the average. Mean Absolute Error gives you a better indication of how the forecast performed period by period, because the absolute value function ensures that negative errors in one period are not canceled out by positive errors in another. Mean Absolute Error is most useful for comparing two forecast methods for the same series.

$$\frac{1}{\# \text{ of observations}} \sum_{i=1}^{\# \text{ of observations}} |Actual_i - Forecast_i|$$

Root Mean Squared Error

The square root of the Mean Squared Error. The Root Mean Squared Error is one of the most commonly used measures of forecast accuracy because of its similarity to the basic statistical concept of a standard deviation. It evaluates the magnitude of errors in a forecast on a period-by-period basis, and is best used to compare alternative forecasting models for a given series.

$$\sqrt{\frac{\sum_{i=1}^{\# \text{ of observations}} (Actual_i - Forecast_i)^2}{\# \text{ of observations}}}$$

Mean Absolute Percentage Error

The percentage error of a forecast observation is the difference between the actual POS value and the forecasted value, divided by the actual POS value. The result of this calculation expresses the forecast error as a percentage of the actual value. The Mean Absolute Percentage Error statistic measures forecast accuracy by taking the average of the sum of the absolute values of the percentage error calculations across all observations. This method is useful when comparing the accuracy of forecasts for different volume products (it normalizes error by volume).

$$\frac{100}{\# \text{ of observations}} \sum_{i=1}^{\# \text{ of observations}} \frac{|Actual_i - Forecast_i|}{Actual_i}$$

Percentage Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecasted value and the actual POS value. The Percentage Absolute Error statistic measures forecast accuracy by calculating the total absolute error as a percentage of the total actual POS. It is calculated by summing the absolute errors for all observations, dividing this value by the absolute value of the sum of all actuals, and dividing the result by the number of observations in the series. Finally, multiply the total by 100 to obtain a percentage result.

$$\frac{100}{\# \text{ of observations}} \frac{\sum_{i=1}^{\# \text{ of observations}} |Actual_i - Forecast_i|}{\sum_{i=1}^{\# \text{ of observations}} |Actual_i|}$$

Final Actuals vs. Forecasts Worksheet

This worksheet displays forecast results and actual point-of-sale values for each product, location, and time period specified in the Forecast Scorecard Wizard. This worksheet allows you to compare the results of both system-generated forecasts and final approved forecasts to actual sales quantities.

	04/10/01	04/17/01	04/24/01	05/01/01	05/08/01	05/15/01	05/22/01	05/29/01	06/05/01
Actual	811.00	794.00	813.00	808.00	817.00	813.00	800.00	803.00	##
System Forecast SKU/Store/Week - Final	802.64	802.75	808.51	806.48	809.22	806.98	804.49	806.58	##
actual-system forecast on SKU/Store/Week - Final	8.36	-8.75	4.49	1.52	7.78	6.02	-4.49	-3.58	##

Final Actuals vs. Forecasts Worksheet

The Final Actuals vs. Forecasts Worksheet may contain the following measures:

System on [Forecast Date]

Displays the system-generated forecast quantity for the product, location, and calendar combination displayed. Note that the [Forecast Date] portion of this measure's label will be replaced in your worksheet with the date of the forecast you are evaluating.

Final on [Forecast Date]

Displays the final approved forecast quantity for the product, location, and calendar combination displayed. Note that the [Forecast Date] portion of this measure's label will be replaced in your worksheet with the date of the forecast you are evaluating.

Actual

Displays the actual point-of-sale quantities for the product, location, and calendar combinations displayed.

(Actual) – (Final on [Forecast Date])

Displays the difference between the actual sales quantities and the final approved forecast values for each product, location, and calendar combination displayed. Negative values indicate that the final forecast exceeded actual sales.

(Actual) – (System on [Forecast Date])

Displays the difference between the actual sales quantities and the system-generated forecast values for each product, location, and calendar combination displayed. Negative values indicate that the system forecast exceeded actual sales.

Export Forecasts

This section outlines the purpose of the Export Forecasts Wizard, and the steps required to export approved forecasts from Retek Demand Forecasting to a text file for subsequent processing.

It is often necessary to make approved forecasting results (both the forecasts and the standard deviations) available in the form of a text file. Approved forecasts are generally set up to be exported automatically from Retek Demand Forecasting during the nightly batch process. The Export Forecasts Wizard allows you to manually export approved forecasts from Retek Demand Forecasting to a text file on the server, so that they can be used for subsequent processes such as replenishment programs. Only forecasts that have been approved can be exported from Retek Demand Forecasting.

Setup for the Export Forecasts Wizard must be performed in the Forecast Export Administration Workbook. Approved forecasts can then be exported interactively by accessing the Export Forecasts Wizard.

Note that the default file name for the exported forecasts is frst.out. If you do not set a different name for the output file in the Forecast Export Administration Workbook, then Retek Demand Forecasting will replace the frst.out file with a new version each time that you export forecasts.

Export Forecasts Wizard

The Export Forecasts Wizard is used to manually export a forecast to another program. To perform this process, complete the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Export Forecasts and click OK.
- 3 The Export Forecasts wizard opens and prompts you to select the job to be exported. Make a selection and click Next.
- 4 Select the threshold date of the forecasts to be exported. All forecasts approved on or after this date will be exported to the text file. Click Next.
- 5 Select Yes to export all product/location combinations, or No to choose a subset of combinations from the standard product and location selection wizards. Selection wizards will appear next only if No is selected here; otherwise, all product/location combinations will be included in the export. Click Next.
- 6 Select whether post-processing of the output file will be performed now. The procedure that is called for post-processing is specified in the Forecast Export Administration Workbook. This procedure is generally used to format the text file and/or to transfer the file to a specific system directory. Make the appropriate selection and click Next.
- 7 Select Yes to begin exporting, or No to cancel the request without exporting data. Click Next.
- 8 The final wizard screen displays the results of the export request. Click the Finish button to close the window.

Forecast Export Administration Workbook

This section covers the purpose and content of the Export Administration Workbook. It discusses the steps required to create and access the Forecast Export Administration Workbook, and the definitions of parameters that exist in the Export Parameters Worksheet.

It is often necessary to make approved forecasting results (both the forecasts and the standard deviations) available in the form of a text file. The Forecast Export Administration Workbook allows you to set and alter the values of certain exporting parameters before approved forecasts are converted to text files and exported. When forecasts are exported, they are saved in a directory on the server for use in subsequent processes, such as replenishment programs. Only forecasts that have been approved can be exported from Retek Demand Forecasting.

Approved forecasts can be set up to be exported as part of a batch process, or they can be exported interactively by accessing the Export Forecasts Workbook in Retek Demand Forecasting. To change the format of the exported text file, you must first access and commit the Export Parameters Worksheet of the Forecast Export Administration Workbook.

Access the Forecast Export Administration Workbook

To access Forecast Export Administration, select Open from the File menu to open an existing Forecast Export Administration Workbook, or perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Export Administration and click OK.
- 3 The Forecast Export Administration Workbook opens so you can make the necessary changes to the export parameters. For the changes to take effect during the next forecast export, you must commit your changes to the master database.

Export Parameters Worksheet

The Export Parameters Worksheet is the only worksheet in the Forecast Export Administration Workbook. It allows you to set and alter the values of various forecast export parameters.

	1 - UCSH/prbr/Week (26 week horizon)	2 - Category/prbr/Week (26 week horizon)
Label	UCSH/prbr/Week (26 week horizon)	Category/prbr/Week (26 week horizon)
Forecast Level	UCSh/PrBr/Week	UCSh/PrBr/Week
Forecast Intersection	UCSh/PrBr/Week	ProdCat/PrBr/Week
Export Start Date [Sys Date]	6/29/00	6/29/00
Number of Periods	26	26
Approval Threshold Date [Sys Date]	6/28/00	6/28/00
Outfile Name ["frst"]	final	catlevl
File Format Type	Entire Horizon per Record	Entire Horizon per Record
Number of Deviations [0]	13	0
Date Field Format	mm/dd/yyyy	mm/dd/yyyy
Numeric Field Width [9]		
Include Decimal Point [Yes]	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Number of Decimal Places [2]		
Zero Pad [No]	<input type="checkbox"/>	<input type="checkbox"/>
Post-Processing Procedure Name		

Export Parameters Worksheet

The following is a description of the parameters that exist in the Export Parameters Worksheet. The values in brackets after each measure name refer to the default value for that parameter if it is left unchanged. The default values for your system may differ from those shown above, as they are specified by the customer during implementation.

Label

The name that is associated with a specific export job. When you start an export job interactively by accessing the Forecast Export Wizard, this name is displayed as an option wizard screen that prompts you to select a job to export.

Forecast Level

A drop-down list from which you can select the forecast level to be exported. Since only approved forecasts can be exported, this list only provides final level forecasts as valid options.

Forecast Intersection

A drop-down list from which you can select the dimensional intersection of the product, location, and calendar hierarchies at which forecasting results are to be exported. It is possible to approve forecasts at the SKU/Store/Week level, for example, but to export them at the Style/Chain/Week level. In this case, the Forecast Level parameter would be set to the SKU/Store/Week intersection, and the Forecast Intersection parameter would be Style/Chain/Week.

Export Start Date [Sys Date]

The start date of the time horizon of the forecast that is being exported. The default value for this parameter is today's date.

Number of Periods

The number of periods (starting with the Export Start Date) in the time horizon for which a forecasted value should be exported. For example, even if forecasts were approved at the week level, they can be exported at the month level; in this case, Number of Periods should be the number of months to export.

Approval Threshold Date [Sys Date]

Only series with at least one forecast value (in the specified time horizon) approved on or after this date are exported. The approval threshold can be set to any date in the domain's time hierarchy from the first day (to export all approved series) up to and including yesterday or today (to export only recent approvals). The default value for this parameter is the system date.

Note: If nothing was approved on or after the Approval Threshold Date, then no file will be created.

Outfile Name ["frcst"]

The name (without the extension) of the text file to be created. The text file is created in the domain's output directory and has a .out extension. Any existing file of the same name is overwritten when the export is performed. The default value for this parameter is frcst (output/frcst.out) if left blank.

File Format Type

A drop-down list from which you can select the basic format type of the text file. There are three types available:

- Entire Horizon per Record – Each record contains the location position name, the product position name, a forecast value for each period in the specified time horizon, and possibly a standard deviation value for each period in the specified time horizon. Records look like the following:

```
loc1 prod1 1 2 5 4 3 0 2
```
- Single Value per Record – Each record contains the location position name, the product position name, the time period, the forecast value, and possibly the standard deviation value. Records look like the following:

```
loc1 prod1 time1 1
loc1 prod1 time2 2
loc1 prod1 time5
loc1 prod1 time4
loc1 prod1 time3
loc1 prod1 time0
loc1 prod1 time2
```
- Horizon Total per Record – Each record contains the location position name, the product position name, a forecast total for the specified time horizon, and possibly a standard deviation total. Records look like the following:

```
loc1 prod1 17
```

Note: Regardless of the format type selected, field separators are not displayed in the final text file output. The examples above show separators for readability only. Also, any blank data cells are exported to the text file as zeros.

Number of Deviations [0]

The number of forecast values (starting with the Export Start Date) for which a corresponding standard deviation value should be exported. For Entire Record per Horizon and Horizon Total per Record file format types, the Number of Deviations value can be any number between zero and Number of Periods inclusive. For the Single Value per Record file format type, the Number of Deviations value can only be zero or Number of Periods. In the exported data file, this field (if present) is displayed after (to the right of) the forecast value field.

Note: If you are exporting at an aggregate level or exporting using the Horizon Total per Record format, the deviations will be summed up (as are the forecasts). This is probably not the result you want, so it is recommended that you omit standard deviations in these cases.

Date Field Format

The date format type (necessary only if the File Format Type is set to Single Value per Record). The choices are mmddyyyy, mm/dd/yyyy, mm-dd-yyyy, or yyyyymmdd. The latter is the default value.

Numeric Field Width [9]

The number of columns for each numeric field (forecasts and standard deviations). The default value is 9 unless otherwise specified.

Include Decimal Point [Yes]

A check box in which you can specify whether to include a decimal point in the numeric data fields of your text file. The precision of the decimal place can be controlled by the Number of Decimal Places parameter below. A checkmark in this field causes a decimal point to be included; no checkmark eliminates the decimal. The default value is Yes.

Number of Decimal Places [2]

The number of places to keep after the decimal point (see Include Decimal Point above). The default value for this parameter is 2.

Zero Pad [No]

A check box in which you can specify whether numeric fields should be padded with zeros. A checkmark in the check box pads the field with zeros; no checkmark causes the field to be padded with blanks. The default value is No.

Post-Processing Procedure Name

The name of the procedure (if any) to be called upon completion of the export. This procedure is typically used to further format the text file and/or to transfer the file to another directory in the system.

Procedures in Statistical Forecasting

The following is a list of standard procedures performed in the context of statistical forecasting. These procedures utilize the workbook templates and wizards associated with Retek Demand Forecasting's Predict module.

- Set or change default parameters for a forecast
- Adjust forecast parameters for specific products or locations
- Manually request an unscheduled forecast
- Review, modify, and approve forecast values
- Manually export a forecast to another program
- Set or change forecast export parameters
- Monitor the accuracy of Retek Demand Forecasting forecasts
- Delete an unnecessary forecast

Set or change default parameters for a forecast

Follow this procedure to review or modify the default values for parameters affecting forecast generation. Such values include the forecast's data source measure, the default generation method, the default training window, the default approval policy, and the horizon start date, among others.

- 1 To open an existing Forecast Administration Workbook, select Open from the File menu, select a workbook, click OK, and go to step 5. Or to open a new workbook, select New from the File menu.
- 2 Click the Predict tab.
- 3 Select Forecast Administration and click OK.
- 4 Click a radio button to select the final forecast level whose default parameters you want to modify. Click Finish.
- 5 On the Final Level Worksheet, review each cell and set its contents or value as needed.
- 6 On the Source Level Worksheet, review each cell and set its contents or value as needed.
- 7 Commit your changes to the master database by selecting Commit Now from the File menu. If desired, you may also save your workbook by selecting Save from the File menu.
- 8 To close the workbook, select Close from the File menu.

Adjust forecast parameters for specific products or locations

Use this procedure to set forecast parameters for individual products and locations to values other than the default values. You may wish to change such parameters as the forecast method, the training window, the approval method, or the forecast generation level, among others.

To access Forecast Maintenance, select Open from the File menu to bypass the Forecast Maintenance Wizard and open an existing workbook, or perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Maintenance and click OK.
- 3 The Forecast Maintenance Wizard opens and prompts you to select the level of the final forecast. The final forecast level is a level at which approvals and data exports can be performed. Depending on your organization's setup, you may be offered a choice of several final forecast levels. Make the appropriate selection and click Next.
- 4 Select the products whose parameters you want to update. Highlight the appropriate products on the Available Products side of the screen and click the right arrow button to move them to the Selected Products side. Click Next.
- 5 Now perform the same operation for locations in the location hierarchy. Highlight the appropriate locations on the Available Locations side of the screen and click the right arrow button to move them to the Selected Locations side. Click Next.
- 6 An additional wizard screen may prompt you to select any additional measures (that is, measures not standard in the Forecast Maintenance Workbook) that you would like included. Make the appropriate selections (if any) and click the Finish button to display the workbook.
- 7 On the Final Level Worksheet, review each cell and change its value or setting as needed. For a description of the measures contained on this worksheet, see "Forecast Maintenance - Final Worksheet" on page 27.
- 8 On the source level worksheets, review the Forecast Method and Training Window measures for each product/location combination and change any values as needed.
- 9 Commit your changes to the master database by selecting Commit Now from the File menu. If desired, you may also save your workbook by selecting Save from the File menu.
- 10 To close the workbook, select Close from the File menu.

Manually request an unscheduled forecast

The Run Batch Forecast wizard allows you to manually start the forecast generation process. To kick off a forecast run at a time other than the regular overnight batch job, you must do the following:

- 1 Create or access a Forecast Administration Workbook. In the Horizon Start Date measure, enter the starting date of the forecast horizon.
- 2 Access the Run Batch Forecast Wizard to begin the batch forecast process for all product/location combinations set to some method of forecast generation other than No Forecast.

Note: Before performing this procedure, make sure there are no other users logged onto Retek Demand Forecasting. No other user can access Retek Demand Forecasting until the forecast generation process is finished.

To access and complete the Run Batch Forecast wizard, perform the following steps.

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Run Batch Forecast and click OK.
- 3 The Run Batch Forecast Wizard opens and prompts you to select the level of the final (approval level) forecast. Click Next.
- 4 Select which batch type to run.
 - **Scheduled Runs and Overrides.** Generates all regularly scheduled batch forecasts set to be generated on today's date. When a scheduled batch forecast is run, the Next Run Date measure in the Forecast Administration Workbook is updated with an appropriate value, based on the established Review Time. After a forecast is generated on its scheduled date and the Next Run Date measure is updated, the forecast won't be generated again during its regularly scheduled batch process time. If you select this option (Scheduled Runs and Overrides) and no batch runs are scheduled for this date, the system reads the Horizon Start Date field of the Forecast Administration Workbook. If this field contains a value, the Next Run Date value will be overridden and a forecast will be generated now (based on the established Horizon Start Date value). If no value is set for the Horizon Start Date (that is, Horizon Start Date = NA), then a forecast will be generated using today's date as the Horizon Start Date, but only if the Next Run Date is less than or equal to today. If the Next Run Date is set to some future date, then no forecast will be run.

Note: After a forecast is generated, the Horizon Start Date field is cleared. This ensures that the same forecast will not be generated again on the same date.

- **Overrides Only.** Causes the system to consider only the Horizon Start Date field of the Forecast Administration Workbook. If this field contains a value, the Next Run Date value will be overridden and the forecast (based on the horizon start date) will be generated now. The Next Run Date value is not updated; as a result, any regularly scheduled run for this date will be performed during the overnight batch process. If no value is set for the Horizon Start Date (that is, Horizon Start Date = NA), then no forecast will be generated.

- 5 Click Next to begin the forecast generation process.

Review, modify, and approve forecast values

After the forecast is generated, the next steps in the forecasting process are analysis and approval. Approval of forecasts is required before the forecasted data can be exported to other processes.

To access Forecast Approval, select Open from the File menu to bypass the Forecast Approval Wizard and open an existing Forecast Approval Workbook, or perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Approval and click OK.
- 3 The Forecast Approval wizard opens and prompts you to select the final level at which to approve forecast values. Make your selection and click Next.
- 4 Select the forecast level(s) to include in the Forecast Approval Workbook. The final forecast level selected on the previous screen will appear in your workbook, plus any source-level forecasts chosen here. If the forecast was not run at a source level as well as a final level, then this screen will not appear. Select as many forecast levels as necessary for comparison, then click Next.
- 5 Now select the generation date of the forecast you wish to approve. Click 'Last Chronological Forecast' to build a workbook containing the most recent forecast values, or click 'List of all generated forecast dates' to select from a list of all previously generated forecasts. Click Next.
- 6 Select the specific products whose forecasts you want to view. Highlight the products on the Available Products side of the screen and click the right arrow button to move them to the Selected Products side. Click Next.
- 7 Perform the same operation for locations that you want to view. Highlight the locations on the Available Locations side of the screen and click the right arrow button to move them to the Selected Locations side. Click Next.

- 8 The next wizard depends on your selection in step 5. If you selected a forecast from the 'List of all generated forecast dates', then select the time periods for which you want to view the forecast data. Highlight the time periods on the Available Dates side of the screen and click on the right arrow button to move them to the Selected Dates side. If you selected, 'Last Chronological Forecast', then enter the number of periods you want to view prior to the forecast. Click Next.
- 9 Select whether you want to view data for any time periods before the forecast horizon. If you select Yes on this wizard screen, you will be prompted to select the specific time periods you want to include. If you select No, a calendar selection wizard will not be displayed. Select the appropriate radio button and click Next.
- 10 From the list provided, check any additional forecasting measures specific to this forecast that you would like to view. Click Next.
- 11 Next, highlight any additional registered measures you would like to view in your workbook. The measures available on this wizard screen are not specific to this forecast, and can be chosen for any new workbooks.
- 12 Click Finish to build and open your workbook.
- 13 On the Final Forecast Worksheet, review all forecast values.
- 14 In the Final Forecast measure, overwrite any forecast values that you wish to change.
- 15 Display the Forecast Approval Worksheet.
- 16 The Approval Method measure of the Forecast Approval Worksheet defines the method that Retek Demand Forecasting uses to automatically assign and authorize final forecast quantities for individual product/location combinations. For each item, specify the manual approval method by clicking in each cell and selecting a method from the list that displays.
- 17 Commit your changes to the master database by selecting Commit Now from the File menu. If desired, you may also save your workbook by selecting Save from the File menu.
- 18 To close the workbook, select Close from the File menu.

Manually export a forecast to another program

The Export Forecasts Wizard is used to manually export a forecast to another program. To perform this process, complete the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Export Forecasts and click OK.
- 3 The Export Forecasts wizard opens and prompts you to select the job to be exported. Make a selection and click Next.
- 4 Select the threshold date of the forecasts to be exported. All forecasts approved on or after this date will be exported to the text file. Click Next.
- 5 Select Yes to export all product/location combinations, or No to choose a subset of combinations from the standard product and location selection wizards. Selection wizards will appear next only if No is selected here; otherwise, all product/location combinations will be included in the export. Click Next.
- 6 Select whether post-processing of the output file will be performed now. The procedure that is called for post-processing is specified in the Forecast Export Administration Workbook. This procedure is generally used to format the text file and/or to transfer the file to a specific system directory. Make the appropriate selection and click Next.
- 7 Select Yes to begin exporting, or No to cancel the request without exporting data. Click Next.
- 8 The final wizard screen displays the results of the export request. Click the Finish button to close the window.

Set or change forecast export parameters

To change the format of the exported text file, you must first access and commit the Export Parameters Worksheet of the Forecast Export Administration Workbook.

To access Forecast Export Administration, select Open from the File menu to open an existing Forecast Export Administration Workbook, or perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Export Administration and click OK.
- 3 The Forecast Export Administration Workbook opens so you can make the necessary changes to the export parameters. For a description of these measures, see “Forecast Export Administration Workbook” on page 58.
- 4 For the changes to take effect during the next forecast export, you must commit your changes to the master database. Select Commit Now from the File menu.
- 5 Select Close from the File menu to close the workbook.

Monitor the accuracy of Retek Demand Forecasting forecasts

Forecast Scorecard allows you to monitor the accuracy of both system-generated forecasts and approved final forecasts. To access Forecast Scorecard, select Open from the File menu to bypass the Forecast Scorecard Wizard and open an existing Forecast Scorecard, or perform the following steps:

- 1 Select New from the File menu.
- 2 Select the Predict tab to display a list of workbook templates for statistical forecasting. Highlight Forecast Scorecard and click OK.
- 3 The Forecast Scorecard wizard opens and prompts you to select the final forecast level that you want to evaluate. Make a selection and click Next.
- 4 Select the generation date of the forecast you want to evaluate. Click the radio button corresponding to your selection and click Next.
- 5 Select the forecasts you want to evaluate (such as System Forecast, Final Forecast, etc.). Click Next.
- 6 Check the boxes corresponding to the error statistics that you want calculated for the forecast data. Click Next.

- 7 Select the specific products whose forecasts and actual sales data you want to view. Highlight the products on the Available Products side of the screen and click the right arrow button to move them to the Selected Products side. Click Next.
- 8 Perform the same operation for locations that you want to view. Highlight the locations on the Available Locations side of the screen and click the right arrow button to move them to the Selected Locations side. Click Next.
- 9 Select any additional measures (that is, measures not standard in the Forecast Scorecard Workbook) that you would like included. Click the Finish button to display the workbook.
- 10 The Forecast Scorecard Actuals vs. Forecasts Worksheet displays forecast results and actual POS values for each product/location/time period selected in the wizard. Compare the results of both system-generated forecasts and final approved forecasts to actual sales quantities.
- 11 Use the Final Error Measure Worksheet to review statistical information that reflects the accuracy of the forecast. For detailed information about the measures contained in this worksheet, see “Forecast Scorecard Final Error Measure Worksheet” on page 53.
- 12 To save your changes, select Save from the File menu.
- 13 To close the workbook, select Close from the File menu.

Delete an unnecessary forecast

- 1 From the File menu, select New.
- 2 Select the Predict tab.
- 3 Select Delete Forecasts. Click OK.
- 4 Select the generation date of the forecast to delete. Click Next.
- 5 Select the Yes button to confirm the forecast deletion.

Note: Deleted forecasts are permanently removed from Retek Demand Forecasting.

- 6 Click Finish to process the request.

Chapter 3 – Promote (Promotional Forecasting)

Overview

Promote is an optional add-on automated predictive solution that allows you to incorporate the effects of promotional and causal events, such as radio advertisements and holiday occurrences, into your time series forecasts. The promotional forecasting process uses both past sales data and promotional information to forecast future demand.

This chapter provides an introduction to promotional forecasting and explains how it differs from the traditional statistical forecasting methodology. It discusses the advantages and limitations of both statistical and promotional forecasting models, and outlines the use of Retek's promotional method of forecasting demand. It describes terminology used in the context of promotional forecasting, and concludes with detailed descriptions of the workbook templates contained in the Promote module.

What is Promotional Forecasting?

Traditional statistical forecasting methods provide significant benefits to the process of forecasting consumer demand, because they are good at predicting level, trend, and seasonality based on sales history. The limitation of traditional statistical methods is that they forecast with less accuracy when there are special events that cause significant deviations in selling patterns.

For example, the Easter holiday, for which companies often run promotions, occurs on a different date each year. Traditional statistical forecast methods can identify seasonality in sales history, but this seasonality is based on periodic similarities in the sales pattern. Since Easter occurs on different dates from year to year (that is, its period of recurrence is not regular), manual intervention is required to predict change in demand using the traditional statistical forecasting method. Events like this are called promotion events. Promotion events, such as advertisements, irregularly occurring holidays, competitor information, free gift with purchase offers, etc., are events that drive businesses from the normal selling cycle. The goal of a promotional forecasting system is to improve time series forecasting by:

- Providing the forecasting system with visibility as to when certain promotion events were active in the past (for example, identifying which weeks of a given year were affected by an Easter promotion).
- Automatically determining the statistical effect, if any, of these events.
- Incorporating significant effects into the future forecasts for time periods also associated with the observed promotion event.

The Promote module combines the automation of statistical time series forecasting with the improved accuracy found in customized causal forecasting. Promote uses both past sales data and promotional information (for example, advertisements, holidays) to forecast future demand. In order to understand the underlying rationale for the promotional forecasting process, it is important to understand the advantages and limitations of its underlying components.

Comparison between promotional and statistical forecasting

Statistical time series forecasting uses past demand to predict future demand. The most basic component of the time series forecast is the level of sales. This is usually determined by looking at demand in the recent past. Often there also exists an underlying trend that can be detected within sales history. This is usually determined by looking at the change in demand in the recent past. A third factor influencing retail demand is seasonality. A forecasting algorithm trying to determine the effects of seasonality can only look for periodic similarities in the sales pattern. For example, December sales from previous years can be used to adjust the forecast for December only because December occurs regularly every 12 months. At every step the time series approach is limited to using historical demand to predict future demand, without regard to the underlying causes that are driving demand in the first place.

Promotional events, however, can create problems in estimating level, trend, and seasonality. Certain events, such as irregularly occurring holidays, newspaper/radio advertisements, free gift with purchase offers, and special discounts can cause significant deviations from the selling pattern.

Promotional forecasting, unlike statistical forecasting, attempts to predict future demand by identifying the factors that influenced spikes and lulls in demand in the past. Once these factors are known, the magnitude and direction of their effect can be evaluated. Their presence can then be incorporated into forecasting models for use during times when the causal factors are again expected to be present.

Developing promotional forecast methods

This section describes how custom promotional forecast models have been developed in the past, leading to the discovery of several consistent findings. These findings have been incorporated into Retek's development of the Promote forecasting module.

Promotional forecasting uses promotional factors to predict future demand. The first step is to determine all of the pertinent information affecting sales and transform this information into variables that the system understands. Seasonality, for instance, can be represented by a single seasonal continuous variable, such as the number of daylight hours or average daily temperature.

Alternatively, it can be represented by 12 different indicator variables representing each of the months. An indicator variable consists only of 0's and 1's, a 1 indicating that the event is "on". For example, a monthly indicator variable for January would consist of a 1 during the first month of the year and 0's for the remaining months.

Once a list of variables is determined, the model needs to represent the promotion events in terms of their influence on overall demand. For example, if a set of promotional variables has a multiplicative promotional effect on demand, a log transformation may be needed to improve the model. After a suitable model is developed, it must be implemented using multivariate linear regression or neural network architecture, with custom code handling the data loading and variable transformations. The final custom model may be quite accurate over the data set on which it was developed. However, this model may not be general enough to be used universally across all data sets, thus requiring the development of multiple custom models to cover a client's entire domain. This has been found to be very time consuming and costly.

The process of developing custom promotional models has, however, brought to light a number of consistent patterns:

- Level, trend, and seasonality are universal components of almost any forecast.
- Including a time series forecast as an input variable often improves promotional models.
- Indicator variables are robust in that they can represent both additive and multiplicative effects.

These findings have led Retek to develop a novel approach to promotional forecasting that combines the automation and generalization of time series forecasting with the improved (albeit data set specific) accuracy met through customized causal forecasting.

Retek's approach to promotional forecasting

Retek's approach to promotional forecasting is somewhat unusual. Retek combines time series forecast methods with causal forecast methods, resulting in a new forecast method supported by the Promote module. Promote uses the AutoES method of forecast generation to determine a baseline time series forecast, and then uses indicator variables to represent promotional events of interest. By giving the forecasting routine visibility as to when certain events occurred in the past, the system can correlate changes in the sales demand with these events and determine each promotional event's effect on sales. Then whenever these events occur in the future, the promotional effects can be incorporated into the forecast.

The Promote module has been developed to produce generalized promotional models automatically, with little human intervention. Combined with the system's ability to allow you to develop your own data loading routines, Promote provides a cost-effective means for producing forecasts using promotional information.

More detailed information on promotional forecasting methods is in Chapter 6.

Promotional forecasting terminology and workflow

Promote is designed to produce sales forecasts using both past sales history and event on/off information, both of which you provide. Using the sales data, the system first determines a seasonal time series model to describe the purchasing behavior of consumers. Differences between the seasonal model and the actual sales are then correlated with known events. Those events that are found to have a statistically significant impact on sales are then included in a promotional forecast model as promotion events. For each promotion event, its promotion effect on sales is determined. The final promotional model consists of the seasonal model, promotion event on/off information, and each promotion event's resulting effect on sales. By combining these three, a final promotional forecast is computed.

Examples of promotion events

The following are examples of promotional variables that could be created and the manner in which their associated on/off event status is specified:

Example 1: Christmas Day applies to all products/locations. The Christmas promotional variable will therefore have only one dimension – Day. Because Christmas Day falls on the 359th day of each non-leap year, the Day359 variable will be set to TRUE for every such year (all other days will be set to FALSE).

Example 2: A television advertisement is run locally in the New York/New Jersey area for the four weeks at the beginning of the spring fashion season. The TVAD promotional variable will have two associated dimensions – State and Week. Week13, Week14, Week15, and Week16 will be set to TRUE only for states NY and NJ (all other states/weeks will be set to FALSE). Since no product dimension exists, the TV ad is assumed to have an effect on all products.

Example 3: A holiday promotion is run involving all sporting goods items for the two weeks prior to Father's Day. The Father's Day promotional variable has two associated dimensions – SKU and Week. For this year, Week23 and Week24 will be set to TRUE only for individual SKUs related to sporting goods items (all other weeks and all other SKUs will be set to FALSE). Since no location dimension exists, the Father's Day promotion is assumed to apply to all stores.

After promotional variables have been loaded into Retek Demand Forecasting, you can use the Promotion Planner workbook to view, edit, and update associated values without having to reload new data.

Promote Workbooks and wizards

The Promote Workbooks and wizards allow you to manage the promotion events used in the system's promotional forecasting processes, and view/edit the system's analysis of the effects of these events on demand. The Promote Workbook templates include:

- Promotion Planner - Allows you to specify when certain promotional events were active in the past, and when they will be active in the future.
- Promotion Selector – A wizard that allows you to specify the subset of loaded promotional events that the forecasting routine evaluates during the forecast generation process.
- Causal Maintenance – Allows you to review the system-calculated promotional lift effects, edit these effects and determine how changes will be factored into the promotional model.
- Promotion Analysis and Planning – Our promotional what-ifing workbook. Allows the use to analyze the performance of previous promotions and simulate future promotions by editing when promotional events will be active at for an item/location, and modifying the system-calculated lift effects.

Promotion Planner Workbook Template

In order to correlate deviations from the seasonal forecast with the occurrence of historic promotion events, the system needs visibility as to when these events were active. The system must also be informed of dates on which the status of upcoming promotion events will again be “on,” so the anticipated promotion effects can be built into the forecasting model.

The Promotion Planner Workbook allows you to indicate to the system when certain events were active in the past, and when they will be active in the future. All promotional events should be represented as accurately as possible so the modeling routine can more precisely detect correlations between event occurrences and changes in sales values.

The Promotion Planner Workbook consists of as many worksheets as are necessary to represent all unique dimensional intersections associated with the promotion events contained in the workbook. A separate worksheet is constructed for each of the required intersections. For example, promotion events such as Advertisement and Gift with Purchase are typically loaded at the [SKU/Store/Week] intersection, while an event such as Christmas is loaded at the [Day] level.

In this setup, the Advertisement and Gift with Purchase measures would appear on one worksheet, and Christmas would appear on another. Whenever a hierarchy is missing from an intersection specification, as in the case of the Christmas promotional event, the event is assumed to apply to all positions in that hierarchy. Thus, Christmas is assumed to apply to all products and all locations, but only to the Day-level calendar position(s) specified in the Promotion Planner Worksheet.

Promotion Planner wizard

The Promotion Planner wizard steps you through the process of creating a new Promotion Planner Workbook from a template. To access the Promotion Planner, select Open from the File menu to bypass the wizard and open an existing Promotion Planner Workbook, or perform the following steps:

- 1 Select New from the File menu.
- 2 On the Promote tab, select Promotion Planner and click OK.
- 3 The Promotion Planner wizard opens and prompts you to select the promotion events whose on/off status you wish to specify or edit. Use Ctrl-Click or Shift-Click to select multiple measures. Click Next.
- 4 Subsequent two-tree wizard screens appear to allow you to select the calendar positions (and possibly locations and/or products) to include in your workbook. A selection screen only appears for the hierarchies that are associated with the promotional event measures selected on the previous screen. On each selector, highlight the appropriate items on the Available side of the screen and click the right arrow button to move them to the Selected side. Click Next.
- 5 When the Next button is no longer enabled, click Finish to build the workbook.

Promotion Planner Workbook and Worksheet

The Promotion Planner Workbook allows you to view and edit the on/off information associated with each loaded promotional event. This workbook provides an interface in which you can specify the time periods (and possibly products and/or locations) for which certain promotional variables are active.

	07/24/99	07/31/99	08/07/99	08/14/99	08/21/99	08/28/99
AD	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Gift with Purchase	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Purchase with Purchase	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Radio	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
TV	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Promotion Planner Worksheet

A check box exists for each calendar/promotion event measure intersection. If a check mark appears in the checkbox for a given cell, then the associated promotion event's status is "on" for that calendar period. If no check mark exists, then the event's status is "off."

Place check marks as required to most accurately describe each product/location/calendar combination's event status. When you have made your changes, commit the workbook by selecting a Commit option from the File menu.

Promotion Selector Template

The Promotion Selector template is a wizard that allows you to specify a subset of loaded promotional events to be considered by the modeling routine as it evaluates the effects of various promotions on baseline sales.

A large number of promotional variables may be loaded into the system, and thus are available for consideration in the generation of promotional forecast models. To save processing time, you may wish to exclude certain promotional events that you know you do not wish to evaluate. The Promotion Selector consists of a standard two-tree wizard selection screen that allows you to update the set of events that you want to keep registered. If you use the Promotion Selector to remove an event from consideration, that event will no longer be available for selection in the workbook build templates for Promotion Analysis and Planning and for the Promotion Planner.

Promotion selection is generally performed soon after system installation, and may not be required on a frequent basis, if at all.

Promotion Selector Wizard

The Promotion Selector wizard steps you through the process of modifying the subset of loaded promotional events considered by the modeling routine. Based on your selections in this wizard, the system registers any selected (but currently unregistered) promotion events, and unregisters any unselected (but currently registered) events.

To access the Promotion Selector, perform the following steps:

- 1 Select New from the File menu.
- 2 On the Promote tab, select Promotion Selector and click OK.
- 3 The Promotion Selector wizard opens and prompts you to select the promotion subset that you want to keep registered. Use the right arrow button to move all desired items to the right side of the selection screen. Use the left arrow button to move all undesired items to the left.
- 4 When you have made your selections, click the Finish button to invoke the registering/unregistering process.

Causal Maintenance Workbook Template

The Causal Maintenance workbook (CMA) provides a view to the system-calculated and adjusted lift effects. You can edit effects at any product/location intersection and determine how these changes will be factored into the promotional models. The Causal Maintenance Workbook contains one worksheet: Final PromoEffects.

Causal Maintenance Wizard

The Causal Maintenance wizard steps you through the process of creating a new Causal Maintenance Workbook from a template. To access the Causal Maintenance Workbook Template perform the following steps:

- 1 Select New from the File menu.
- 2 On the Promote tab, select Causal Maintenance and click OK.
- 3 The Causal Maintenance wizard opens and prompts you to select the forecast level for analysis. Click Next.
- 4 Select the promotion events to analyze. Click Next.
- 5 Subsequent two-tree wizard screens appear to allow you to select the items and locations to include in your workbook. On each selector, highlight the appropriate items on the Available side of the screen and click the right arrow button to move them to the Selected side. Click Finish.

Causal Maintenance Workbook and Worksheets

Final PromoEffects Worksheet

The Final PromoEffects Worksheet allows you to view and modify the system-evaluated additive effects and multiplicative effects of a given promotion.

Product	Location	
Bounty Decorator Towel	30/1	Barcelona
Promotion: Center Page 1 Effects Calculated		3.54
Promotion: Center Page 1 Effects Override		1.00
Promotion: Center Page 1 Type		Automatic
Promotion: Circular 1 Effects Calculated		1.00
Promotion: Circular 1 Effects Override		1.00
Promotion: Circular 1 Type		Automatic
Promotion: Discount Percent 1 Effects Calculated		1.00
Promotion: Discount Percent 1 Effects Override		1.00
Promotion: Discount Percent 1 Type		Automatic
Promotion: Front Page 1 Effects Calculated		1.00
Promotion: Front Page 1 Effects Override		1.00
Promotion: Front Page 1 Type		Automatic
Promotion: In Store Display 1 Effects Calculated		1.00
Promotion: In Store Display 1 Effects Override		1.00
Promotion: In Store Display 1 Type		Automatic

Final PromoEffects Worksheet

Effects

The system-calculated lift effect.

Effects Override

The user-specified lift effect.

Type (Causal Variable Type)

Causal variable types define how causal variables are treated in the causal model fitting process (which includes a call to the lower level regression engine) and the forecast generation process where the model is used to extend the forecast over the forecast horizon. Following are the drop-down list options:

Automatic

The inclusion of the causal variable is decided by regression. If the causal variable is found to be significant on the training set it is included in the model, otherwise it is rejected.

Force In The Model

The causal variable is forced in to the model. Thus regression is not given a choice to reject even if the effect is considered insignificant by regression. As a result we will always return an effect.

Disabled

The variable is excluded from the model. Hence no effect will be returned either.

Override All

This type allows the user to specify a causal effect that will be used during the fitting and forecasting process. This is a directive that is recognized only by the causal engine and not by the lower level regression engine. For causal variables specified as *Override All*, the user also specifies the corresponding causal effect. The causal engine then de-causalizes the training data using the user-specified effect. The variable then is internally set *Disabled* to calculate the fit. During forecast generation the user-specified effect is used to determine the causal forecast.

Override Future

This type allows the user to specify a causal effect that will be used only during forecasting process and not during the fitting process. This is also a directive that is recognized only by the causal engine and not by the lower level regression engine. For causal variables specified as *Override Future*, the user also specifies the corresponding causal effect. The causal engine then internally sets the causal variable to *Automatic* to calculate the fit. The calculate effect is not however written back to the effects array (so as to not overwrite the user specified effect). During forecast generation the calculated effect is ignored and instead the user-specified effect is used to produce a causal forecast.

Override Higher Level

This type allows the system to use an average value of causal effects computed from product/location combination in the same group (the intersection level of the group to be used being specified during system setup) for product/location combinations on whose sales history alone a causal effect is unable to be computed for that causal variable. The way the system handles this is by having the effects array filled in with higher level effects for those variable specified as *Override Higher Level* and which did not have an effect returned in the previous run. The effect array is passed to the causal engine. If the engine succeeds in computing a significant effect it writes back the system computed effect and uses it for generating a forecast. Otherwise it uses the effect passed in (which as mentioned is the average effect among product stores belonging to the same group) for generating the forecast.

Promotion Analysis and Planning Workbook Template

The Promotional Analysis and Planning Workbook (PAP) Template is a historical and future view to the effects of a planned promotion.

This workbook has three worksheets:

- Writable view of data from the Promotion Planner Workbook
- Writable view of the Causal Maintenance Workbook.
- Promotion Effects Analysis Worksheet: Visibility to Actuals, the forecast, future baseline

In this workbook a user views the promotional forecast. There is also the ability to analyze the effects on the forecast if a promotion does or does not occur. This analysis is performed by turning a promotional event or attribute on or off for dates in the future plans. This data along with promotional costs, actual sales and gross margin of sales, helps determine the financial impact of the promotion.

Using these measures, we can calculate after the fact promotional effectiveness, in terms of profit generated, cost of promotions, and promotional payback. This workbook is intended for simulation and analysis purposes only. There is no 'Commit' functionality.

Promotion Analysis and Planning Wizard

The Promotion Analysis and Planning wizard steps you through the process of creating a new Promotion Planning and Analysis Workbook from a template. To access the Promotion Analysis and Planning wizard, perform the following steps:

- 1 Select New from the File menu.
- 2 On the Promote tab, select Promotion Analysis and Planning and click OK.
- 3 The wizard opens and prompts you to select the forecast level on which to perform effectiveness analysis. Click Next.
- 4 Subsequent two-tree wizard screens appear to allow you to select the historic or future dates in the calendar, the items and location to use in analysis and to include in your workbook. Click Next.
- 5 Select the promotion events to analyze. Click Next.
- 6 Select any additional measures to be viewed in the workbook. Click Finish.

Promotion Analysis and Planning Workbook

This workbook is used for the evaluation of historic promotions and to simulate future promotions by modify when promotions are active in the future and the lift effect associated with the promotion. This workbook is intended for promotion analysis and simulation, therefore there is no 'commit' capabilities. The following is an example of how the Promotion Analysis and Planning workbook can be viewed:

The screenshot displays the Retek Promotion Analysis and Planning Workbook with three worksheets visible:

- Promotion Planning Worksheet:** Shows a calendar view for the product '10000141 Tide Regular' at location 'Berlin'. It lists various promotion types (Circular, Center Page, Front Page) and their dates (01/05/02 to 02/16/02).
- SKU/Store/Week - Final Promotion Effects Worksheet:** Shows the effects of the promotion on various measures (e.g., Promotion: Circular 1 Effects, Promotion: Circular 1 EffectsOvr, etc.) for the same product and location.
- SKU/Store/Week - Final Promotion Effectiveness Analysis Worksheet:** Shows a detailed table of promotional effectiveness analysis, including measures like Promo: Ad Block, Promo: Ad Page Number, Promo: Ad Type, Promo: Discount Percent, Promo: Promo Counter, Promo: Promo Retail, Promo: Regular Retail, Actual, Forecasted Baseline, Forecasted Gross Margin, Forecasted Incr GM, Forecasted Lift, Forecasted Sales - Baseline, Forecasted Sales - Lift, Forecasted Sales - Total, and Promotional Forecast, across a range of dates from 01/05/02 to 04/06/02.

Promotion Analysis and Planning Workbook

This workbook is a view to the promotion calendar, the promotional lift effects and the resulting promotional forecast.

- 1 The user may view and modify both the Promotion Planning and Causal Maintenance Worksheets (see previous sections on the use of these worksheets).
- 2 Select Calculate to see how changes to the Promotion Planning and Causal Maintenance Worksheets effect the Promotional Forecast that is viewed in the Promotion Effectiveness Analysis Worksheet.

Procedures in Promotional Forecasting

The following outlines the standard procedures performed in order to set up the system to run a promotional forecast

Set up the system to run a promotional forecast

Follow this procedure to set up your system to run a promotional forecast.

- 1 Select New from the File menu.
- 2 On the Predict tab, select Forecast Administration and click OK.

Note: This procedure sets the default forecast generation method such that a promotional forecast is run for all items (unless an alternative method is specified for individual items in Forecast Maintenance). If you only want to run a promotional forecast for a small number of items, open the Forecast Maintenance Workbook instead and select the relevant product/location combinations.

- a Select the final forecast level and click Finish.
 - b In the Forecast Administration Workbook: Advance Parameter Worksheet, select 'Generate Baseline'
- 3 In the desired Forecast Level Worksheet or item/location intersection (if using Forecast Maintenance Workbook), set the Generation Method measure to "Causal."
- 4 If you intend to run a promotional forecast manually using the Run Batch Forecast wizard, make sure there is a value specified in the Horizon Start Date measure. A forecast will not be generated if the value of this measure is NA.
- 5 Commit your changes to the master database by selecting Commit Now from the File menu.
- 6 Close the workbook by selecting Close from the File menu. You are now ready to access the Run Batch Forecast wizard to generate your forecast at this time.

Chapter 4 – Curve (Profile-based Forecasting)

Overview

Curve is an optional automated predictive solution that can generate ratio arrays from historical data at user-specified intersections. The profiles generated by Curve can be used for various purposes, for example, to convert the organization-level assortment plans into base-level weekly sales forecasts, and to generate seasonal forecasts, daily forecasts, or new product forecasts using life cycle profiles.

Curve meets the need of operational systems (such as Retek Demand Forecasting (Demand Forecasting) and the Retek Merchandising System (RMS)) to have sales predictions at a more detailed level than those provided by corresponding planning programs. The planning process attempts to establish the correct balance between different products in order to maximize sales opportunities in the available sales space. The planning process is supported by the generation of an assortment plan, which provides details of your anticipated sales volumes and stock requirements at aggregated levels. However, operational systems like Retek Demand Forecasting and RMS require base data at the SKU-chain-week or SKU-store-week level, because these systems are responsible for ensuring that the right quantity of each product is in the right store at the right time.

To use Curve with Retek's Enterprise software (interface with the Merchandising system), custom work may be required. For information about performing this custom work, contact your Retek representative.

Profile Management Terminology and Workflow

This section outlines the basic functionality of the Curve module and defines terminology used in the context of generating profile-based forecasts. The profile management process is discussed, as are the workbook templates used to perform profile management functions.

Profiles (spreading ratios)

Profiles are generated using historical data and phase definitions based on the system configuration. Typical profiles include but are not limited to store participation, size distribution, and time profiles. The *aggregated* and *destination levels* of the forecast spreading (transformation) uniquely describe each profile. A concept that is closely related is the *unit-of-measure*, which defines the symbolic ratio of the destination level to the source level. Since a profile is used to spread the source level plan into the destination level, a naive way of calculating a profile is to obtain the ratio of sales at the destination level to the sales at the source level. In most situations, however, the sales at the destination level are sparse, noisy, or do not exist (new product or new store). Instead of generating the profiles at the destination level, the sales can be aggregated, and profiles can be obtained at higher levels as long as the unit-of-measure is appropriate. The aggregated level at which profile is generated is referred to as the *profile source level*. Profiles can be calculated at multiple source levels.

A separate profile is calculated for each hierarchy. Several profiles are used to spread the assortment plan into the final level, the user can generate daily profiles to be used for daily forecasting, seasonal profiles to be used in generating a seasonal forecast, or new product forecasting can be done using life cycle profiles. In order to improve the usability of the system, Curve has the capability of processing several profiles at the same time. Therefore, instead of working on one profile at a time, the user can focus on the *processing plan*, which defines the configuration of one or multiple profiles. For ease of implementation, if no plans are setup the system automatically assigns a plan for each profile level (entry in pgparm).

Steps in profile generation

For each profile / processing plan, you need to perform the following steps:

- 1 Set the parameters for profile building, using the Profile Administration and Profile Maintenance Workbooks.
- 2 Generate and auto-approve the profile(s), using the Run Batch Profile Workbook.
- 3 Adjust and approve the profile(s), using the Profile Approval Workbook.
- 4 Set the parameters for forecasting, using the Forecast Administration and Forecast Maintenance Workbooks, and/or spreading policies.
- 5 Run forecast generation.

The first three steps correspond to the profile generation, while the last two steps are specific to profile-related forecasting. The forecasting process should be set to using the profiles in specific ways, such as profile spread, seasonal indices, daily forecasting.

Also it is important to note that in most cases profiling and forecasting parameters are already set, and you can skip step 1 and step 4. A probable business scenario is that several traders will generate and approve profiles for the products that they are interested in. When this task is accomplished, a batch job will be started to run profile/forecast generation. This chapter addresses only the first three steps related to profile generation.

Profile generation techniques

This section describes the algorithms behind the profile generation process. Each profile typically acts on one dimension. A concept that is closely associated with each profile is the *unit-of-measure*, which defines the symbolic ratio of the destination level to the source level. We provide an example with three profiles. The first one is a *time profile*. The unit-of-measure for the time profile is week/phase. The two other profiles that are considered in the example are *size distribution* and *store participation*. Size distribution determines how much each size is sold in a given option. For a fashion item, option-size typically defines a sku. Therefore, size distribution is the profile to explode an option level plan into a sku level plan. Store participation, on the other hand, is the contribution of a store to the chain sales. The units of measure for these two profiles are sku/option and store/chain, respectively.

Building system profiles based on the historic level

Two techniques are developed for profile generation. The first profile generation technique is built on the assumption that historical information, which is relevant to the planned products, exists and future sales will follow the past sales patterns. Under these assumptions, a profile can be calculated as the ratio of sales at the destination level to the sales at the aggregated level. In the case of a continuity line, this is likely to be true. For a new seasonal product this may not be the case, but in many instances similar style, subclass, or class history will provide an accurate baseline to build the profiles against. This baseline is referred to as the *source level*. Note that, while generating the profiles at the source level, the unit-of-measure should be considered. Most profiles will be built based on historic data. Since the system allows the user to override the profiles, the technique has the unique feature of combining the trader's expertise and computer intelligence without heavily depending on extensive human interaction.

Overriding profiles by adjusting the source measure

Profile generation requires a metric containing historical data as input. In some cases the data is used as it is. At other times, however, it is necessary to make adjustments using overrides to correct seasonal and causal effects. A *source measure* is constructed from an input data metric and one or more override measures. Overrides might be used, for example, to compensate the effect of Easter changing in time each year. Another situation that might require overrides is non-recurrent causal events such as Euro 96. Multiple overrides are particularly useful for improving accuracy and provide better visualization of the effects of independent events. Source measures can be adjusted in the Source Measure Maintenance Workbook of the Analysis template group. Note that an alternative to sales overrides is to edit the profiles directly.

Defining the training window for historical data

Another feature of Curve is that only the part of the historical data that is relevant to the profile generation process is utilized for the calculations. This task is accomplished by specifying the start and end dates of the historic data that will be used for profile generation, that is, by specifying a *training window*. The user has the flexibility of defining default training windows for each profile. In general, a one-year long training window is appropriate for store participation. On the other hand, training window for size-distribution should be set such that the effects of the stock-out periods are minimized. The user can also use system-generated training windows that are based on the start and end dates of the current phase. It is recommended to use the system-generated windows for the time profile. As an example, for calculating the time profile for a *Spring-Summer phase* plan, the training window should be set such that only the data from the last year's Spring-Summer phase is used during the profile generation process. Since the start and end dates, that is, the *phase definition*, of the current phase is already loaded into the system, a simple yet effective rule to generate the training window is to subtract one year from the current phase definitions dates. The user also has the ability to modify the training window definitions using the Profile Maintenance Workbook. If the user does not specify the training window method, the system sets it by default to "use default training window".

Aggregating historical data

Because a profile is often used in forecasting to spread the source level plan into the destination level, a naive way of calculating a profile is to obtain the ratio of sales at the destination (final) level to the historic sales at the aggregated level. Considering the store participation example, this is equivalent to dividing the sku-store sales by sku-chain sales. In most situations, however, the sales at the final level are sparse, noisy, or do not exist (new product or new store). Instead of generating the profiles at the destination level, the sales can be aggregated, and profiles can be obtained at higher levels as long as the unit-of-measure is appropriate.

The store profile, for example, can be calculated as the ratio of subclass-store level sales to the subclass-chain level sales. The unit-of-measure of this ratio is still store/chain, as desired. The aggregated level at which profile is generated is referred to as the *source level*. Multiple source levels can be calculated and used during the explosion stage. The source level for the above example is subclass-store. Note that the time dimension does not exist. This implicitly means that sales are aggregated across the training window, and the store participation is assumed to be constant throughout the phase.

To demonstrate the process, consider the following scenario. A trader is trying to determine the store participation for shirts. The shirts are planned to be sold in four colors (black, navy, white, and blue) in the upcoming phase, which is four weeks long. The color "black" is newly introduced and the color "navy" was only sold in Lenox store previously. The sales history for the equivalent four-week training window last year for the four options and corresponding style is given below.

	Week 1	Week 2	Week 3	Week 4	Phase
Black Shirt	0	0	0	0	0
Navy Shirt	1	0	1	0	2
White Shirt	2	1	1	2	6
Blue Shirt	0	0	1	0	1
Shirts	3	1	3	2	9

Lenox store sales

	Week 1	Week 2	Week 3	Week 4	phase
Black Shirt	0	0	0	0	0
Navy Shirt	0	0	0	0	0
White Shirt	1	1	1	0	3
Blue Shirt	1	0	0	2	3
Shirts	2	1	1	2	6

Cumberland store sales

The ratios obtained at the option-store and style-store levels are provided below.

	Lenox	Cumberland
Black Shirt	NA	NA
Navy Shirt	100%	0%
White Shirt	67%	33%
Blue Shirt	25%	75%
Shirts	60%	40%

Store ratios

If the trader calculates the store participation ratios at the option-store level, then she cannot obtain any ratios for black shirts since no history exists for this color. Moreover, the ratios for navy shirts will be biased since this color hasn't been sold in Cumberland store before. An alternative is to calculate the ratios at aggregated levels such as style-store. The trader can also use multiple source levels (for example., style-store for black, navy and blue shirts, and option-store for white shirts). The motivation is to obtain more accurate ratios for stable products such as "white shirts" while using ratios obtained at higher levels for new or low-volume products. The trader can choose an appropriate source level for each product during the approval process.

Unreliable data, thresholds, and threshold exceptions

In certain situations, data may still be sparse even after aggregation. Sparse data generally causes unreliable ratios. Consider the following hypothetical situation:

	Store 1	Store 2	Store 3	Chain
Subclass A	1	1	1	3
Subclass B	100	100	100	300
Subclass C	100	200	100	400

Total sales

	Store 1	Store 2	Store 3
Subclass A	0.33	0.33	0.33
Subclass B	0.33	0.33	0.33
Subclass C	0.25	0.50	0.25

Ratios

The ratios for the subclass B and C are reliable and minor fluctuations in sales history will not significantly effect the ratios. On the other hand, the store participation ratios for the products belonging to subclass A can drastically change if the sales for the second store were two units instead of one. In fact, it is not clear if the subclass A behaves as subclass B or subclass C. Curve automatically generates a *threshold exception* for unreliable ratios when the sales are sparse at the source level. The generation rule is simple, if the dividend of the ratio is smaller than a threshold, then all ratios using this dividend is marked as exception. The user provides the threshold through the administration workbook. Since the user can view these exceptions during the approval stage, the exception mechanism can be viewed as a decision support system.

Loading default ratios

As previously stated, Retek provides an alternative approach to profile generation: loading default ratios. This option should be utilized if significant mismatches between historic and current sales patterns are anticipated. The profiles can be loaded at any level. The loading level is analogous to the source level from the system perspective.

Approving the generated profiles

In RDF 9.5, profile generation includes an auto-approval process, in which the system profiles are spread down and combined at the final level, and copied to the approved final profiles through an approval mask that is set up in the profile administration and maintenance workbooks

Up to this point we have discussed the algorithms involved with profile generation. The next step is running the Profile Approval Workbook. Profile approval involves viewing and editing the source level profiles, and choosing an appropriate source level for each product. The process will be described later in this documentation.

Once the profiles are approved, the system will run a series of algorithms so that these profiles can be utilized in the forecasting stage. In particular, the profiles at the source level are copied down to the final level. To demonstrate the process, consider the following example. Assume that the following size distribution for skirts is given at the “style-size”-chain level

Skirts-2	Skirts-4	Skirts-6	Skirts-8	Skirts-S	Skirts-M	Skirts-L	Skirts-XL	Skirts-32	Skirts-34
0.05	0.1	0.1	0.05	0.1	0.2	0.2	0.1	0.05	0.05

Size profile for Skirts

Let’s assume that three colors are available for skirts: blue, black, and red. These three options use the same size distribution profile during the forecasting stage. However, it is possible that not all sizes are available for each option. As an example, assume that blue is sold in sizes [2,4,6,8], black is sold in sizes [S,M,L,XL], and red is sold in sizes [32,34]. The size distribution at the style-size level can not be directly used at the option level. Instead, the ratios of the available sizes should be modified for correct forecast results. There are two rules for modification. First, the profile ratios should still add up to one at the option level. Second, the relative ratios of two different sizes should be preserved. The normalized ratios are as follows.

Blue Skirt-2	Blue Skirt-4	Blue Skirt-6	Blue Skirt-8
1/6	1/3	1/3	1/6

Size profile for blue skirts

Black Skirt-S	Black Skirt-M	Black Skirt-L	Black Skirt-XL
1/6	1/3	1/3	1/6

Size profile for black skirts


Red Skirt-32	Red Skirt-34
1/2	1/2

Size profile for red skirts

Similar normalization is required, if certain stores are known to not participate in certain options.


After approving the profiles, you can start the forecasting stage. Curve is algorithmically similar to spreading source level forecast into the final level. Forecast spreading via Curve is typically accomplished in a series of steps as illustrated in the figure. The planned sales level is typically the default generation level. It is at this level that sales plans are loaded. The SKU/Store/Week level in this example is the final forecast level. The two intervening levels do no more than describing the transformation path from the assortment plan level (option/chain) to the sku/store/week level.

Planned Sales at Option/Chain




Apply Option/Chain/Week Profile

Weekly Sales at Option/Chain



Apply SKU/Chain/Week Profile

Weekly Sales at SKU/Chain



Apply SKU/Store/Week Profile

Weekly Sales at SKU/Store

Transforming an assortment plan to an operational sales plan

Profiles are mostly used for spreading, including daily profiles (spreading from week to day) and life cycle (spreading a plan to the week level). Besides spreading, Curve generated Seasonal profiles are used by the forecasting engine for generating profile-based forecasts. The Curve module should be configured for generating Seasonal profiles at the forecast level intersection or higher. The forecasting part is handled by Predict and its setup and use are described in the Predict documents.

Profile Administration Workbook

The Profile Administration Workbook allows you to set default parameters for profile generation, the first step in profile generation. These parameters are typically set during system implementation, and are configured based on your business practices and needs. This configuration can be updated if you need to change certain parameters over time. However, it is not practical to change the configuration on a regular basis. The Profile Administration Workbook gives you the flexibility to change profiling parameters as the need arises to improve both forecasting accuracy and computational efficiency.

Profile Administration Wizard

The Profile Administration wizard steps you through the process of creating a Profile Administration Workbook from the template. This workbook allows you to set default parameters for profile generation.

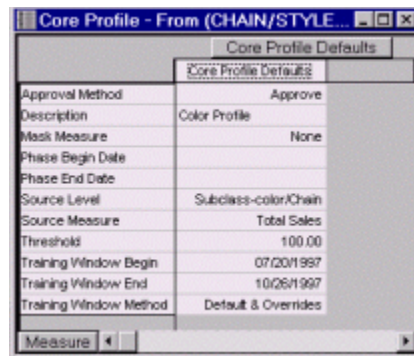
The Profile Administration wizard requires you to select the processing plan that you wish to edit. The number and the names of the available processing plans are specified at system implementation. When the Profile Administration Workbook is created, a separate worksheet will exist for each profile generation level associated with the selected processing plan. For example, for a business with two processing plans available, the wizard screen might display these two choices:

- Profile from Option/Chain/Phase to SKU/Chain/Week: used to generate profiles for the time and product hierarchies. Two Core Profile Worksheets would be created: one for the Option/Chain/Phase to SKU/Chain/Phase transformation, and one for the SKU/Chain/Phase to SKU/Chain/Week transformation.
- Profile from SKU/Chain/Week to SKU/Store/Week: used to generate store profiles. One Core Profile Worksheet would be created, associated with the SKU/Chain/Week to SKU/Store/Week transformation.
- New: used to create new profile templates. For more information, see “Creating new profile templates” on page 98.

When you have selected a processing plan to edit, click the Finish button to open the Profile Administration Workbook. The workbook contains Core Profile Worksheets for each profile generation level associated with the selected processing plan.

Core Profile Worksheet

The Core Profile Worksheet allows you to specify default values for parameters affecting profile generation. Values must be specified for each profile generation level in your selected processing plan. A separate Core Profile Worksheet exists for each transformation, and the worksheet title bar indicates the source and destination level of each.



Core Profile Worksheet

The following is a description of the measures contained on the Core Profile Worksheet:

Approval Method

The default approval strategy for each profile. Choices are Approve and Don't Approve. The default approval method is automatically copied into the Approval Method field for each product brought into the Profile Approval Workbook.

Description

Displays the label for the profile generation level.

Mask Measure

Displays the profiling mask. In many cases, product availability in each store changes from one year to the next. The profiling mask is used to exclude from the exploded plan SKU/Store combinations that are not available in the coming phase. The mask contains the list of SKUs that are available for each store. A Mask Measure of None indicates that no mask exists.

Phase Begin Date

The default phase begin day. Change the value in this field by clicking the pop-up calendar. If phase definitions are not available, the default phase begin date will be used. Also, when calculating time profiles, default dates are used for intermediate computations. For computational efficiency, use the most common phase definition as the default value.

Phase End Date

Displays the default phase end day. Change the value in this field by clicking the pop-up calendar. If phase definitions are not available, the default phase end date will be used. Also, when calculating time profiles, default dates are used for intermediate computations. For computational efficiency, use the most common phase definition as the default value.

Source Level

The default source level. You can choose a source level using a pick list. A list of possible source levels is determined at the implementation for each final profile level.

Source Measure

The source measure to be used for profile generation. The source measure is built by combining historic input data with any associated override measures. You can select the appropriate source measure from a drop-down list.

Threshold

The default historic threshold, or level at which historic sales are considered too sparse to generate reliable profiles. If the aggregated data, used as a dividend in ratio calculations, is not greater than this threshold, an exception will be generated. Exceptions are displayed in the Profile Approval Workbook.

Training Window Begin

The default training window start date. Change the value in this field by clicking the pop-up calendar. The default date will be used only if the training window method is Default and Overrides.

Training Window End

Displays the default training window end date. Change the value in this field by clicking the pop-up calendar. The default date will be used only if the training window method is Default and Overrides.

Training Window Method

Displays the default method used to obtain training window. The options are:

- **Default & Overrides:** Uses the default dates as set in the Training Window Begin and Training Window End measures.
- **Phase Definitions & Overrides:** Calculates the training window dates based on the phase definitions.
- **Last 10 & Overrides:** Uses the last ten periods of data to define the training window.

All methods are overridden if you enter specific training window values for a subset of products in the Profile Maintenance Workbook.

Creating new profile templates

Using the Dynamic Profile Library feature, you can create new profiles through the Profile Administration Workbook. This new RDF feature greatly simplifies the process of extending the available profile templates, which previously had to be created outside of the RDF user interface. For details, see the procedure “Create a new profile template”.

Profile Maintenance Workbook

After setting default parameters for profile generation in the Profile Administration Workbook, the next step in profile generation is to select any subset of items for which the training window values should differ from the defaults. This step is necessary in those situations where it is not efficient to use the same parameters for all products. The act of selecting item subsets is known as ranging. Ranging is performed through the standard two-tree item selection screen in the Profile Maintenance Wizard.

Profile Maintenance Wizard

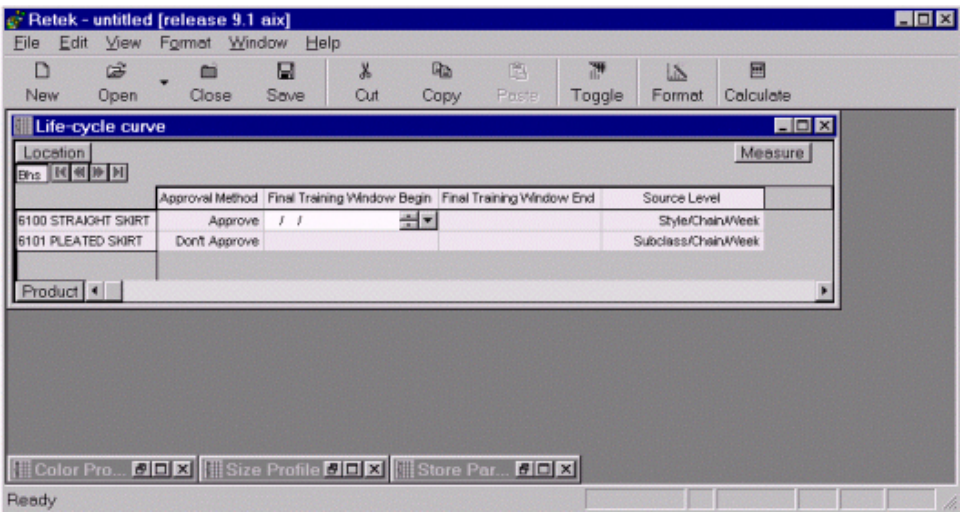
As in Profile Administration, the first Profile Maintenance wizard screen prompts you to select a processing plan to maintain. When the Profile Maintenance Workbook is created, a separate worksheet will exist for each profile generation level associated with the selected processing plan.

The next step in the wizard process is to select the subset of items to appear in the workbook. Depending on your domain, you may be asked to select items from the product and/or location hierarchy. Based on specifications set at system implementation, the system automatically determines the hierarchical level at which ranging is performed. If these specifications were not provided at installation, the system automatically determines the ranging level based on the destination level of the profiles.

When you have made the required selections, click the Finish button to build the Profile Maintenance Workbook. Separate worksheets will be created for each profile generation level in the processing plan, and each worksheet will contain only the items selected in the ranging process.

Profile Maintenance Workbook

The Profile Administration Workbook allows you to set default values for profile parameters. However, it is not always efficient to use the same training window for all items. The Profile Maintenance Workbook allows you to edit the training windows, approval methods, and source levels separately for each profile.



Profile Maintenance Workbook

The following is a description of the measures contained in the Profile Maintenance Workbook:

Approval Method

The default approval strategy for each profile. Choices are Approve and Don't Approve.

Final Training Window Begin

Displays the override training window start date. Change the value in this field by clicking the pop-up calendar. This value overrides any default training window or phase begin dates set in Profile Administration.

Final Training Window End

Displays the override training window end date. Change the value in this field by clicking the pop-up calendar. This value overrides any default training window or phase end dates set in Profile Administration.

Source Level

The default source level. You can choose a source level using a pick list. A list of possible source levels is determined at the implementation for each final profile level.

Run Batch Profile

The Run Batch Profile template consists of a wizard that steps you through the process of manually executing a profile calculation batch job.

A single-select wizard allows you to select the processing plan for which the profiles will be calculated. Your batch job will create a database to hold the profiles, and then will calculate the profiles and exceptions for the levels specified in the processing plan. Both the profiles and exceptions will be created at the proper source levels. Curve allows you to calculate values at multiple source levels for each item profiled. An auto-approval process is run after the source level profiles have been generated. The process consists of spreading the source level profiles to the final level, and then combining them according to the source level selections made in the Profile Administration and Profile Maintenance Workbooks. The final level profiles are approved with the approval methods selected in the Profile Administration and Profile Maintenance Workbooks.

Note: Before you start a Run Batch Profile process, make sure that no other users are logged onto the Retek Demand Forecasting system. Other users cannot access the system until after the batch run has finished.

Profile Approval Workbook

The profiles generated at the historic levels must be viewed, analyzed, revised, and approved using the Profile Approval Workbook. In the approval process, you select the appropriate source level for each product/location combination. After you make any necessary changes to the profiles and commit the workbook, the profiles are normalized to preserve the appropriate ratios. At this time, Curve automatically spreads the source level profiles to the final level and combines them. After you commit your changes, you can refresh the data in your workbook to display the newly generated final level profiles.

Use the Profile Approval Workbook to view, analyze, revise, and approve the profiles generated at the historic levels. This workbook contains three types of worksheets:

- Profile Approval Worksheet
- Source Level Worksheet
- Final Profile Worksheet

Profile Approval Wizard

The Profile Approval wizard steps you through the process of creating a new Profile Approval Workbook.

The Profile Approval wizard first prompts you to select the processing plan to be approved. When the Profile Approval Workbook is created, a separate worksheet will exist for each profile generation level associated with the selected processing plan. The next step is to select the subset of items to appear in the workbook. Depending on your domain, you may be asked to select items from the product and/or location hierarchy.

When you have made your selections, click the Finish button to build the workbook.

Profile Approval Worksheet

The Profile Approval Worksheet allows you to review and approve profiles by product and location. This worksheet also allows you to select the appropriate historic level to be used when profiles are spread down to the destination level.

The following is an example Profile Approval Worksheet for the store profile.

Location	Product
Bricks & Mortar	10000043 Kangaroo Pock 10000051 Bohemian 5 Pox 10000064 Carpenter Jean 10000067 5-Disk CD Playe
Approval Date	
Approval Method	Don't Approve Don't Approve Don't Approve Don't Approve
Approved By	
Source Level	Store/Subclass Store/Subclass Store/Subclass Store/Subclass
Measure	

Profile Approval Worksheet

The Profile Approval Worksheet contains the following measures.

Approval Method

The approval status for each profile. Choices are “Approve” and “Don’t Approve”. The default approval method is automatically copied into the Approval Method field for each product brought into the Profile Approval Workbook. When you commit the workbook, all product/location combinations with an Approval Method set to ‘Approve’ will be approved.

Approved By

Displays the user who approved the profile for a given product/location combination.

Approval Date

Displays the date on which a profile was approved.

Source Level

Displays the source level at which a profile should be approved. The source level is the aggregated level at which the profile is generated. The source level is sometimes equivalent to the destination level, but in instances where sales are sparse, historic sales are aggregated and profiles obtained at a higher level. Profiles for multiple source levels can be calculated; a separate Source Level Worksheet exists for each in the Profile Approval Workbook. In addition, you can choose different source levels for different products.

Source Level Worksheet

The Source Level Worksheet displays the profiles generated at the source level for all product/location combinations selected to appear in the workbook. Editing the values on this worksheet modifies the final level profiles. That is, the values that are edited are changed in the final profile. The remaining values are normalized such that the profiles preserve the ratio property. This worksheet also displays the source system profiles, that is, the profiles that are calculated by the system, and the threshold exceptions detected during the profile generation process, indicating values that may be unreliable due to sparse historic data conditions.

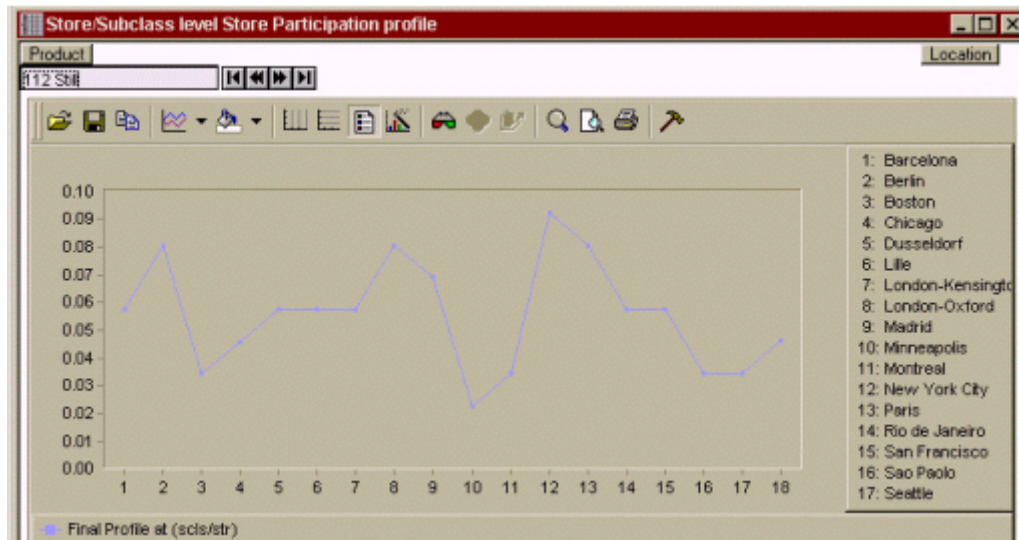
A separate Source Level Worksheet exists for each historic level at which profiles can be generated. You can view and compare the profiles generated at the different levels in order to select the one that should be spread down to the destination level.

Following is an example Source Level Worksheet for a life-cycle curve profile. In this example, the source level is style-chain-week. The values for week 4 and week 5 are edited. The example depicts the worksheets after these edits are entered to the system. Note that the values for relative week 1 through relative week 52 add up to one for final and system profiles.

Product	Location	Barcelona	Berlin	Boston	Catalog Store	Chicago	Dusseldorf	E-Os
T2 Still	Final Profile at (scis/str)	0.06	0.08	0.03	1.00	0.05	0.06	1.00
	System Profile at (scis/str)	0.06	0.08	0.03	1.00	0.05	0.06	1.00
	Threshold Exception at (scis/str)							

Source Level Worksheet for Profile Approval

An alternative representation of this worksheet is in the graph or chart form. The source level system profiles for the Still water product are provided in the next figure. You can modify the source profile field in the chart mode as well. At the source level, the time dimension is represented in terms of relative weeks. The first relative week corresponds to the first week of the phase.



Source Level Worksheet for Profile Approval – chart form

The following measures are contained in the Source Level worksheet:

Final Profile

A read-only measure that displays the profile value that will be approved when the workbook is committed. When a new Profile Approval Workbook is built, the values in this field match those in the System Profile measure.

System Profile

A read-only measure that displays the system-generated profiles calculated at the historic level for each product/location combination displayed.

Threshold Exception

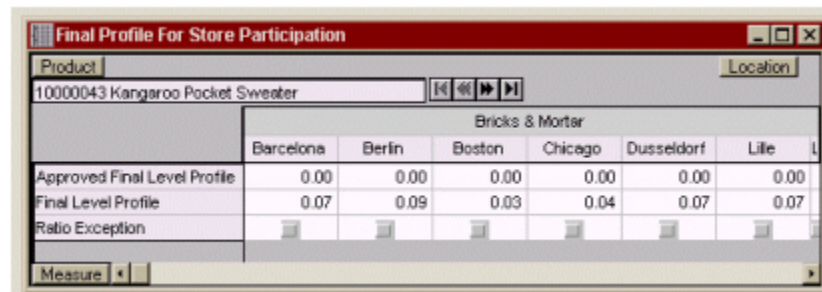
A checkbox measure that warns you when the system has detected a profile value that may be considered unreliable due to sparse historic data. A threshold for profile calculations is set in the Profile Administration Workbook. When source measure values at the historic level fall below this threshold, a threshold exception is generated.

Final Profile Worksheet

The final type of worksheet is the Final Profile worksheet, which allows you to view the profiles at the final level. This worksheet is initially empty when the Profile Approval Workbook is first built, and is updated when the Approval Workbook is committed to the master database and refreshed. .

Through this workbook, you can view the final profiles and ratio exceptions. A ratio exception is generated if a particular SKU is available in a store (where availability is determined with respect to the mask measure) and the system generates a zero ration for that SKU. Note that all the options in the given style have the same ratios.

Following is an example Final Profile Worksheet for a store participation profile.



Product	Bricks & Mortar						Location
	Barcelona	Berlin	Boston	Chicago	Dusseldorf	Lille	
10000043 Kangaroo Pocket Sweater							
Approved Final Level Profile	0.00	0.00	0.00	0.00	0.00	0.00	
Final Level Profile	0.07	0.09	0.03	0.04	0.07	0.07	
Ratio Exception							

Final Profile Worksheet

The following measures are contained in the Final Profile Worksheet:

Approved Final Level Profile

Displays the approved profiles at the destination level. The destination level in the example is PrBr/UCSh (Pure Branch – Unit Code/Shade). These profiles will be used in the assortment plan explosion process. Approved final level profiles only appear for product/location combinations whose approval method was set to “Approve” (on the Profile Approval Worksheet) before the workbook was committed and refreshed.

Final Level Profile

Displays the final level profiles at the destination level, whether the profiles were approved or not. For product/location combinations whose approval method was set to “Don’t Approve” (on the Profile Approval Worksheet) before the workbook was committed and refreshed, the Final Level Profile is still calculated, but this value is not copied over to the Approved Final Level Profile measure.

Ratio Exception

A checkbox measure that warns you when the system has detected a ratio exception. A ratio exception is generated if a product is available in a given store (as determined with respect to the mask measure), and the system generates a zero ratio for that item.

Procedures in Profile-based Forecasting

The following is a list of standard procedures performed when conducting profile-based forecasting:

- Set default parameters for a profile
- Modify the training window for a subset of items in a profile
- Manually request a generated profile
- Approve a profile
- Set Retek Demand Forecasting to perform plan explosion
- Create a new profile

Set default parameters for a profile

Use this procedure to set the default parameters for profile generation. Typically, values for these parameters are set during system installation and need not be modified on a regular basis.

- 1 To open an existing Profile Administration Workbook, select Open from the File menu, select a workbook, click OK, and go to step 5. Or to open a new workbook, select New from the File menu.
- 2 Click the Curve tab.
- 3 Select Profile Administration and click OK.
- 4 Click on a radio button to select the processing plan that you want to edit. Click Finish.
- 5 A worksheet is created for each profile generation level in your selected processing plan. On each worksheet, review each measure and set its contents or value as needed. For descriptions of all measures, see “Profile Administration Workbook” on page 95.
- 6 Commit any changes that you make to the master database by selecting Commit Now from the File menu. If desired, you may also save your workbook by selecting Save from the File menu.
- 7 To close the workbook, select Close from the File menu.

Modify the training window for a subset of items in a profile

Use this procedure to change the training window values for a subset of items in a profile.

- 1 To open an existing Profile Maintenance Workbook, select Open from the File menu, select a workbook, click OK, and go to step 6. Or to open a new workbook, select New from the File menu.
- 2 Click the Curve tab.
- 3 Select Profile Maintenance and click OK.
- 4 Click on a radio button to select the processing plan that you want to edit. Click Next.
- 5 Select the subset of items whose training windows you want to change. Depending on your domain, you may be presented multiple selection screens from which to choose products and/or locations. Use the arrow buttons to move list items between the Available list and the Selected list as needed. When you have made your selections, click the Next button. On the final selection screen, click Finish.
- 6 Make any necessary changes to the training window begin dates or end dates for individual product/location combinations.
- 7 Commit any changes that you make to the master database by selecting Commit Now from the File menu. If desired, you may also save your workbook by selecting Save from the File menu.
- 8 To close the workbook, select Close from the File menu.

Manually request a generated profile

Use this procedure to manually generate a set of Curve profiles, before the next scheduled automatic batch process.

Note: Before performing this procedure, make sure there are no other users logged onto Retek Demand Forecasting. No other user can access Retek Demand Forecasting until the profile generation process is finished.

- 1 From the File menu, select New.
- 2 Click the Curve tab.
- 3 Select Run Batch Profile and click OK.
- 4 Click on a radio button to select the processing plan whose profiles you want to generate.

- 5 Click Finish.

Approve a profile

Use this procedure to view, analyze, revise, and approve the profiles generated at the historic levels. In the approval process, you select the appropriate historic level for each product/location combination and commit the approved profile values to the master database.

- 1 To open an existing Profile Approval Workbook, select Open from the File menu, select a workbook, click OK, and go to step 6. Or to open a new workbook, select New from the File menu.
- 2 Click the Curve tab.
- 3 Select Profile Approval and click OK.
- 4 Click on a radio button to select the processing plan containing the profiles that you want to approve. Click Next.
- 5 Select the items whose profiles you want to analyze and approve. Depending on your domain, you may be presented multiple selection screens from which to choose products and/or locations. Use the arrow buttons to move list items between the Available list and the Selected list as needed. When you have made your selections, click Next. On the final selection screen, click Finish to build the workbook.
- 6 The Historic Level Worksheet displays the historic level profile values and lets you edit those values. Review the values in the Final Profile field and System Profile fields. Pay special attention to any item for which there is a checkmark in the Threshold Exceptions field, indicating that sales are sparse at the item's historic level, and the profile ratio might be unreliable. To change a value in the Final Profile measure, type in the revised value in the corresponding Edit History Level field.

Note: Historic Level Worksheets for various levels may be displayed for your system. These levels are specified at system implementation, depending on users' needs. You can compare and edit the values on any Historic Level Worksheet that is displayed.

- 7 Display the Profile Approval Worksheet to approve the profiles by product and location. If the default approval method in the Profile Administration Workbook is set to Approve, each Approval Method field will already display "Approve", each Approved By field will display the approver's user ID, and each Approval Date will display the proper date.
- 8 If the default approval method in the Profile Administration Workbook is set to Don't Approve, each Approval Method field will display "Don't Approve." The Approved By and Approval Date measures will be unfilled.

- 9 To approve the profile value for an item, click the Approval Method drop-down list and select “Approve” (or leave “Approve” displayed). Your user ID will automatically display in the Approved By field and today’s date will display in the Approval Date field.
- 10 For each item, if you want to change the historic level that you want to use for the profile spread, click the Historic Level field (drop-down list) and select the historic level to use.
- 11 If you want to leave a comment about any item, type your comments in the Approval Comment field.
- 12 From the File menu, select Commit Now to commit the approved profile values.
- 13 To review the profiles at the destination level, display the Final Profile Worksheet, with no values displayed in the fields.
- 14 From the File menu, select Refresh.
- 15 On the Refresh window, click the Refresh All button and click OK. After refreshing, the Final Profile Worksheet will display the final level profile data.
- 16 Review the final profiles and ratio exceptions.
- 17 To save the workbook and changes, select a Save option from the File menu.
- 18 To close the workbook, select Close from the File menu.

Set Retek Demand Forecasting to perform plan explosion

Use this procedure to set up your Retek Demand Forecasting system to perform plan explosion. If you are performing profile-based spreading alone without time series forecasting, perform this procedure only. If you are performing both plan explosion and time series forecasting, you must also make changes to the Forecast Maintenance Workbook, corresponding to the changes made in this procedure, by adjusting forecast parameters.

- 1 To open an existing Forecast Administration Workbook, select Open from the File menu, select a workbook, click OK, and go to step 5. Or to open a new workbook, select New from the File menu.
- 2 Click the Predict tab.
- 3 Select Forecast Administration and click OK.
- 4 Select the level of the final forecast and click Finish.
- 5 Display the Final Level Worksheet.

- 6 In the Default Generation Level field, click the drop-down list and select the source level that loads the assortment plan.
- 7 In the Default Generation Method field, click the drop-down list and select Profile Spread.
- 8 In the Horizon Start Date field, click to display the pop-up calendar. Select a date for the start of your plan explosion period.
- 9 In the Number of Periods field, type the number of periods. Adding this number of weeks to the Horizon Start Date you selected should result in the end date for the plan explosion period.
- 10 Display the Source Level Worksheet.
- 11 In the Default Generation Method field, click the drop-down list and select the Load Plan option.
- 12 Commit the changes to the master database by selecting Commit Now from the File menu. If desired, you may also save your workbook by selecting Save from the File menu.

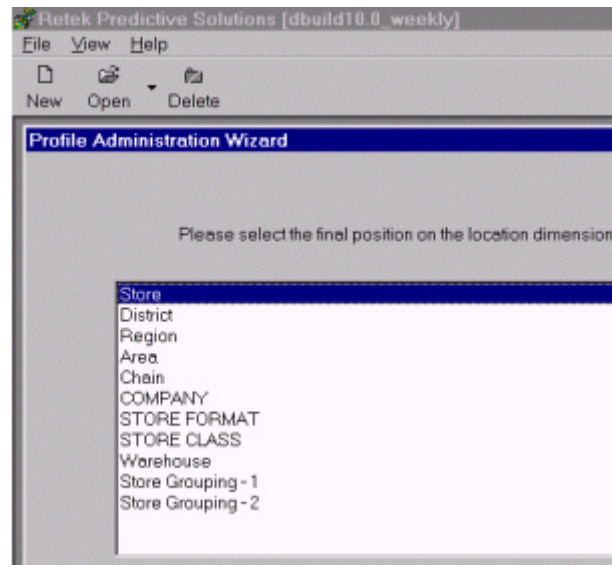
To close the workbook, select Close from the File menu.

Create a new profile template

- 1 Select New from the File menu.
- 2 Click the Curve tab.
- 3 Select Profile Administration and click OK.
- 4 In the first dialog of the Profile Administration Wizard, select New and click Next.
- 5 In the next dialog, specify a name for the new profile and click Next..
- 6 In the next dialog, Select the type of the profile and click Next.
- 7 The next dialogs displayed depend on the type of profile selected in the previous step.

If you chose Store Participation Profile, Color Profile, or Size on the previous dialog, the next two dialogs displayed are for selecting the final and aggregate dimensions on the corresponding hierarchy, that is, the product hierarchy for color size, or the location hierarchy for store participation.

For example, the following positions would be displayed for the Store Participation Profile:



Store Participation Profile – Location hierarchy selections

If you chose Seasonal Profile or Life Cycle Profile, the system automatically sets up a profile with “week” as the time dimension for the final profile, and “all weeks” as the aggregate dimension.

For Daily profiles, the final time dimension is “day” and the aggregate dimension is “all.”

- 8 The next dialogs are dedicated to choosing the rest of the profile dimensions: location for color or size profiles, product for store participation profile, or both product and location for the time profiles.
- 9 The system automatically assigns one source level for the dynamically created profile. The intersection of the source level profiles is automatically set to the profile intersection.

Chapter 5 – Preprocessing

This chapter describes the features and functions available in the RDF Preprocessing module. It covers the following topics.

- An overview of the Preprocessing module
- Descriptions of the Preprocessing Workbook templates and wizards
- Descriptions of the preprocessing methods, including detailed descriptions of the mathematics involved in the methods, and graphical or chart-form examples of the effects of preprocessing methods on forecast data.

Overview

The purpose of the Preprocessing module, which may also be referred to as “Lost sales,” is to correct past data points that represent unusual sales values not representative of a general demand pattern. Such correction may be necessary either when an item is out of stock and cannot be sold, resulting in usually low sales. Conversely, correction of data may also be necessary in a period when demand is unusually high. The Preprocessing module allows you to automatically make adjustments to the raw POS (Point of Sales) data, so that subsequent demand forecasts do not replicate undesired patterns caused by lost sales or unusually high demand.

Preprocessing Workbook templates and wizards

The Preprocessing module has three workbook templates and wizards:

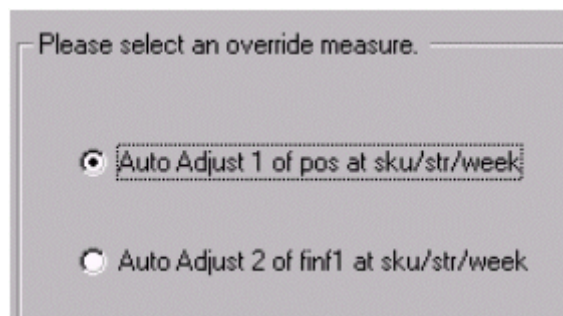
- Administration: For setting up parameters.
- Run Batch: For starting the batch process that preprocesses the raw data
- Run Selection Update: For updating the non-sparse and outage product/location selections.

Administration Workbook

At configuration, preprocessing levels are set up in a similar ways as the forecast levels. Each level is meant to be associated with an override measure (also referred to as source measure). The Administration Workbook enables a user to setup some preprocessing parameters for a given preprocessing level. This workbook consists of a wizard and one worksheet.

Administration Wizard

The Preprocessing Administration Wizard enables you to select the level you want to set up for preprocessing. The following figure shows such a wizard.

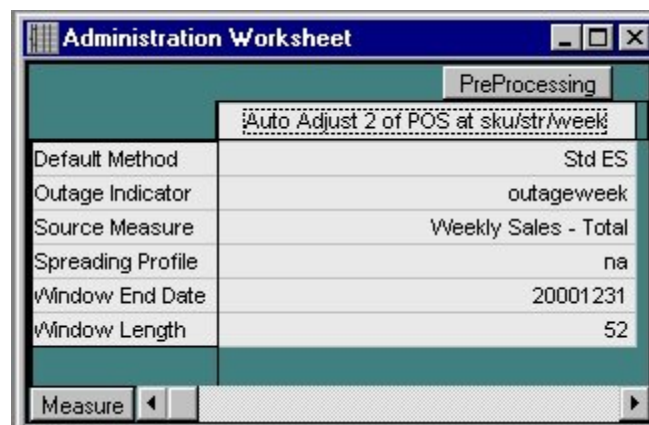


Preprocessing Administration Wizard

After this wizard is completed, the Administration Worksheet is displayed.

Administration Worksheet

The Preprocessing module's Administration Worksheet looks similar to the following:



Preprocessing Administration Worksheet

The worksheet contains the following parameters, which you can adjust as needed.

Default Method

The default generation method. This parameter defines the algorithm to be used by the preprocessing run. The list of available methods is included later in this document, along with descriptions of each method.

Outage Indicator

A measure representing outages or/and any other special events (for example, promotion, bad weather). Your choices here depend on the events that were defined at system configuration.

Source Measure

The measure holding the data to be preprocessed. This module can preprocess the raw point of sales (for example, POS) or a measure resulting from a previous preprocessing run (for example, the override measure POSOVER).

Spreading Profile

The preprocessing calculations may be carried at a higher level (for example, style/district/week) than the intersection of interest (for example, sku/str/week). This field represents a selection of profiles generated by the Curve module that perform the spreading from the calculation level (for example, style/district/week) to the lower intersection of interest (for example, sku/str/week).

Window End Date

This value is meant to be the same as the forecast generation date. It refers to the point following the last point to be adjusted.

Window Length

This value defines how far back in time from the window end date the raw data measure will be preprocessed. The number of points that it represents is in time-units of the base calendar intersection (for example, 'week' if the preprocessing results are reported at the week level).

Run Batch Workbook

The Run Batch template enables you to launch a preprocessing run, using the last committed parameters setup in the Administration Workbook. The following wizards sequentially guide you through setting up the run batch process.

- Selection of preprocessing level (same as in administration WB).
- Two-tree selection wizards to limit the preprocessing run to a subset of product/location items as desired.
- Clear alert wizard. An alert is generated for time-series that are automatically adjusted. You may decide whether to reset that alert. (See Preprocessing Related Measure below for more information on that alert.)
- Confirmation to actually run the batch process on the selection. If you confirm the batch run, the batch run starts and may take a long time depending upon the selection size. For a given data set size, a preprocessing run is significantly faster than a forecast generation run.
- Finally, a wizard alerts you when the batch run has completed.

Run Selection Update

The run selection update template enables you to run a batch routine that creates product/location selections that may in turn be used in the run batch selection process. Two selection sets may be generated by this routine:

- Non-sparse items – This feature is particularly important for the median filters because median filters should not process sparse time-series. An item (either product or location) belongs to this selection if there is at least one non-sparse time-series containing that item (“OR” aggregation of TSMASK_DENSE measure; see Preprocessing Related Measure below for more information on TSMASK_DENSE). Its label is “Non-sparse TS of ...”.
- Items with an outage indicator – “ON” on the time frame defined by the generation date and the number of points in the administration WB. This selection is created if and only if an outage indicator measure exists for the corresponding preprocessing level. An item (either product or location) belongs to this selection if there is at least one time-series containing that item that has an outage (“OR” aggregation of OUTAGE measure, see Preprocessing Related Measure below for more information on OUTAGE). Its label is “Outage of ...”.

The following figure shows a list of product selections that were generated by the run selection update algorithm.

The screenshot shows a software window with two tabs: 'Selected Products' and 'Selections'. The 'Selections' tab is active, displaying a list of product selections. The list has a header 'Name' and contains the following entries:

- Non-sparse TS of finf1 at prbr/ucsh/week (01/10/98 - 01/02/99) -
- Non-sparse TS of pos at prbr/ucsh/week (04/11/98 - 04/03/99) --
- Outage of finf1 at prbr/ucsh/week (01/10/98 - 01/02/99) -- PROD
- Outage of pos at prbr/ucsh/week (04/11/98 - 04/03/99) -- PROD

Below the list is a horizontal scrollbar. Underneath the scrollbar is a checkbox labeled 'Load name only' which is currently unchecked. To the right of the checkbox are two dropdown menus: 'Load in selection list by:' with 'Adding' selected, and 'Save access:' with 'User' selected. At the bottom of the window are three buttons: 'Load', 'Save', and 'Delete'.

Product selections generated by Run Selection Update

Preprocessing methods

The following preprocessing methods are implemented in RDF 9.5. The algorithms that have the prefix Lost Sales in their label / name are specific to lost sales calculations in the two following ways:

- The adjustment override (LSOVER) is always greater than or equal to the raw source data (SRC). That is, positive spikes are not adjusted using lost sales specific algorithms.
- For Lost Sales algorithms taking as an input an outage indicator, the data points that are adjusted are those corresponding to an outage event and those following an outage. Points following an outage indicator are corrected because they very likely correspond to a period of partial outage.

Available preprocessing methods

Following is the list of currently available preprocessing methods, with recommendations and notes on their use. Further details on the mathematics used in these methods are on page 122.

Std Median

- Recommended for getting data baselines on long time ranges when promo indicators are not available.
- A standard median filter implementation.
- Does not take outage info as an input.
- One optional parameter: window length.

Retek Median

- Recommended for getting data baselines on long time ranges when promo indicators are not available.
- A sophisticated median filter that takes into consideration trends and improves side effects over the standard median filter. It makes 5 standard median filter passes.
- Does not take outage info as an input.
- Five optional parameters: window length for each pass.

Seasonal Median 1

- Recommended for getting data baselines on long time ranges when promo indicators are not available. To note, it is recommended to use that method when several historical cycles are available (for example, a year of history for preprocessing at a daily level).
- A median filter type of algorithm that focuses on seasonal patterns.
- Does not take outage info as an input.
- One optional parameter: window length.

Std ES

- Recommended for removing specific spikes of low or non-seasonal data when spike indicators are available.
- Based on standard Exponential Smoothing calculations of past and future sales velocities.
- Must have “unusual event” measure info as input (that is, also referred to as outage).
- Does not take into consideration seasonal components.
- 3 optional parameters:
 - Alpha (ES parameter used to evaluate past and future velocities).
 - Maximum number of historical points to calc past velocity.
 - Maximum number of future points to calc future velocity.

Seasonal ES 1

- Recommended for removing specific spikes of seasonal data when spike indicators are available. To note, it is recommended to use that method when several historical cycles are available (for example., a year of history for preprocessing at a daily level).
- Based on standard Exponential Smoothing calculations of past and future sales velocities taking into consideration seasonal cycles.
- Must have “unusual event” measure info as input (that is, also referred to as outage).
- Does take into consideration seasonal components.

- Three optional parameters:
 - Alpha (ES parameter used to evaluate past and future velocities).
 - Maximum number of historical points to calculate past velocity.
 - Maximum number of future points to calculate future velocity.

Causal Median

- Recommended for promotional data.
- Processes non-promotional and promotional data points independently.
- Does not take outage info as an input.
- Does take a list of causal/promotional events to process independently.
- Optionally takes pre-calculated causal effects in.

Causal ES

- Recommended for promotional data.
- Processes non-promotional and promotional data points independently.
- Must have “unusual event” measure info as input (that is, also referred to as outage).
- Does take a list of causal/promotional events to process independently.
- Optionally takes pre-calculated causal effects in.

Lost Sales – Std Median

- Functions the same as Std Median, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Retek Median

- Functions the same as Retek Median, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Std ES

- Functions the same as Std ES except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Seasonal ES 1

- Functions the same as Seasonal ES 1 except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Seasonal Median 1

- Functions the same as Seasonal Median 1 except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Causal Median

- Functions the same as Causal Median, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Causal ES

- Functions the same as Causal ES, except that it only adjusts lost sales (that is, negative spikes).

Clear

- Used for canceling the effect of some former preprocessing adjustments.
- Does not take outage info as an input.
- No parameter.

No Filtering

- If you select No Filtering, it turns the preprocessing off. This enables you to turn preprocessing off from the RDF client without modifying the batch code script.
- Does not take outage info as an input.
- No parameter.

Mathematical formulation of preprocessing methods

The following discussion provides an in-depth explanation of the mathematics involved behind the preprocessing algorithms.

Notations used in descriptions

The descriptions of preprocessing mathematics use the following special notations:

- $SRC(t)$: Actual source data to preprocess at time index t .
- $LSOVER(t)$: Preprocessing result of SRC at time index t .
- $LS(t)$: Preprocessing adjustment at time index t , that is, $LS(t) = LSOVER(t) - SRC(t)$.

Std Median

Std Median is the standard median filter.

$LSOVER(t) = \text{median value of } SRC \text{ over } [t-\text{window}/2, t+\text{window}/2]$,

Where: “window” is the parameter window length of the filter.

Retek Median

Retek Median is a sophisticated median filter that takes into consideration trends and improves side effects over the standard median filter.

The filter makes five passes.

- The first two passes recursively apply the standard median filter. The result is denoted by $MEDIAN_2(t)$. The one-step difference of $MEDIAN_2(t)$ is calculated, that is, $DIFF_1(t) = MEDIAN_2(t) - MEDIAN_2(t-1)$. Then, the standard median filter is applied to $DIFF_1(t)$. The result is denoted by $MEDIAN_DIFF_1(t)$.
- Using $MEDIAN_DIFF_1(t)$, a first smoothed version (that is, baseline) of the source data is calculated at the third step: $SMOOTH_1(t) = SMOOTH_1(t-1) + MEDIAN_DIFF_1(t)$ on points where the absolute deviation of $SRC(t)$ over its mean is larger than half of the global absolute standard deviation, $SMOOTH_1(t) = SRC(t)$ otherwise.
- To prepare for the fourth pass, the one-step difference of $SMOOTH_1(t)$ is calculated, that is, $DIFF_2(t) = SMOOTH_1(t) - SMOOTH_1(t-1)$. An average version of $DIFF_2(t)$ is calculated using the standard median filter. The result is denoted by $AVG_DIFF_2(t)$. The result of the fourth pass is $SMOOTH_2(t) = SMOOTH_2(t-1) + AVG_DIFF_2(t)$.

- Finally, LSOVER(t) is the result of applying the standard median filter to SMOOTH_2(t).

Seasonal Median 1

Seasonal Median 1 applies the Standard Median filter to sub-series built from data points belonging to the same seasonal relative position (for example., same day of week).

Std ES

Std ES is the standard Exponential Smoothing filter. It preprocesses a subset of points as predetermined by an input measure. For every contiguous sequence of points to adjust, say between t_f and t_l , a past velocity and a future velocity are calculated using an exponentially weighted average. For the points between t_f and t_l , the adjustment is calculated as a linear interpolation of the past and future velocities.

$$Past_Velocity = \frac{\sum_{i=1}^{np} (1-\alpha)^{i-1} * SRC(t_f - i)}{\sum_{i=1}^{np} (1-\alpha)^{i-1}}$$

$$Future_Velocity = \frac{\sum_{i=1}^{nf} (1-\alpha)^{i-1} * SRC(t_l + i)}{\sum_{i=1}^{nf} (1-\alpha)^{i-1}}$$

$$LSOVER(t) = Past_Velocity + \frac{Future_Velocity - Past_Velocity}{t_f - t_l + 2} * (t - t_f + 1), \forall t \in [t_f, t_l]$$

Where:

“ α ” is the exponential coefficient used to evaluate past and future velocities.

“ np ” is the maximum number of historical points to calc past velocity.

“ nf ” is the maximum number of future points to calc future velocity.

Seasonal ES 1

Seasonal ES 1 applies the Standard ES filter to sub-series built from data points belonging to the same seasonal relative position (for example, same day of week).

Causal Median

Causal Median applies the Std Median filter independently for non-causal and causal points.

Causal ES

Causal ES applies the Std ES filter independently for non-causal and causal points.

Lost Sales – Std Median

Functions the same as Std Median, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Retek Median

Functions the same as Retek Median, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Std ES

Functions the same as Std ES, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Seasonal ES 1

Functions the same as Seasonal ES 1, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Seasonal Median 1

Functions the same as Seasonal Median 1, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales – Causal Median

Functions the same as Causal Median, except that it only adjusts lost sales (that is, negative spikes).

Lost Sales -- Causal ES


Functions the same as Causal ES, except that it only adjusts lost sales (that is, negative spikes).

Clear



$$\text{LSOVER}(t) = \text{SRC}(t).$$


Chart-form examples of preprocessing effects

The following examples show the effects of several preprocessing methods on data displayed in chart form.

Std Median

Std Median with “window” = 13 points

Retek Median*Retek Median with default parameters***Std ES***Std ES with " α " = 0.2, " n_p " = 2 weeks, and " n_f " = 2 weeks*

Lost Sales -- Std ES

Lost Sales -- Std ES with “ α ” = 0.2, “np” = 2 weeks, and “nf” = 2 weeks

Chapter 6 – Retek Demand Forecasting methods

This chapter discusses the forecasting methods used in Retek Demand Forecasting in detail.

Forecasting techniques used in RDF

Retek Demand Forecasting uses a variety of predictive techniques to generate forecasts of demand. The technical methods used are driven by the goal to provide the most accurate forecasts possible in an automatic and efficient manner. These methods have been analyzed, optimized, and refined over years of research on retail specific data.

The primary techniques RDF uses include:

- Exponential smoothing
- Regression analysis
- Bayesian analysis
- Prediction intervals
- Automatic method selection
- Source level forecasting
- Promotional forecasting

Exponential smoothing

Exponential smoothing models fit basic features of the demand pattern such as level, trend, and seasonality, and project these into the future. These models provide computational benefits and have been chosen for their ability to handle different types of time series, including short and/or noisy series that are characteristic of retail sales. They are "smoothing" models because they use weighted averages on historic data. They are "exponential smoothing" models because the weighting used decays at an exponential rate. That is, more recent data is weighted more heavily than the past.

Regression analysis

Regression analysis is another standard technique used in prediction. Regression uses a least-squares estimator to fit a model of predictor variables to another set of target variables. Seasonal Regression is a Retek specific extension of this procedure for use in seasonal models with between one and two years of history. Causal Forecasting uses stepwise regression to determine which causal variables are significant.

Bayesian analysis

Bayesian analysis considers a priori information as a starting point in development of a prediction. Bayesian forecasting, as developed by Retek, uses a sales plan as this starting point which is adjusted based on observed data. This method fills a gap in standard time series forecasting when new, short life cycle, or products with significant life cycles are being forecast.

For more information on Bayesian forecasting, see page 152.

Prediction intervals

Prediction from these various models gives the estimated mean outcome. By using standard statistical distributional assumptions, RDF develops measures of uncertainty associated with forecast point estimates from these models. While this is of key concern for various optimization applications of the forecast, the technical details are beyond the scope of this document. For further details on prediction interval calculations, see Char&Yatfield, *International Journal of Forecasting*, March 1992.

Automatic method selection

Providing multiple forecasting methods is only valuable if the appropriate model can be selected in an accurate and efficient manner. In order to make this feasible in a retail environment Retek has developed a number of different Meta-Methods that can automatically select the best method among a number of competing models. Automatic Exponential Smoothing (AutoES) is an example of one such method that clients can select. The final selection between the competing models is made according to a performance criterion that involves a tradeoff between the model's fit over the historic data and its complexity (a description of the competing models used within AutoES is described in section two of this document).. In academia this discipline is known as *Information Theory* and is used in the combination and selection of various competing models.

Source level forecasting

Sometimes it is difficult to capture seasonality, trend, or causal effects on the final level (SKU/Store) due to scarcity of the data. Also, time series are often too noisy at that level. To overcome these issues RDF utilizes source level forecasting. In source level forecasting, data is aggregated first to a higher level across the product or location hierarchy (or both). Then the forecast is generated and proportionally spread down to the final level. We have experimentally proven that source level forecasting technique very often improves the accuracy on the final level.

Promotional forecasting

In some instances, especially in retail, pure time series techniques are inadequate for forecasting demand. Instead of using only historic demand patterns to forecast future demand, additional causal or promotional factors are used to better explain past performance. With the help of a promotional calendar, an indication of when promotions will be run in the future, these promotional forecasting techniques can better predict demand in the future.

For more information on promotional forecasting methods, see page 155.

Time series (statistical) forecasting methods

This section describes those techniques within RDF which generate forecasts directly from only a single time series. Generally the time series provided is past sales history for a given SKU/Store that is used to predict what future demand might be. In actual practice these algorithms have been and can be used to forecast a myriad of different data streams at any product/location level (shipment data at item/warehouse, financial data at dept./chain, etc.).

The following topics present fundamentals of the Retek Demand Forecasting statistical forecasting processes. Included is a discussion of the importance of confidence intervals and confidence limits, the time series methods used to generate forecasts, and how the best forecasting method is selected from a list of candidate models.

A wide variety of statistical forecasting techniques are available, ranging from very simple to very sophisticated. All of these methods attempt to best capture the statistical probability distribution discussed above, and they do this by fitting quantitative models to statistical patterns from historical data. Put simply, the better the history of the variable being forecasted, the stronger these statistical patterns will be. Increased forecast accuracy depends on the strength of these patterns in relation to background irregularities.

Retek Demand Forecasting is able to use several time series methods to produce forecasts. Time series methods extrapolate features from the past (in this case, past sales data) to the future. The time series methods that the system offers include:

- Auto Exponential Smoothing Forecasting (AutoES)
- Seasonal Exponential Smoothing Forecasting (SeasonalES)
- Simple Moving Average
- Simple Exponential Smoothing
- Croston's Method
- Holt Exponential Smoothing
- Multiplicative Winters Exponential Smoothing
- Additive Winters Exponential Smoothing
- Seasonal Regression
- Bayesian Information Criterion

Why use statistical forecasting?

The purpose of statistical forecasting is to make the process of predicting future events both objective and quantitative. Statistical forecasting utilizes information from the past (such as sales data) to predict what will happen in the future. Forecast accuracy depends on the degree to which a mathematical model can detect and extract statistical patterns from historic data. The most common statistical methodologies used are univariate. This means that they are based solely on the history of one variable, such as sales. Each forecast observation reflects a future value of the sole input variable. Statistical forecasting processes are relatively easy to implement, and the better the historical data, the better the resulting forecasts.

Businesses benefit greatly from the use of systematic statistical forecasting techniques that aim to accurately predict product demand, enabling these businesses to maintain sufficient product inventory levels. When inventory levels are optimized, lost sales due to product stock-outs are greatly reduced, as are the costs incurred by overstocking.

Confidence limits

Uncertainty in statistical forecasting is represented as a probability distribution. A probability distribution uses two components: a point forecast, or the best estimate of a future value; and a confidence interval, which describes the distribution spread above and below the point forecast. Confidence limits are the endpoints of a confidence interval. Confidence intervals are typically represented as percentages; a 90% confidence interval indicates that the actual value of the forecasted event should fall inside the interval 90% of the time. The forecast value is centered between the upper and lower confidence limits; these limits describe the spread of the probability distribution above and below the forecast.

A practical example of upper and lower confidence limits and their relationship to the actual forecast is a weather forecast. Even if a meteorologist predicts that tomorrow's temperature at noon will be 80 degrees, it is unlikely that the temperature tomorrow will be exactly 80 degrees. However, it is very likely (for example, a 90% chance) that the temperature will fall somewhere between 70 and 90 degrees. Additionally, it is even more likely (for example, a 99% chance) that the temperature will fall between 40 and 120 degrees. Therefore, for this weather forecast, the 90% confidence limits are 70 and 90, while the 99% confidence limits are 40 and 120.

Now, suppose you move to San Diego (a climate that sees far less variations in daily temperatures than the original location), and tomorrow's forecast is still 80 degrees. In the new climate, the 90% confidence limits may be 78 and 82 degrees, and the 99% confidence limits may be 76 and 84 degrees. Because the San Diego confidence limits are closer to the forecast, you are much more confident in the forecast's accuracy. Note that changing the confidence interval from 90% to 99% does not affect the value of the forecast; the forecast for tomorrow's temperature is still 80 degrees.

Auto Exponential Smoothing (AutoES) forecasting

The primary process by which RDF automatically fits an exponential smoothing model to a time series is called Automatic Exponential Smoothing (*AutoES*). When *AutoES* forecasting is chosen in RDF, a collection of candidate models is initially considered. The models in the candidate list include

- Simple (One Parameter) Exponential Smoothing
- Croston's Method (Intermittent ES)
- Holt (Two Parameter) Exponential Smoothing (Trend ES)
- Winters (Three Parameter) Exponential Smoothing (Seasonal ES)
- Seasonal Regression

These models include level information, level and trend information, and level, trend and seasonality information, respectively. The optimal smoothing parameters for each model form are determined automatically (that is, greater smoothing will be applied to noisier data). The final selection between the resulting models is made according to a performance criterion that involves a tradeoff between the model's fit over the historic data and its complexity.

The amount of available historic information can affect the complexity of the model that can be fit. For example, fitting a seasonal model would not be appropriate without a complete year of historic data. In fact, one prefers to see each "season" occur multiple times. For a particular series, even if the amount of available history allows one to fit a complex model (that is, one with seasonal components), the resulting model is not necessarily superior to a simpler model. If a simpler model (for example, a model with only a level component, or level and trend components) fits "as well" as a seasonal model, then the *AutoES* forecasting process will find the simpler model to be preferable. In such a case, the simpler model captures the basic features supported by the data without over fitting and therefore will generally project better forecasts.

Definitions of equation notation used in this section

The following notation is used in equations throughout this section.

- Y_t - Observed value at time t .
- h - Number of periods ahead for which to forecast.
- L_t - Smoothed level component at end of time t .
- T_t - Smoothed trend component at end of time t .
- S_t - Smoothed seasonal index component at end of time t .
- α - Smoothing parameter for level of series.
- γ - Smoothing parameter for trend.
- δ - Smoothing parameter for seasonal indexes.
- ϕ - Damped trend constant.
- p - Number of periods per year.
- $\hat{Y}_t(h)$ - Forecast for time $t+h$ from base t .
- \tilde{S}_{t+h} - Most recent seasonal index for time $t+h$

Components of exponential smoothing

Several features of the data can be incorporated into exponential smoothing models using the following structural components:

The *level* (L_t) of the series estimates the non-seasonal, slowly changing process of the time series. This data feature represents a flat value for the data when noise, trend and seasonality effects are excluded from the data.

The *trend* (T_t) of a series reflects the rate of change of the series from one point to another. RDF uses a damping factor with its trend component to curtail exaggerated effects over long forecast horizons.

Smoothed seasonal indexes (S_t) estimate periodic patterns in the demand data. Using these estimated values, a model can extrapolate the seasonal features of a series. A seasonal index can be derived as an additive or multiplicative feature of the data.

For a particular time series, a more complex model might consider seasonality and/or trend. For all models, a form of the level of the series is used.

Simple moving average

A simple moving average forecast involves taking the average of the past n time periods and using that average as the forecast for all future time periods (where n is the length of fitting period). Simple moving average forecasts are frequently used in the system because they make few assumptions about the historical time series, they can be generated with little historical data, and because they are very fast to generate. Typically, moving average forecasts are generated at the final forecast level (for example, SKU/Store) and their results used to spread more sophisticated higher-level forecasts (for example, those generated with exponential smoothing).

A Simple Moving Average model assumes that historical data is too short or noisy to consider seasonal effects or local trend, and is based on the level of the series. Since this model does not use a smoothing parameter to place added weight on more recent historic values, a Simple Moving Average model is not actually in the exponential smoothing family. However, it is an adequate model to use when low-level (final forecast) ratios are needed for RDF's spreading of high-level (aggregate) forecasts. That is, when aggregate forecasts can be calculated for long and less noisy aggregate time series, Simple Moving Average models provide an adequate (and computationally quick) forecast to determine the ratios needed for RDF spreading. User input in overriding the automatic training horizon further enhances the simple robustness of this model for base-level data.

Overall, a forecast is evaluated as:

$$\hat{Y}_t(h) = \frac{1}{n} \sum_{k=0}^{n-1} Y_{t-k}$$

Confidence intervals constraining in Average:

The confidence interval is capped using the following rule:

if $\text{interval}(0) = \text{stdev}(\text{frst error}) < 0.3 \cdot \text{level}$, then $\text{interval}(i) = 0.3 \cdot \text{level} \cdot \sqrt{(i+1)}$

$$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$$

if $\text{interval}(0) = \text{stdev}(\text{frst error}) > \text{level}$, then $\text{interval}(i) = \text{level} \cdot \sqrt{(i+1)}$

$$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$$

Otherwise, $\text{interval}(i) = \text{interval}(0) \cdot \sqrt{(i+1)}$

$$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$$

Simple exponential smoothing

Simple Exponential Smoothing does not consider seasonality or trend features in the demand data (if they exist). It is the simplest model of the exponential smoothing family, yet still adequate for many types of RDF demand data. Forecasts for short horizons can be estimated with Simple Exponential Smoothing when less than a year of historic demand data is available and acts-like associations are not assigned in RDF.


In its recursive form, the Simple Exponential Smoothing equation is:

$$L_t = \alpha Y_t + (1 - \alpha) L_{t-1}$$

A forecast point estimate is evaluated as:

$$\hat{Y}_t(h) = L_t$$

The following figure is an example of a forecast whose data seems to be untrended and unseasonal; note the flat appearance of the forecast.



Simple exponential smoothing

Confidence intervals constraining in SimpleES:

The confidence interval is capped using the following rule:

if $\text{interval}(0) = \text{stdev}(\text{frst error}) < 0.3 \cdot \text{level}$, then $\text{interval}(i) = 0.3 \cdot \text{level} \cdot \sqrt{(i+1)}$

$$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$$

if $\text{interval}(0) = \text{stdev}(\text{frst error}) > \text{level}$, then $\text{interval}(i) = \text{level} \cdot \sqrt{(i+1)}$


$$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$$

Otherwise, the intervals and cumints are extended in horizon using the ES formula for intervals from SimpleES (¹)

Croston's method

Croston's method is used when the input series contains a large number of zero data points (that is, intermittent demand data). The method involves splitting the original time series into two new series: (i) magnitude series and (ii) frequency series. The magnitude series contains all the non-zero data points, while the frequency series consists of the time intervals between consecutive non-zero data points. A Simple Exponential Smoothing model is then applied to each of these newly created series to forecast a magnitude level as well as a frequency level. The ratio of the magnitude estimate over the frequency estimate is the forecast level reported for the original series.

The following figure shows a sales history of data where the demand for a given period is often zero.



Croston's method

Simple/Intermittent Exponential Smoothing

This method is a combination of the Simple ES and Croston's (Intermittent ES) methods. The Simple ES model is applied to the time series unless a large number of zero data points are present, in which case the Croston's model is applied.

Holt exponential smoothing

Holt exponential smoothing treats data as linearly trended but non-seasonal. The Holt model provides forecast point estimates by combining an estimated trend (for the forecast horizon - h) and the smoothed level at the end of the series. RDF uses a damped Holt model that decays the trend component so that it disappears over the first few weeks. This improves forecasts created using Holt over longer forecast horizons.

Overall, a forecast is evaluated as:

$$\hat{Y}_t(h) = L_t + \left[\sum_{i=1}^h \varphi^i \right] T_t$$


where the Level at the end of the series (time t) is:

$$L_t = \alpha Y_t + (1 - \alpha)(L_{t-1} + T_{t-1}), \text{ and}$$

the Trend at the end of the series (time t) is:

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}$$

When this forecasting method is selected, the forecasts are seen as trending either up or down, as in the following example:



Holt exponential smoothing

Multiplicative Winters exponential smoothing

Winters exponential smoothing treats data as both trended and seasonal.

From sufficient data, RDF extracts seasonal indexes which are assumed to have multiplicative effects on the deseasonalized series. Note that the component describing the deseasonalized values (which is multiplied by the seasonal index $\hat{S}_t(h)$) is just the Holt model described above. In this case, three parameters are used to control smoothing of the components of level, trend and seasonality.

Overall, a forecast point estimate is evaluated as:

$$\hat{Y}_t(h) = \left(L_t + \left[\sum_{i=1}^h \varphi^i \right] T_t \right) \hat{S}_t(h),$$

where the Level at the end of the series (time t) is:

$$L_t = \alpha \frac{Y_t}{S_{t-p}} + (1 - \alpha)(L_{t-1} + T_{t-1}), \text{ and}$$


the Trend at the end of the series (time t) is:

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}, \text{ and}$$

the Seasonal Index for the time series (applied to the forecast horizon) is:

$$S_t = \delta \frac{Y_t}{L_t} + (1 - \delta)S_{t-p}$$

When this forecasting method is selected, the forecasts tend to look “squiggly”, as shown in the following figure:



Multiplicative Winters exponential smoothing

Additive Winters exponential smoothing

Additive Winters is similar to Multiplicative Winters, but the deseasonalized values are adjusted by adding the seasonal index (for the forecast horizon). Three parameters are evaluated to control the smoothing of components of level, trend and seasonality.

Overall, a forecast point is evaluated as:

$$\hat{Y}_t(h) = L_t + \left[\sum_{i=1}^h \varphi^i \right] T_t + \hat{S}_t(h)$$

where the Level at the end of the series (time t) is:

$$L_t = \alpha(Y_t - S_{t-p}) + (1 - \alpha)(L_{t-1} + T_{t-1}), \text{ and}$$

the Trend at the end of the series (time t) is:

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}, \text{ and}$$

the Seasonal Index for the time series (applied to the forecast horizon) is:

$$S_t = \delta(Y_t - L_t) + (1 - \delta)S_{t-p}$$

Confidence intervals constraining in HoltES, AWinters, MWinters:

The confidence interval is capped using the following rule:

$$\text{if } \text{interval}(0) < 0.3 \cdot \text{forecast}(i), \text{ then } \text{interval}(i) = 0.3 \cdot \text{forecast}(i) \cdot \sqrt{(i+1)}$$

$$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$$

if $\text{interval}(0) > \text{forecast}(i)$, then $\text{interval}(i) = \text{forecast}(i) \cdot \sqrt{(i+1)}$

$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$ from the first constrained interval until the end of the forecast horizon

Otherwise, the intervals and cumints are extended in horizon using the interval ES formulas (i, ii, iii)

Seasonal Exponential Smoothing (SeasonalES)

In certain instances, it is known that seasonal models will generally outperform non-seasonal models in forecast accuracy. When this is true (generally used for known seasonal items or for forecasts with a long horizon), it is advantageous to prevent AutoES from selecting from the Simple/Holt/Croston's methods. Choosing SeasonalES does just this. In instances where there exists too little data to create a seasonal forecast (less than 52 weeks of history), the SeasonalES method will cascade to select among Simple/Holt/Croston. In all other instances, a seasonal model (Winter's or Seasonal Regression) will be chosen. See the following section for more information on Seasonal Regression.

Confidence intervals constraining in Seasonal method:

The confidence interval is constrained using the following rule:

if $\text{interval}(0) = \text{stdev}(\text{first error}) < 0.3 \cdot \text{forecast}(i)$, then $\text{interval}(i) = 0.3 \cdot \text{forecast}(i) \cdot \sqrt{(i+1)}$

if $\text{interval}(0) > \text{forecast}(i)$, then $\text{interval}(i) = \text{forecast}(i) \cdot \sqrt{(i+1)}$

Otherwise, $\text{interval}(i) = \text{interval}(0) \cdot \sqrt{(i+1)}$

$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$

Seasonal Regression

A common benchmark in seasonal forecasting methods is sales last year. A sales last year forecast is based entirely on sales from the same time period of last year. Forecasting using only sales last year involves simple calculations and often outperforms other more sophisticated seasonal forecasting models. This method performs best when dealing with highly seasonal sales data with a relatively short sales history.

The seasonal models used in earlier releases of Retek Demand Forecasting (Additive and Multiplicative Winters) were designed to determine seasonality. However, they were not designed to work with sales histories of shorter than two years. Because sales histories of longer than two years are often difficult to obtain, many retail environments need a seasonal forecast that can accommodate sales data histories of between one and two years. In addition, the Additive and Multiplicative Winters models search for short-term trends and have difficulties with trends occurring inside the seasonal indices themselves. The current Retek Demand Forecasting Seasonal Regression forecasting model is designed to address these needs.

The Seasonal Regression Model uses simple linear regression with last year's data as the predictor variable and this year's sales as the target variable. The system determines the multiplicative and additive weights that best fit the data on hand. When optimizing the Seasonal Regression Model, the sales last year forecast is inherently considered, and will automatically be used if it is the model that best fits the data. If there have been significant shifts in the level of sales from one year to the next, the model will learn that shift and appropriately weight last year's data (keeping the same shape from last year, but adjusting its scale).

As with other seasonal models, you can forecast demand for products with insufficient sales histories using this method if:

- You paste in fake history as needed, providing a seasonal profile for the previous year.
- You also forecast for a source level (with the same seasonality profile as the forecasted item and with more than one year of history), using seasonal regression and spread these forecast results down to the member products.

The Seasonal Regression Model is included in the AutoES family of forecasting models, and is thus a candidate model that will be selected if it best fits the data.

More formally, for a series \mathcal{X} of length N , define $x(j)$ as the value of \mathcal{X} at time j , $1 \leq j \leq N$, and let p be the length of the seasonal cycle in series \mathcal{X} , $1 \leq p \leq N$, which is assumed to be known apriori. The forecast for the series at time $l + h$ is represented by $\hat{x}_l(h)$ where l , $1 \leq l \leq N$, is the location in the series at which forecasting begins and h is the position in the forecast window. Using this terminology the forecast is defined as:

$$\hat{x}_l(h) \equiv A \cdot x(l + h - p) + B,$$

where A and B are the constants calculated by linear regression.

Plotting the series where one axis is the value of the series at time t (that is, $x(t)$) and the other axis is the value of the series at time $t - p$ (that is, $x(t - p)$) yields a scatter plot. As can be determined by visual inspection the series plotted against itself shifted by a single cycle falls along a rough line. It is this line that is captured by the linear regression where A is the slope of the line and B is the point of intersection.

Thus, this method captures the trend of a series through the slope of the regression line while the series shifted by a cycle provides its seasonal profile. Using this method the resulting forecast for the original series is calculated. The regression method provides a much better forecast of the series than was possible with the other exponential smoothing algorithms.

Based on the assumptions of the model that this method is trying to describe, versus the noisy data it is likely to receive, several exceptions to this regression technique are caught and corrected. First, since it is logically impossible to receive a negative value for the slope (such a value suggesting an inverse seasonality) any time a negative slope is detected the regression is rerun with the intercept fixed to zero. This guarantees that a positive slope will be calculated and thus a more logical forecast will be given.

The second noise-driven concession is to check the slope to determine if it is either too slight or too great. If this is the case, then the method rejects itself out of hand and allows one of the other competing methods to provide the forecast.

Thus, the full forecasting algorithm is:

- 1 For the series X of length N with seasonality occurring with cycle length p form the pairs $X = (x(1), x(1 + p)), \dots, (x(N - p), x(N))$.
- 2 Perform regression on the pairs in X to retrieve the slope, A , and the intercept, B , using the equation above.
- 3 Check the slope of the regression. If the slope is negative then rerun the regression on X fixing the intercept to zero.
- 4 Check the slope of the regression to determine if it is in the accepted range. If not this method fails to produce a forecast.
- 5 Build the forecast for the horizon of length l as

$$\hat{x}_N(h) = A \cdot x(N - p + h) + B, \quad 1 \leq h \leq l.$$

Bayesian Information Criterion (BIC)

Within AutoES, the model that minimizes the Bayesian Information Criterion (BIC) is selected as the final model. The BIC criterion attempts to balance model complexity with goodness-of-fit over the historical data. The BIC criterion rewards a model for goodness-of-fit and penalizes a model for its complexity. The complexity penalty is necessary to avoid over fitting.

There are various equivalent versions of the Bayesian Information Criterion, but RDF minimizes the following:

$$BIC = s \cdot n^{k/2n}$$

where n is the number of periods in the available data history, k is the number of parameters to be estimated in the model (a measure of model complexity), and s is the root mean squared error computed with one-step-ahead forecast errors resulting from the fitted model (a measure of goodness-of-fit). Note that since each member of the model candidate list is actually a family of models, an optimization routine to select optimal smoothing parameters is required to minimize s for each model form (that is, to select the best model).

Within RDF, a few modifications to the standard selection criteria have been made. These include reducing the number of parameters the Winter's model is penalized by discounting seasonal indices that have little impact on the forecast (multiplicative indices close to 1, additive indices close to 0). These changes tend to favor the seasonal models to a slightly higher degree which improves the forecasts on retail data, especially for longer forecast horizons.

AutoES Flowchart

The following outlines the processing routine the system runs through to evaluate each time series set to forecast using the AutoES method. Please see Chapter 2: Forecast Administration/Advance Settings for more information on adjusting the parameters used to qualify a time series for Croston's, Holt or Winters methods:

Step 1: Filter all leading zeros in the input data that is within the training window. Go to Step 2.

Step 2: Does the time series contain the minimum data points to qualify to forecast using the Croston's method? If yes, generate the forecast and statistics using the Croston's method, and move on to the next time series. If no, move on to Step 3.

Step 3: Does the time series contain enough relevant data to generate a forecast? If yes, generate a forecast and statistics using the Simple ES method and move on to Step 4. If no, do not forecast, then go to the next time series.

Step 4: Does the time series contain the minimum data points to qualify to forecast using the Holt method? If yes, generate a forecast and statistics using the Holt method and move on to Step 5. If no, move on to Step 5.

Step 5: Does the time series contain more than 52 weeks of input data? If yes, generate a forecast and statistics using the Seasonal Regression method and move on to Step 6. If no, move on to Step 9.

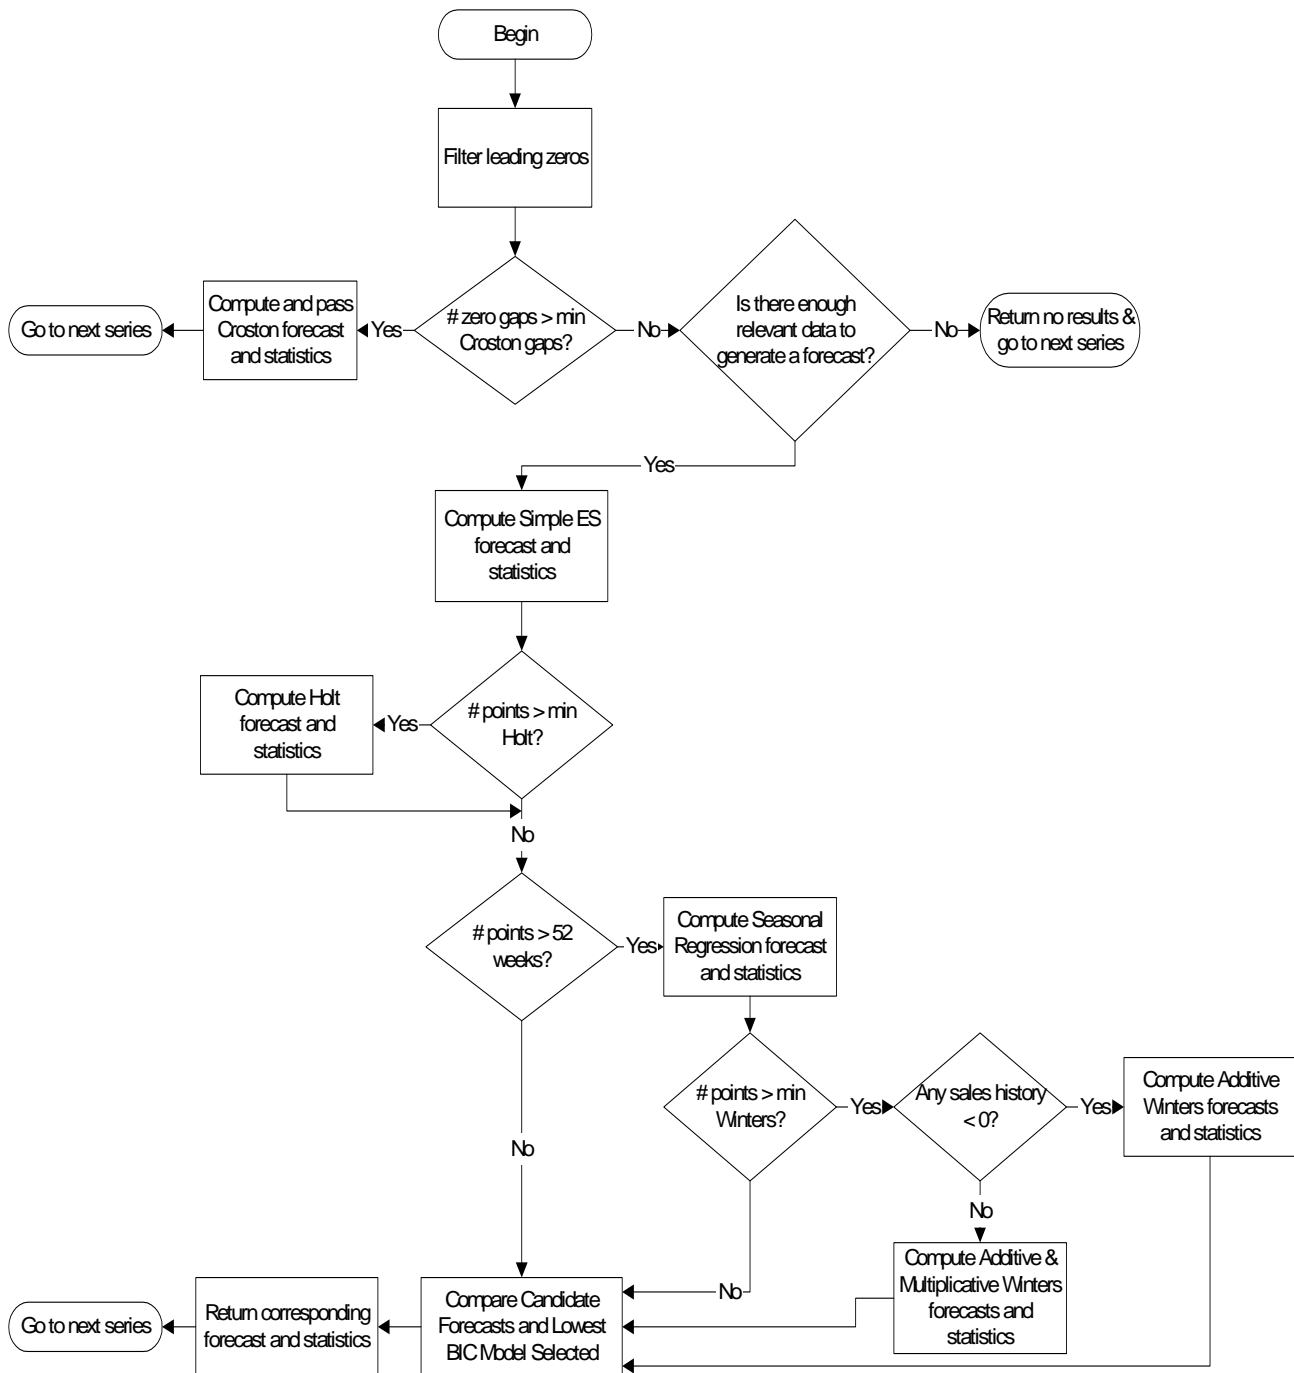
Step 6: Does the time series contain the minimum data points to qualify to forecast using Winters methods? If yes, move on to Step 7. If no, move on to Step 9.

Step 7: Does the time series contain any data point with sales equal qualify to forecast using Additive Winters method? If yes, generate the forecast and statistics using the Additive Winters method, and move on to Step 9. If no, move on to Step 8.

Step 8: Does the time series qualify to forecast using the Multiplicative Winters method? If yes, generate the forecast and statistics using both the Additive Winters and Multiplicative Winters methods, and move on to Step 9.

Step 9: Compare all candidate forecasts using BIC Criterion.

Step 10: Return the corresponding forecast and statistics for the system-selected forecast method, and move on to the next time series.



Automatic forecast level selection (Auto-Select)

This section describes how the automatic forecast level selection (Auto-Select) could help improve the accuracy of your forecasts.

In the system, one of the key elements to producing accurate forecasts is using the system's ability to aggregate and spread sales data and forecasts across the product and location hierarchies. Low selling or relatively new products can use aggregated data from similar products/locations at a higher level in the hierarchy, generate forecasts using this data, and then spread these higher level forecasts back down to provide more accurate forecasts. The difficulty comes in deciding which products/locations will benefit from this technique, and from what level in the hierarchy these source level forecasts should be spread.

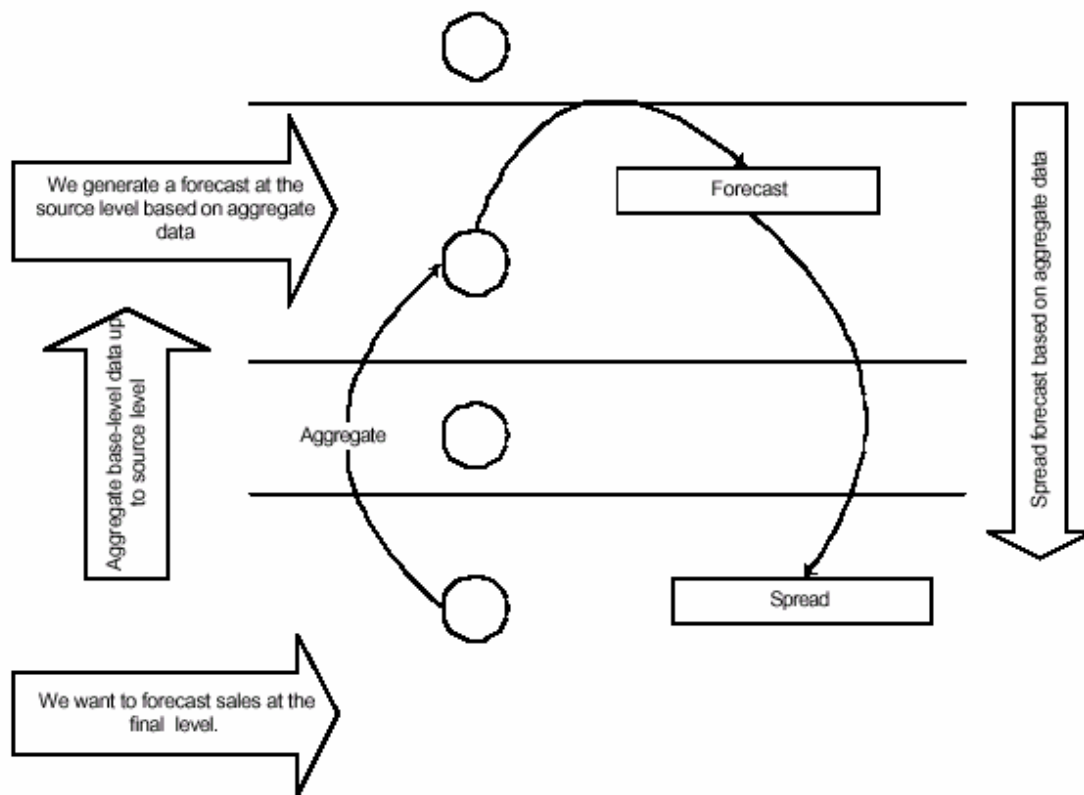
The Automatic Forecast Level Selection feature of the system automates the selection of best aggregation level (forecast source level) for each product/location combination. While providing invaluable information regarding the best aggregate level for source forecasts, the Automatic Forecast Level Selection process may be very CPU intensive. To solve this problem, the task of selecting best aggregation levels for product/location combinations is decomposed and processed piecemeal during times when the computer would normally be idle. Identifying the best aggregation levels for sets of products and locations can be divided into a number of sub-problems:

- Forecasting
- Determining the best source level forecast
- Status and Scheduling

The forecast level selection process

The automatic source generation level selection subsystem selects the best source generation level for each product/location in a given final forecast level. In order to determine the best level, a final forecast is generated for each product/location using each candidate source generation level. As illustrated in the figure below, a final forecast is generated by 1) aggregating up from the base level to the source level, 2) generating a source level forecast, and 3) spreading the source-level forecast down to the final level.

For example, assume base-level sales data is at the SKU/STORE level, the final forecast level is at the SKU/STORE level, and the candidate source generation level is at the STYLE/STORE level. Then, base-level sales data is aggregated from the SKU/STORE level up to the STYLE/STORE level. A STYLE/STORE forecast is generated, and the forecast data is then spread back down to the SKU/STORE level. This forecast represents the final forecast.



Final forecast generation flow chart

Determining the best source level forecast

The selection of the best level is based on a train-test approach. In this process, historical data is used to generate a forecast for a test period for which actual sales data already exists. The forecast, generated over the train period, can be compared to the actual sales figures in the test period to calculate the percent absolute error (PAE) between the two.

A final-level forecast is generated for each product/location combination using each potential source generation level. Each time a source level forecast is generated, a PAE is calculated for that level. If that PAE is better than the current best PAE (corresponding to the current best source generation level), then the source generation level that generated that better PAE becomes the new best level.

Status & scheduling

Identifying the best aggregation level for a given set of products and locations may take a significant amount of time (that is, an amount of time that is greater than the duration of the computer's shortest idle period). This task, however, can be partitioned; that is, the problem of selecting the best aggregation levels can be decomposed into smaller sub problems to be solved piecemeal during times when the computer would normally be idle.

For each product/location combination at the final forecast level, the problem consists of 1) generating forecasts at each unique aggregation level and 2) using the train-test approach to evaluate the percent absolute error statistics for each. One or more of these subtasks will be performed during each period that the computer is idle. The best aggregation status keeps track of which sub problems have been performed and which sub problems remain. In this way, when the best aggregation procedure is run, the procedure knows what the next sub problem is.

Best aggregation level procedures are run during idle computer periods. The scheduling of the Automatic Forecast Level Selection process must be integrated with the schedules of other machine processes. In general, you should select a schedule so that source generation level selection does not conflict with other activities. The following is an example of a typical schedule for the Automatic Forecast Level Selection process: Monday through Thursday, the selection process starts at midnight and runs for 8 hours; on Friday and Saturday, the process is allowed to run for 20 hours; Sunday is reserved for generating forecasts.

Using the system-selected forecast level

You have the option of accepting the system-generated source level selection or manually selecting a different source level to be used. The value for the source forecast level can be manipulated in the Final Level Worksheet of the Forecast Maintenance Workbook. For each product/location combination, the best source forecast level identified by Retek Demand Forecasting will appear in the Suggested Generation Level measure on this worksheet. You can enable the use of this level by placing a checkmark in the Use Suggested Generation Level measure for that product/location. The absence of a checkmark in this measure causes the system to select the source level chosen in the Generation Level measure, which you can specify using a drop-down list.

Profile-based forecasting

The Profile-based forecasting method generates a forecast based on a seasonal profile. The profile may be loaded, manually entered, or generated by Curve. It can also be copied from another profile and adjusted.

Forecast method

The Profile-based forecasting method proceeds as follows:

- 1 The historical data and the profile are loaded.
- 2 The data are desesonalized using the profile and then fed to Simple method.
- 3 The alpha is capped by 0.5 (this number will be editable from a workbook in RDF 9.5).
- 4 The Simple forecast is then reseasonalized using the profiles.

Confidence intervals constraining in Profile Based method:

In Profile Based method, the intervals are taken from the baseline SimpleES forecast. However, the Profile Based forecasts are obtained by multiplying the baseline forecasts with the profiles, thus a new interval constraining is required after the multiplication. The confidence interval is constrained using the following rule:

if $\text{interval}(0) = \text{stdev}(\text{frst error}) < 0.3 \cdot \text{forecast}(i)$ (= Simple forecast * profile),
then $\text{interval}(i) = 0.3 \cdot \text{forecast}(i) \cdot \sqrt{(i+1)}$

if $\text{interval}(0) > \text{forecast}(i)$, then $\text{interval}(i) = \text{forecast}(i) \cdot \sqrt{(i+1)}$

Otherwise, $\text{interval}(i) = \text{interval}(0) \cdot \sqrt{(i+1)}$

$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$

Profile based method and new SKUs

The Profile-based forecasting method can be successfully used to forecast new SKUs. In order to do that we need to have a profile (which can be copied from a SKU that shares the same seasonality), and a number that specifies the deseasonalized demand (DD value). The forecast is calculated using the DD value multiplied by the profile. The confidence interval is set to 1/3 of the DD value.

If the DD value is used to forecast, the history (if exists) of the product is ignored. Once we have enough history (number of data points exceed a global parameter), the forecast stops using the DD value and it defaults to the “normal” Profile Based method.


Example

Consider a product, sunglasses, with monthly data from January 1999 through March 2000.

In grid form, the profile-based forecast for this product is:

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Profile	1.1	0.8	0.7	0.5	0.8	1.1	1.3	1.1	0.8	0.9	1.2	1.7	1.1	0.8	0.7	0.5	0.8	1.1	1.3	1.1	0.8	0.9	1.2	1.7
Historical Data	89	99	49	40	60	96	91	84	80	72	87	137	77	67	98									
Deseas. Data	80	73.75	70	80	75	78.18	70	76.36	75	80	72.5	90.59	70	80.75	80									
Simple Forecast																	76.9	76.9	76.9	76.9	76.9	76.9	76.9	76.9
Re seas. Forecast																	37.88	80.8	83.33	99.48	83.33	60.8	68.18	90.91

Here is the same forecast displayed in chart form:



Bayesian forecasting

The Bayesian Forecasting method is based on combining historic sales data with sales plan data. It is especially effective for new products with little or no historic sales data.

Your sales plan can incorporate expert knowledge in two ways, shape and scale. Shape is the selling profile or life cycle that can be derived from a sales plan. For example, the shape for certain fashion items might show sales ramping up quickly for the first four weeks and then trailing off to nothing over the next eight weeks. The scale, or magnitude, of a sales plan is the total quantity expected to be sold over the plan's duration.

Bayesian Forecasting assumes that the shape that sales will take is known, but the scale is uncertain. In Bayesian Forecasting, when no sales history is available, the sales forecast figures are equal to the sales plan figures (at this point, there is no reason to mistrust the sales plan). As point of sale data becomes available, the forecast is adjusted and the scale becomes a weighted average between the initial plan's scale and the scale reflected by known sales history. Confidence in the sales plan is controlled by the amount of sales data on hand and a Bayesian sensitivity constant (which you can set between zero and one).

Unlike standard time series forecasting, which requires only sales history to produce a forecast, Bayesian Forecasting requires a sales plan and sales history (if available). Because of this difference, Bayesian Forecasting is not included in AutoES. You must select it manually as a forecasting method, in Forecast Administration or Forecast Maintenance.

Obtaining accurate short life-cycle product forecasts is very difficult, and standard statistical time series forecasting models frequently do not offer an adequate solution for many retailers. Major problems in automatically developing these forecasts include:

- The lack of substantial sales history for a product (which especially makes obtaining seasonal forecasts very difficult),
- The difficulty of automatically matching a new product to a previous product or profile
- The inability to include planners' intuition (for example, the overall sales level of the product, how quickly the product will take off, how the product's sales will be affected by planned promotions) into a forecasting model.

Using a Bayesian approach, a short life-cycle forecasting algorithm has been developed that begins with a product's seasonal sales plan (that is developed externally to the system by the planner). As sales information arrives during the first few days or weeks of the season, the model generates a forecast by merging the information contained in the sales plan with the information contained in the initial sales data. These forecast updates can be critical to a company's success and can be used to increase or cancel vendor orders.

As forecasting consultants and software providers, Retek assists clients in obtaining “good” forecasts for future demands for their products based upon historical sales data and available causal information. Depending upon the information available, Retek’s software supports various forms of exponential smoothing and regression based forecasting. Frequently, however, clients already have some expectations of future demands in the form of sales plans. Depending upon the quality of their plans, they can provide very useful information to a forecasting algorithm (especially when only limited historical sales data is available). A forecasting algorithm has been developed that merges a customers sales plans with any available historical sales in a Bayesian fashion (that is, it uses *new information* to update or revise an existing set of probabilities).

Sales plans vs. historic data

In most retail situations, clients are interested in obtaining good product forecasts automatically with little or no human intervention. For stable products with years of historic sales data our time series approaches (Simple, Holt, Winters, Regression based Causal, etc.) produce adequate results. The problem arises when attempting to forecast products with little or no history. In such instances expert knowledge is required, generally in the form of sales plans. Given that both sales plans and time series forecasts are available, an obvious question exists: *When should the transition from sales plan to time series forecasting occur?* Suppose that in answer to that question (in a particular scenario) we have determined that thirteen weeks of history is the transition point. Does that mean that at 12 weeks the time series results are irrelevant and that at 14 weeks the sales plan has no value? Our intuition tells us that instead of there existing a hard-edge boundary, there is actually a steady continuum where the benefits from the sales plan decrease as we gather more historic sales data. This was the motivation for developing an approach that would combine the two forecasts in a reasonable manner.

Forecasting algorithm

The forecasting algorithm is provided only to give the technical reader insight into exactly how actual sales are combined with a plan to produce a forecast. What is important for all users to understand is that the algorithm is dependent on a sensitivity parameter which can be set by the client. A higher sensitivity setting will make the forecasting process increasingly reactive to actual sales. This parameter should be set based on the business goals of each individual client.

Let:

N be the number of periods in the current season (and $N = \infty$ for staple products),

M be the current period,

$p(j)$ be the sales plan for periods $j = 1, \dots, N$,

$x(j)$ be the achieved sales for period $j = 1, \dots, M$, and

α be a constant between 0.0 and 1.0 (This parameter influences the balance between model sensitivity and robustness, that is, how responsive the model is to new sales data).

Additionally, define

$P' \equiv \sum_{j=1}^M p(j)$ as the sum of the sales plan up to the current period,

$P'' \equiv \sum_{j=1}^N p(j)$ as the sum of the sales plan over the entire season, and

$X' \equiv \sum_{j=1}^M x(j)$ as the sum of the achieved sales up to the current period.

Finally, compute the forecasts as

$$\hat{x}(j) \equiv p(j) \left(\frac{X'}{P'} \right) \left(\frac{P'}{P''} \right)^{\alpha} + p(j) \left(1 - \left(\frac{P'}{P''} \right)^{\alpha} \right) \text{ for periods } j = M+1, \dots, N.$$

The motivation for this forecast is relatively clear. The forecast is a convex combination of a scaled version of the sales plan, $p(j) \left(\frac{X'}{P'} \right)$, and the original

sales plan itself. The scaled version of the sales plan is scaled based upon the ratio of the achieved historical sales to the historical sales plan (for example, if we had sold twice what we had planned to sell in the past, then the scaled plan would be twice the original plan). Thus, if we had no confidence in the magnitude of the original sales plan (but we still believed in its time profile or shape), then the scaled plan would probably be a good forecast on its own. On the other hand, if we really still believe in our plan and we don't believe that the recent past performance is indicative of future performance, then it would make sense to stick with the original sales plan as our forecast.

In our forecasting algorithm, the weights assigned to the scaled and original plans represent the confidence in the respective portions. As $\left(\frac{P'}{P''} \right)$ becomes larger

(that is, the portion of the plan that we can compare to historical sales increases), we tend to have more confidence in the scaled plan. For example, if $\left(\frac{P'}{P''} \right) =$

0.01, then we really don't have much information upon which to base the scaled plan.

On the other hand, for example, as the quantity approaches 0.5 (that is, the season is half way over), then we really should start seriously considering why the plan was incorrect and we may have greater belief in the scaled plan. Additionally, the α parameter is used to tweak the sensitivity of the forecasting method. As α increases the forecast will tend to stick closer to the original plan, and for small values of α the forecast will move rapidly towards the scaled plan as historical sales data becomes available. Clients should use their own data and judgement to determine an appropriate α for their particular business problem at hand.

Guidelines

Bayesian forecasting is primarily designed for use with new product/location positions. The following guidelines should be followed:

- 1 No more than one plan should exist for a given product/location position. If multiple plans are to be set up for different time periods, the domain should be set up with different forecasting levels for each time period of interest.
- 2 Any time period with non-zero actuals for a given product/location position should have a corresponding plan component (otherwise the system will assume a plan exists and equals zero and will act accordingly).
- 3 Any non-zero actuals not within the time period of interest should be overridden to zero.

Causal (promotional) forecasting methods

Causal, or promotional, forecasting requires three input streams:

- Time Series Data
- Historical Promotional Calendar
- Future Promotional Calendar

Promote decomposes the problem of promotional forecasting into two sub-tasks: estimating the effect promotions have on demand and forecasting baseline (that is no promotions) demand.

To accomplish the first task a stepwise regression sub-routine is used. This routine will take a time series and a collection of promotional variables and determine which variables are most relevant and what effect those relevant variables have on the series. Thus, the output from the algorithm is a selection of promotional variables and the effects of those variables on the series. In the second step the time series is de-causalized using the promo effects. Then, AutoES is used to calculate the baseline demand. Once we know the effects and we have the baseline demand we can generate a promotional forecast by applying the effects wherever the promotion is active in the future.

It should be noted that just because promotional forecasting is selected it doesn't necessarily imply that a "promotional" forecast will result. In some instances no promotional variables will be found to be statistically significant. In these cases the forecast ends up equivalent to a standard time series forecasts. (If users wish to force in certain promotional variables into the model this can be managed through forecasting maintenance parameters.)

Retek's experience in promotional forecasting has led us to believe that there are a few requirements that are necessary to successfully forecast retail promotions:

- Baseline forecasts need to consider seasonality; otherwise "normal" seasonal demand will be attributed to promotional effects.
- Promotional Effects need to be able to be analyzed at higher levels in the retail product and location hierarchies. This produces cleaner signals and alleviates issues involved in forecasting new SKUs and new Stores and issues involving data sparsity.
- Users need to be aware that the forecasting models cannot tell the difference between causal effects and correlated effects. What this means is that users should be wary of promotional effects attributed to an event that occurs at the same time every year. The system cannot distinguish between the promotional effect and the normal seasonality of the product. The same can be said for any two events that always occur at the same time (the combined effect will most likely be attributed to one or the other event).

Following is a description of the causal, or promotional forecasting, including the following topics:

- A description of the causal forecasting algorithm, including the causal variables used in RDF versions 9.1 and 9.5
- A description of the causal forecasting algorithm process
- The array interface between the Acumate environment and the AliAutoES binary
- Causal forecasting at the daily level
- Final considerations about causal forecasting

The causal forecasting algorithm

For purposes of understanding the algorithm, a promotional variable is defined as a causal event that is only active for certain series at certain points in time. It is assumed that these events are entered into the system independently of the forecasting algorithm and that a vehicle exists for determining which promotional variable are relevant to a particular series.

The core of the causal forecasting algorithm uses a stepwise regression sub-routine. This routine takes a time series and a collection of promotional variables, and determines which variables are most relevant and what effect those relevant variables have on the series. Thus, the output from the algorithm is a selection of promotional variables and the effects of those variables on the series.

Causal variable types define how causal variables are treated in the causal model fitting process (which includes a call to the lower level regression engine) and the forecast generation process where the model is used to extend the forecast over the forecast horizon.

Causal Variable Types

The RDF 9.5 Causal Variables have functionality to allow users to specify causal effects that the system uses to generate forecasts.

Automatic

The inclusion of the causal variable is decided by regression. If the causal variable is found to be significant on the training set it is included in the model, otherwise it is rejected.

Force In The Model

The causal variable is forced in to the model. Thus regression is not given a choice to reject even if the effect is considered insignificant by regression. As a result we will always return an effect.

Disabled

The variable is excluded from the model. Hence no effect will be returned either.

Override All

This type allows the user to specify a causal effect that will be used during the fitting and forecasting process. This is a directive that is recognized only by the causal engine and not by the lower level regression engine. For causal variables specified as *Override All*, the user also specifies the corresponding causal effect. The causal engine then de-causalizes the training data using the user-specified effect. The variable then is internally set Disabled to calculate the fit. During forecast generation the user-specified effect is used to determine the causal forecast.

Override Future

This type allows the user to specify a causal effect that will be used only during forecasting process and not during the fitting process. This is also a directive that is recognized only by the causal engine and not by the lower level regression engine. For causal variables specified as *Override Future*, the user also specifies the corresponding causal effect. The causal engine then internally sets the causal variable to *Automatic* to calculate the fit. The calculate effect is not however written back to the effects array (so as to not overwrite the user specified effect). During forecast generation the calculated effect is ignored and instead the user-specified effect is used to produce a causal forecast.

Override Higher Level

This type allows the system to use an average value of causal effects computed from product/location combination in the same group (the intersection level of the group to be used being specified during system setup) for product/location combinations on whose sales history alone a causal effect is unable to be computed for that causal variable. The way the system handles this is by having the effects array filled in with higher-level effects for those variable specified as *Override Higher Level* and which did not have an effect returned in the previous run. The effect array is passed to the causal engine. If the engine succeeds in computing a significant effect it writes back the system computed effect and uses it for generating a forecast. Otherwise it uses the effect passed in (which as mentioned is the average effect among product stores belonging to the same group) for generating the forecast.

Accounting for trend in the model

A trend is accounted for in the model by means of one specialized causal variable.

Types of causal models

Two different types of causal models are possible. The first type is an additive model in which each effect is considered to add a constant amount to the sales of the product. The second type of model is the multiplicative model in which the system determines a baseline value that is multiplied by each effect to get the final forecast value.

Causal forecasting algorithm process

The causal forecasting algorithm itself lies in the AutoES binary code, and executes in the following manner.

- 1 The binary reads the history of the time series.
- 2 The binary reads the type of each promotional variable into the system.
- 3 The binary reads in all the promotional variables that apply to the series.
- 4 The binary creates the internal promotional variable to allow the modeling of trend.
- 5 Promotional variables, internal promotional variables, promotional variable types, and the series itself are passed to the stepwise regression routine with the historic data serving as the dependent variables. Stepwise regression is run for both the additive and multiplicative models [Note: for the multiplicative model the logarithm of the historic data is used].
- 6 If the regression finds no significant promotional variables for either the additive or the multiplicative models then the casual method is considered to have failed to fit. In this case, the standard time series methods are used to generate a forecast and we skip to step 15.
- 7 The fit region is calculated for both models (additive and multiplicative). The fit at time t , $\text{fit}(t)$, is defined in terms of β_0 , the intercept of the regression, β_i , the effect corresponding to promotional variable i , and $p_i(t)$, is the value of promo variable I , in time t as:

$$\text{fit}(t) = \beta_0 + \sum \beta_i * p_i(t) \quad \text{for the additive model}$$

$$\text{fit}(t) = \beta_0 * \prod \beta_i^{p_i(t)} \quad \text{for the multiplicative model}$$

- 8 The RMSE is determined across the two fit regions and the winning model is the one that produces the least RMSE.
- 9 The Time Series is de-causalized in the fit region of the history, by removing the causal effects (that is subtracting the additive causal effects, respectively dividing by the multiplicative causal effects for the product/location/time positions where the corresponding promo variables are “on”)

- 10 A seasonal model is fitted to the de-causalized series. The seasonal model is the winner of a competition between Seasonal ES, Additive Winters and Multiplicative Winters. The RMSE is determined across the fit region on the de-causalized series and the winning model is the one that produces the least RMSE.
- 11 The winning seasonal model is then used to determine the seasonal forecast by applying the $\text{fit}(t)$ function across the forecast horizon. This forecast is also exported as the *baseline* forecast (that is forecast without any causal influences).
- 12 The forecast is obtained by re-causalizing the seasonal forecast. This is done by adding back the additive causal effects, respectively multiplying by the multiplicative causal effects for the product/location/time positions where the corresponding promo variables are “on” in the forecast region.
- 13 The binary writes the winning promotional variables effects back to the database.
- 14 The selected model is recorded in the database.
- 15 The binary records the forecast and the baseline in the database.

Confidence intervals constraining in Causal:

The confidence interval is capped using the following rule:

if $\text{interval}(0) = \text{stdev}(\text{frst error}) < 0.3 * \text{forecast}(i)$, then $\text{interval}(i) = 0.3 \cdot \text{forecast}(i) \cdot \sqrt{(i+1)}$

if $\text{interval}(0) > \text{forecast}(i)$, then $\text{interval}(i) = \text{forecast}(i) \cdot \sqrt{(i+1)}$

Otherwise, $\text{interval}(i) = \text{interval}(0) \cdot \sqrt{(i+1)}$

$\text{cuminterval}(i) = \sqrt{\text{cuminterval}(i-1)^2 + \text{interval}(i)^2}$

Besides the interval constraining, in Causal.cpp the interval calculation was changed. Specifically, to the extension formula:

$\text{interval}(i)^2 = \text{mse} \cdot (i+1)$

Causal forecasting array interface description

The promotional variables are stored in the database named promo. In the database named pmaint, the following arrays are stored:

- The promotional variables aggregated to the forecast level, along with promotional effects arrays. Both the additive and multiplicative effects are saved for promotional effectiveness assessment.
- The type arrays (arrays in which the type of each variable – see “Algorithm Description” - is stored for each product/location intersection at the forecast level).
- The list of variables to be considered in the stepwise regression.

The promotional variables (located in the promo database) are aggregated at the forecast level and saved in pmaint by the forecasting wrapper, fbatch.wrapautoes.

For causal forecasting the path to the database “pmaint” is passed as input to the AliAutoES binary. From pmaint the binary reads the promotion list and the promotional variables, and writes back the additive and multiplicative causal effects and intercept in the right arrays.

Forecast and baseline forecast are stored in separate databases.

Causal forecasting at the daily level

The causal forecasting at the daily level is calculated by spreading the weekly causal forecast down to days. The spreading utilizes causal daily profiles thus obtaining a causal forecast at the day granularity. This functionality does not exist in RDF 9.1.

The daily casual forecast process executes in the following manner.

- 1 Preprocess the day level promotional variables by multiplication with daily profiles. Aggregate the preprocessed continuous day level promotional variables to the week level.
- 2 Calculate the causal forecast at the weekly level. Set promotional effects if desired. Use the RDF causal engine to generate the forecast.

- 3 Calculate the multiplicative promotional effects at the sku/str level for every promo variable. The effects can be either
 - Manually preset (see step 1)
 - Calculated. When calculating the causal forecast, the calculated causal effects are written back to the database. If the effects are calculated at higher level than sku/str, the effects will be replicated down to sku/str – reasonable assumption since the effects are multiplicative. If source level forecasting is used and causal method is used both at the source level and at the final level then the effects from the final level will be used.
- 4 Daily profiles are calculated, using the Curve module. Since we use as much history as possible and we average it over seven days we assume these profiles are decausalized. The decausalized daily profiles capture the day of week effect and should be quite stable.
- 5 Causal effects are applied to the daily profiles. We multiply the profiles by the causal effects. Then, we have to renormalize the profiles.

Example:

For every SKU/Store combination, calculate a normal week-to-day profile based on historic data (note that this profile is already computed for spreading the weekly forecasts to the day level). Suppose for a certain product, the profile is as follows:

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday	Week
10%	10%	10%	10%	20%	30%	10%	100%

Suppose that in the past, the promotion was held on Wednesday, Thursday, and Friday of week w6:

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
		P	P	P		

Then the continuous weekly indicator for this promotion in w6 should be set to 0.4 which is the sum of the weights of Wednesday, Thursday, and Friday

Now assume that the same promotion will be held in a future week, say w36, but only on Thursday:

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
			P			

Then the continuous weekly indicator for w36 should be set to 0.1 which is the weight of Thursday only.

The approach to use the continuous promotion indicators to generate an accurate causal forecast at the day level is as follows:

- Calculate the weekly multiplicative effect for the promotion using the standard causal forecasting system with continuous indicators.
- Calculate the forecast for w36 using the standard causal forecasting system with continuous indicators. Note that with the multiplicative model and continuous promotional variables the multiplicative factor for a promotional event is obtained by raising the multiplicative effect to the power of the promotional variable. Let's say this multiplicative factor for w32 was calculated to be 2 (the promotion doubles the regular sales).
- Update the week-to-day profile of w36 so that the weight of Thursday is doubled (the multiplicative factor is 2):

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday	Week
10%	10%	10%	20%	20%	30%	10%	110%

Normalize the profile for w36:

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday	Week
9%	9%	9%	18%	18%	27%	9%	100%

Finally, spread the forecast of w36 using the normalized profile.

Final considerations about causal forecasting

In Retek's approach to causal forecasting, the causal effects are obtained by fitting a stepwise linear regression model that determines which variables are most relevant and what effect those relevant variables have on the series. The data used to fit the regression is the fit history of each time series, so basically we fit a model per time series. A problem arises due to potential lack of significant data, that is, when a promotional variable is not represented in the history but is present in the forecast region. In that case, the effect for that variable would not be computed at all, thus affecting the accuracy of the forecast. There are a few solutions that make use of the effects from other similar time series. One solution would be to do source level causal forecasting and then spread down to the final level. This would be equivalent to using the effects at the source level for time series that have no causal variable instances in the history. However, this has a serious conceptual drawback: by aggregating the promotional variables at the source level, we would force the effects on the other time series in the same aggregation class that otherwise would not have the causal variables on at the same time. A better solution will be incorporated in RDF version 9.5: whenever a causal effect cannot be computed because of lack of significant data an averaged effect from the other TS in the same aggregation class is going to be used instead – see Override Higher Level type.

In RDF 9.1, source level causal forecasting used binary promotional variables aggregated to the source level with OR. It is possible, however, to use continuous promotional variables and aggregate them by conventional methods such as add and average. This capability is important for RDF 9.5 and subsequent future versions, and would be useful for causal daily forecast and improve the performance of source level causal forecasting.

Glossary

Note: With a few exceptions, this glossary contains definitions of terms specific to Retek Demand Forecasting. For further definitions of terms and concepts relating to the user interface on which Retek Demand Forecasting runs, see the Retek Predictive Application Server online help or User Guide.

additive seasonal method

Also referred to as Additive Winters Model, this model is similar to the Multiplicative Winters model, but is used when zeros are present in the data. This model adjusts the un-seasonalized values by adding the seasonal index (for the forecast horizon).

alert

A notice displayed to system users that a forecasted value is above or below user defined limits (an exception).

Alert Manager window

A window that displays the alerts assigned to you.

This dialog provides a list of all identified instances in which a monitored measure's values fall outside a set of defined limits. You may pick an alert from this list and have RCS automatically build a workbook containing the measure values that triggered the alert.

AutoES Method

See Automatic Exponential Smoothing.

Automatic Exponential Smoothing Method

Retek Demand Forecasting fits the sales data to a variety of exponential smoothing (time series) models of forecasting, and the best model is chosen for the final forecast. The candidate methods considered by AutoES are: Simple ES, Intermittent ES, Trend ES, Multiplicative Seasonal, Additive Seasonal and Seasonal ES. The final selection between the models is made according to a performance criterion (Bayesian Information Criterion) that involves a tradeoff between the model's fit over the historic data and its complexity.

Bayesian Method

Useful for short life cycle forecasting and for new products with little or no historic sales data. The Bayesian method requires a product's known sales plan (created externally to RDF) and considers a plan's shape (the selling profile or life cycle) and scale (magnitude of sales based on Actuals). The initial forecast is equal to the sales plan but as sales information comes in, the model generates a forecast by merging the sales plan with the sales data. The forecast is adjusted so that the sales magnitude is a weighted average between the original plan's scale and the scale reflected by known history.

Causal Method

Causal is a forecasting method used for promotional forecasting and can only be selected if Promote is implemented (see Chapter 3 for more information on Promote). Typically, the Causal method is used at the Final Levels (i.e., item/str/week). Causal uses a Stepwise Regression sub-routine to determine the promotional variables that are relevant to the time series and their lift effect on the series. AutoES utilizes the time series data and the future promotional calendar to generate future baseline forecasts. By combining the future baseline forecast and each promotion's effect on sales (lift), a final promotional forecast is computed.

confidence interval

The confidence percentage value used to calculate the value of the upper confidence limit. A different confidence level can be set for every forecast level appearing in the workbook

confidence limit

The upper bound of a normal probability distribution that is centered about the forecast value. It is calculated based on the Default Confidence % parameter value entered in the Forecast Administration Workbook, or on the Confidence Level value entered in the Confidence Worksheet of the Forecast Approval Workbook (if such a value has been entered to override the default).

Croston's model of exponential smoothing

See Intermittent Exponential Smoothing

Curve

An optional automated predictive solution that transforms organization-level assortment plans into base-level weekly sales forecasts.

exception

A forecast value that is greater than or less than a user-defined limit.

exponential smoothing

A form of a weighted moving average. It weights decline in data exponentially. Most recent data weighted more heavily. Requires smoothing constant (α). Ranges from 0 to 1. Subjectively chosen.

final forecast level

A low level in a hierarchy from which a forecast is generated, and at which approvals and data exports can be performed. Often, data from forecasts at a low level is insufficient to generate reliable forecasts without first aggregating the data to a higher level and then spreading the data back to the low level.

forecast-driven planning

Planning that keys off of forecasts fed directly into a planning system. Connection to Retek Demand Forecasting (RDF) is built directly into the business process supported by Retek Predictive Planning through an automatic approval of a forecast that is fed directly in the planning system. This allows you to accept all or part of Sales Value forecast. Once that decision is made, the balance of business measures are planned within Retek Predictive Planning.

Holt's model of exponential smoothing

See Trend Exponential Smoothing

interactive forecasting

A workbook in RDF used to simulate forecast by modifying parameters such as Forecast Method and History Start Date

IntermittentES

See Intermittent Exponential Smoothing.

Intermittent Exponential Smoothing.

Retek Demand Forecasting fits the data to the Croston's model of exponential smoothing. This method should be used when the input series contains a large number of zero data points (that is, intermittent demand data). The original time series is split into a Magnitude and Frequency series, then the Simple ES model is applied to determine level of both series. The ratio of the magnitude estimate over the frequency estimate is the forecast level reported for the original series.

like item

See like SKU

like SKU

An item that will be used as a model to forecast a new item introduction.

lost sales

Periods in sales data in which there was no inventory to meet consumer demand

measure

Any item of data that can be represented on a grid in worksheets.

measure description

The description of the measure that can be viewed in a workbook. This description may contain relationships and calculations.

measure function

Internal functions that can be used to simplify building calculations for a measure.

measure identifier

The combination of role, version, metric, and units that uniquely specifies a single measure.

metric

A measure definition with the role, version, and units omitted.

Multiplicative Seasonal

(also referred to as Multiplicative Winters Model) This model extracts seasonal indices that are assumed to have multiplicative effects on the un-seasonalized series.

Predict

In Retek Demand Forecasting (RDF), Predict refers to RDF's statistical forecasting capabilities.

Preprocessing

In Retek Demand Forecasting (RDF), Preprocessing refers to a module that processes data before forecasts are generated to adjust for situations such as lost sales and unusually high demand.

profile

Spreading ratios that are used in the Curve process. Typical profiles can include store participation, size distribution, and time (phase-to-week) profiles, as well as other information. Profiles are generated using historical data and phase definitions, based on your system configuration.

Profile Based

Retek Demand Forecasting generates a forecast based on a seasonal profile that can be created in Curve or a legacy system. Profiles can also be copied from another profile and adjusted. Using historic data and the profile, the data is de-seasonalized and then fed to the Simple ES method. The Simple forecast is then re-seasonalized using the profiles

Profile Spread

Used at the final level to utilize a profile (either generated externally or with Curve) to determine the spreading ratios from the Source level forecast down to the Final level.

Promote

Promote is an optional add-on automated predictive solution that allows you to incorporate the effects of promotional and causal events, such as radio advertisements and holiday occurrences, into your time series forecasts. The promotional forecasting process uses both past sales data and promotional information to forecast future demand.

promotion planning

A workbook template and simulation process used within the context of promotional forecasting. Promotion planning involves specifying whether the event status for a particular promotional variable is active (on) or inactive (off) for a specific product/location/calendar combination. When past promotional events are represented as accurately as possible, the modeling routine can more precisely detect correlation between event occurrences and changes in sales values.

promotional effectiveness

A workbook template used in the context of promotional forecasting. This workbook allows you to analyze the effects of promotions on SKUs at both the micro and the macro level. “What if” analysis can also be performed on the results of promotional forecasts, as you can modify future and past promotional inputs, the system-estimated effects of promotions, and the promotional forecasts themselves.

promotional forecasting

Promote’s forecasting technique (also referred to as Causal forecasting) uses promotional factors and events to predict future demand. Promotion events are events such as advertisements, holidays, competitor information, and other factors that affect the normal selling cycle for a business.

promotion group

A set of products or locations that are believed to exhibit similar effects during common causal events. Promotion groups should be established to maximize the number of time series for each group (so each promotional event can be evaluated from as many different observations as possible) while ensuring that each time series is affected by causal events to the same degree.

Seasonal ES Method

A combination of several Seasonal methods. This method is generally used for known seasonal items or forecasting for long horizons. This method applies the Multiplicative Seasonal model unless zeros are present in the data, in which case the Additive Winters model of exponential smoothing is used. If less than 2 years of data is available, then a Seasonal Regression model is used. If there is too little data to create a seasonal forecast (in general, less than 52 weeks), then the system will select from the Simple ES, Trend ES and Intermittent ES methods.

Seasonal Regression

Seasonal Regression cannot be selected as a forecasting method, but is only a candidate model used when the Seasonal ES method is selected. This model requires a minimum of 52 weeks of history to determine seasonality. Simple Linear Regression is used to estimate the future values of the series based on a past series. The independent variable is the series history one-year or one cycle length prior to the desired forecast period, and the dependent variable is the forecast. This model assumes that the future is a linear combination of itself one period before plus a scalar constant.

Simple/Intermittent ES Method

A combination of the Simple ES and Intermittent ES methods. This method applies the Simple ES model unless a large number of zero data points are present, in which case the Croston's model is applied.

SimpleES

See Simple Exponential Smoothing.

Simple Exponential Smoothing Method

Retek Demand Forecasting uses a simple exponential smoothing model to generate forecasts. Simple ES ignores seasonality and trend features in the demand data and is the simplest model of the exponential smoothing family. This method can be used when less than 1 year of historic demand data is available.

simple moving average

See Average Method

sister store

A store that will be used as a model to forecast a new store.

source level forecast

The level at which the aggregate, more robust forecast is run.

Time series

Set of evenly spaced numerical data obtained by observing response variable at regular time periods. This data is used to forecast based only on past values. It assumes that factors influencing past and present will continue influence in future

training window

The number of weeks of historical sales data to use in generating a forecast.

TrendES

Trend Exponential Smoothing.

Trend Exponential Smoothing

(also referred to as Holt's Model) Retek Demand Forecasting fits the data to the Holt model of exponential smoothing. The Holt model is useful when data exhibits a definite trend. This method separates out base demand from trend, then provides forecast point estimates by combining an estimated trend and the smoothed level at the end of the series.

wizard

A set of screens that guide you through the process of creating a new workbook or performing other actions in an application, by asking you various questions and having you select values.

workbook

The framework used for displaying data and user functions. Workbooks are task-specific and may contain one or more worksheets. Users define the format of their workbooks.

See also workbook template, worksheet.

workbook template

The framework for creating a workbook. You build each new workbook from an existing workbook template, such as Pre-Season Financial Plan or Forecasting Administration. Several workbook templates are supplied with the Retek Predictive Solutions, and are available for selection when you choose File+New to create a new workbook.

worksheet

A multidimensional spreadsheet used to display workbook-specific information. Worksheet data can also be displayed in chart format.

Index

- Additive Winters exponential smoothing, 140
- Adjusting forecast parameters, 64
- Administration Workbooks, 12
- alerts, 5, 165
- Approval Worksheet, 44
- Approve a profile, 108
- assortment planning, 4
- Auto Exponential Smoothing (AutoES)
 - forecasting, 133
- AutoES forecasting, 133
- Automatic Exponential Smoothing (AutoES)
 - Bayesian Information Criterion, 144
- Automatic forecast level selection, 147, 149
- Automatic method selection, 130
- Basic Settings workflow tab, 13
- Batch Forecast
 - Manually request an unscheduled forecast, 65
 - Run Batch Forecast, 37
- Bayesian analysis
 - Overview, 130
- Bayesian Forecasting
 - Forecasting algorithm, 153
 - Guidelines for using, 155
 - Overview, 152
 - Sales plans vs. historic data, 153
- Bayesian Information Criterion (BIC), 144
- BIC, 144
- Causal ES preprocessing method, 120, 124
- Causal forecasting methods. *See* Promotional forecasting methods
- Causal Median preprocessing method, 120, 124
- Changing forecast parameters, 64
- Clear preprocessing method, 121, 125
- Confidence limits, 133
- Confidence Worksheet, 47
- Core Profile Worksheet, 96
- Create a new profile template, 98, 110
- Croston's method, 137
- Curve, 166
 - About profiles, 86
 - Approve a profile, 108
 - Core Profile Worksheet, 96
 - Create a new profile template, 110
 - Create new profile templates, 98
 - Dynamic Profile Library, 110
 - Forecasting methods used in, 150
 - Manually request a generated profile, 107
 - Modify training window values for a subset of items in a profile, 107
 - Overview, 7, 85
 - Procedures, 106
 - Profile Administration Workbook, 95
 - Profile Approval Workbook, 101
 - Profile generation techniques, 87
 - Profile Maintenance Workbook, 98
 - Run Batch Profile, 100
 - Set default parameters for a profile, 106
 - Set Retek Demand Forecasting to run plan explosion, 109
 - Workbook and worksheet descriptions, 85
- Default parameters
 - Setting for SKU Explosion, 106
- default ratios, 92
- Delete Forecasts Wizard, 38
- Deleting forecasts, 70
- Developing promotional forecast methods, 72
- Dynamic Profile Library, 98, 110
- exception, 166
- exponential smoothing
 - automatic (AutoES), 133
 - components of, 135
 - seasonal (SeasonalES), 141
 - simple, 136
- Exponential smoothing, 129, 136
 - Additive Winters, 140
 - Holt, 138
 - Multiplicative Winters, 139
- Export Forecasts
 - Manually export a forecast to another program, 68
 - Overview, 56
- Export Forecasts Wizard, 56, 57, 68
- Export Parameters Worksheet, 59
- Final Actuals vs. Forecasts Worksheet, 55
- Final Error Measure Worksheet, 53
- final forecast level, 167
- Final Forecast Worksheet for Forecast Approval, 41
- Final Forecasting Parameter Worksheet, 49
- Final Interactive Forecasting Worksheet, 50
- final level forecasts, 6
- Final Level Worksheet – Advanced Settings for Forecast Administration, 20
- Final Level Worksheet – Basic Settings for Forecast Administration, 14

- Final Level Worksheet – Basic Settings for Forecast Maintenance, 28
- Final Parameters Worksheet, 45
- Final Profile Worksheet, 105
- Forecast Administration
 - Forecast Administration Workbook, 12
- Forecast Administration
 - Forecast Administration Wizard, 13
- Forecast Administration
 - Forecast Administration Wizard, 13
- Forecast Administration
 - Basic Settings workflow tab, 13
- Forecast Administration
 - Final Level Worksheet – Basic Settings, 14
- Forecast Administration
 - Source Level Worksheet – Basic Settings, 18
- Forecast Administration
 - Final Level Worksheet – Advanced Settings, 20
- Forecast Administration
 - Source Level Worksheet – Advanced Settings, 24
- Forecast Approval
 - Approval Worksheet, 44
 - Confidence Worksheet, 47
 - Final Forecast Worksheet, 41
 - Final Parameters Worksheet, 45
 - Forecast Approval Wizard, 40
 - Forecast Approval Workbook, 39
 - Review, modify, and approve forecast values, 66
 - Source Level Worksheet, 43
 - Wizard, 66
- Forecast Approval Wizard, 40
- Forecast Approval Workbook, 39
- Forecast default parameters, 63
- Forecast Export Administration
 - Export Parameters Worksheet, 59
 - Set or change forecast export parameters, 69
- Forecast Export Administration Workbook, 58, 69
- Forecast Maintenance
 - Adjust forecast parameters, 64
 - Advanced Settings workflow tab, 29
 - Basic Settings workflow tab, 27
 - Final Level Worksheet – Basic Settings, 28
 - Forecast Maintenance Wizard, 26
 - Forecast Maintenance Workbook, 26
 - Sister Store Worksheet, 30
 - Source Level Worksheet – Basic Settings, 29
- Forecast Maintenance Wizard, 26, 64
- Forecast parameters, 64
- Forecast Scorecard, 51, 52, 69
 - Final Actuals vs. Forecasts Worksheet, 55
 - Final Error Measure Worksheet, 53
 - Forecast Scorecard Workbook, 51
 - Wizard, 52
- Forecast Scorecard Wizard, 52, 69
- forecast-driven planning, 167
- Forecasting challenges and solutions, 2
- Forecasting demand for new products and locations, 3
- Forecasting methods used in RDF
 - Additive Winters exponential smoothing, 140
 - Auto Exponential Smoothing (AutoES), 133
 - Automatic forecast level selection, 147
 - Automatic method selection, 130
 - Bayesian analysis, 130
 - Bayesian Forecasting, 152
 - Bayesian Information Criterion (BIC), 144
 - Causal (promotional), 155
 - Confidence limits, 133
 - Croston's method, 137
 - Exponential smoothing, 129
 - Holt exponential smoothing, 138
 - Multiplicative Winters exponential smoothing, 139
 - Overview, 129
 - Prediction intervals, 130
 - Profile-based forecasting, 150
 - Promotional forecasting, 131
 - Regression analysis, 130
 - Seasonal Exponential Smoothing (SeasonalES), 141
 - Seasonal regression forecasting, 142
 - Simple exponential smoothing, 136
 - Simple moving average forecasts, 135
 - Source level forecasting, 131
 - Statistical (time series), 131
- Forecasting overview, 9
- Forecasting Workbooks, 9
 - Deleting, 70
 - Overview, 9
- Forecasts
 - Deleting, 70
- frst.out file, 57
- Generation of final level and source level forecasts, 6
- Handling items with limited demand histories, 2
- Handling lost sales and unusually high demand, 3
- high demand, 3
- Holt exponential smoothing, 138
- Incorporating effects of promotions on demand, 4
- Interactive Forecasting
 - Final Forecasting Parameter Worksheet, 49

- Final Interactive Forecasting Worksheet, 50
- Interactive Forecasting Workbook, 47
- intermittent demand, 137
- items with limited demand histories, 2
- loading default ratios, 92
- Lost sales, 113, *See also* Preprocessing
- Lost Sales – Causal ES preprocessing method, 121, 125
- Lost Sales – Causal Median preprocessing method, 121, 124
- Lost Sales – Retek Median preprocessing method, 120, 124
- Lost Sales – Seasonal ES 1 preprocessing method, 121, 124
- Lost Sales – Seasonal Median 1 preprocessing method, 121, 124
- Lost Sales – Std ES preprocessing method, 120, 124, 127
- Lost Sales – Std Median preprocessing method, 120, 124
- lost sales and unusually high demand, 3
- Manually export a forecast to another program, 68
- Manually request a generated profile, 107
- Manually request an unscheduled forecast, 65
- measure description, 168
- measure function, 168
- measure identifier, 168
- measures, 168
- metric, 168
- Modify training window values for a subset of items in a profile, 107
- Monitor the accuracy of Retek Demand Forecasting forecasts, 69
- Multiplicative Winters exponential smoothing, 139
- new products and locations, 3
- No Filtering preprocessing method, 121
- non-sparse items, 116
- plan explosion, 109
- Plan promotions
 - overview of promotional forecasting, 71
 - workbooks and wizards, 75
- Predict
 - Adjust forecast parameters for specific products or locations, 64
 - definition, 168
 - Delete a forecast, 70
 - Delete Forecasts, 38
 - Export Forecasts, 56
 - Forecast Administration, 12
 - Forecast Approval Workbook, 39
 - Forecast Export Administration, 58
 - Forecast Maintenance, 26
 - Forecast Scorecard, 51
 - Forecasting methods used in, 131
 - Interactive Forecasting, 47
 - Manually export a forecast to another program, 68
 - Manually request an unscheduled forecast, 65
 - Monitor the accuracy of Retek Demand Forecasting forecasts, 69
 - Overview, 7, 9
 - Procedures, 63
 - Review, modify, and approve forecast values, 66
 - Run Batch Forecast, 37
 - Set or change default parameters for a forecast, 63
 - Set or change forecast export parameters, 69
 - Workbook and worksheet descriptions, 9
- Prediction intervals, 130
- Preprocessing
 - Administration Wizard, 114
 - Administration Workbook, 114
 - Administration Worksheet, 114
 - definition, 168
 - Examples, 125
 - Handling items with an outage indicator, 116
 - Handling non-sparse items, 116
 - Methods implemented in RDF, 118
 - Overview, 8, 113
 - Run Batch Update, 116
 - Run Batch Workbook, 116
 - Workbook and worksheet descriptions, 113
- preprocessing methods
 - Causal ES, 120, 124
 - Causal Median, 120, 124
 - Clear, 121, 125
 - Lost Sales – Causal ES, 121, 125
 - Lost Sales – Causal Median, 121, 124
 - Lost Sales – Retek Median, 120, 124
 - Lost Sales – Seasonal ES 1, 121, 124
 - Lost Sales – Seasonal Median 1, 121, 124
 - Lost Sales – Std ES, 120, 124, 127
 - Lost Sales – Std Median, 120, 124
 - mathematical formulation of, 122
 - No Filtering, 121
 - Retek Median, 118, 122, 126
 - Seasonal ES 1, 119, 124
 - Seasonal Median 1, 119, 123
 - Std ES, 119, 123, 126
- preprocessing methods
 - Std Median, 118
- preprocessing methods
 - Std Median, 122

- preprocessing methods
 - Examples, 125
- preprocessing methods
 - Std Median, 125
- Profile Administration Workbook, 95
- Profile Approval Wizard, 101
- Profile Approval Workbook, 101
- Profile Approval Worksheet, 102
- Profile Maintenance Wizard, 98
- Profile Maintenance Workbook, 98, 99
- Profile-based forecasting
 - About profiles, 86
 - Approve a profile, 108
 - Core Profile Worksheet, 96
 - Create a new profile template, 110
 - Creating new profile templates, 98
 - Dynamic Profile Library, 110
 - Example, 151
 - forecasting new SKUs, 151
 - Manually request a generated profile, 107
 - Method, 150
 - Modify training window values for a subset of items in a profile, 107
 - Overview, 150
 - Procedures, 106
 - Profile Administration Workbook, 95
 - Profile Approval Workbook, 101
 - Profile generation techniques, 87
 - Profile Maintenance Workbook, 98
 - Run Batch Profile, 100
 - Set default parameters for a profile, 106
 - Set Retek Demand Forecasting to run plan explosion, 109
- profiles
 - creating new, 98, 110
 - definition, 168
 - overview, 86
- Promote
 - definition, 169
 - Forecasting methods used in, 155
 - Overview, 7, 71, 169
 - Procedures, 84
 - Promotion Planner Workbook Template, 76
 - Promotion Selector Template, 78
 - Set up the system to run a promotional forecast, 84
 - Workbook and worksheet descriptions, 71, 169
 - workbooks and wizards, 75
- Promotion Group, 170
- Promotion Planner Workbook Template, 76
- promotion planning, 71, 75, 77, 169
- Promotion Selector Template, 78
- promotional effectiveness, 4, 169
- promotional forecasting, 72, 75
- Promotional forecasting. *See also* Promote, *See also* Promote
 - Comparison to Statistical forecasting, 72
 - definition, 169
 - Overview, 71, 169
 - Retek's approach, 74
 - Set up the system to run a promotional forecast, 84
 - When to use, 131
 - workbooks and wizards, 75
- Promotional forecasting methods
 - Algorithm, 157
 - Causal forecasting array interface description, 161
 - Causal forecasting at the daily level, 161
 - Causal variable types, 157
 - Considerations and guidelines, 163
 - forecasting algorithm process, 159
 - Overview, 155
- Regression analysis, 130
- Retek Demand Forecasting, 85
 - Curve (profile-based forecasting), 7, 85
 - feature summary, 5
 - Forecasting methods used in, 129
 - Forecasting workbooks, 9
 - interfaces with other products, 1
 - modules, 7
 - Predict (statistical forecasting), 7
 - Promote (promotional or causal forecasting), 7
 - Promotional Forecasting workbooks and wizards, 75
- Retek Median preprocessing method, 118, 122, 126
- Review, modify, and approve forecast values, 66
- Run Batch Forecast, 37
- Run Batch Forecast Wizard, 37, 65
- Run Batch Profile, 100
- sales predictions based on assortment plan, 4
- Seasonal ES 1 preprocessing method, 119, 124
- Seasonal Exponential Smoothing (SeasonalES) forecasting, 141
- Seasonal Median 1 preprocessing method, 119, 123
- Seasonal regression forecasting, 142
- SeasonalES forecasting, 141
- Selecting the best forecasting method, 2
- Set default parameters for a profile, 106
- Set forecast default parameters, 63
- Set or change forecast export parameters, 69

- Set Retek Demand Forecasting to run plan explosion, 109
- Set up the system to run a promotional forecast, 84
- Simple exponential smoothing, 136
- Simple moving average forecasts, 135
- Simple Moving Average forecasts, 135
- Sister Store Worksheet, 30
- Source level forecasting, 6, 131
- Source Level Worksheet – Advanced Settings for Forecast Administration, 24
- Source Level Worksheet – Basic Settings for Forecast Administration, 18
- Source Level Worksheet – Basic Settings for Forecast Maintenance, 29
- Source Level Worksheet for Forecast Approval, 43
- Source Level Worksheet for Profile Approval, 103
- spreading ratios, 86
- Statistical forecasting
 - Procedures, 63
 - reasons for using, 132
- Statistical forecasting methods, 131
 - Additive Winters exponential smoothing, 140
 - Auto Exponential Smoothing (AutoES) forecasting, 133
 - Bayesian Information Criterion (BIC), 144
 - Croston's method, 137
 - Holt exponential smoothing, 138
 - Multiplicative Winters exponential smoothing, 139
 - Seasonal Exponential Smoothing (SeasonalES) forecasting, 141
 - Seasonal regression forecasting, 142
 - Simple exponential smoothing, 136
 - Simple moving average forecasts, 135
- Std ES preprocessing method, 119, 123, 126
- Std Median preprocessing method, 118, 122, 125
- Time series forecasting methods, 131
- training window, 171
- Training window values, 107
- wizard, 171
- workbook, 171
- workbook template, 171
- worksheet, 172

ⁱ Chatfield, C. and M.Yar, 1990, “Prediction intervals for the Holt-Winters forecasting procedure”, *International Journal of Forecasting*, 6, 127-137

ⁱⁱ Chatfield, C. and M.Yar, 1990, “The Holt-Winters forecasting procedure”, *Applied Statistics*, 27, 264-279

ⁱⁱⁱ Chatfield, C. and M.Yar, 1990, “Holt-Winters forecasting: theory and practice”, *The Statistician*, 37, 129-140