Oracle® Retail Demand Forecasting

User Guide for the RPAS Classic Client

Release 16.0.3

F12586-03

July 2019

ORACLE®

Oracle Retail Demand Forecasting User Guide for the RPAS Classic Client, Release 16.0.3

F12586-03

Primary Author: Melissa Artley

reconfigurations, reassembly or reverse assembly, re-engineering or reverse engineering and recompilations or reverse compilations of the VAR Applications or any derivatives of the VAR Applications. You acknowledge that it shall be a breach of the agreement to utilize the relationship, and/or confidential information of the VAR Applications for purposes of competitive discovery.

The VAR Applications contain trade secrets of Oracle and Oracle's licensors and Customer shall not attempt, cause, or permit the alteration, decompilation, reverse engineering, disassembly or other reduction of the VAR Applications to a human perceivable form. Oracle reserves the right to replace, with functional equivalent software, any of the VAR Applications in future releases of the applicable program.

Contents

3 Setting Forecast Parameters

4 Generating and Approving a Forecast

5 Preprocessing Administration Workbook

6 Source Measure Maintenance Workbook

7 Preprocessing Review and What-if Workbook

12 Forecast Analysis Tools

13 Promote (Promotional Forecasting)

14 Grouping Management

15 Cross Promotional Effects Module (CPEM)

16 Curve

17 Grade

A Appendix: Preprocessing

Glossary

List of Tables

List of Figures

Send Us Your Comments

Oracle Retail Demand Forecasting User Guide for the RPAS Classic Client, Release 16.0.3.

Oracle welcomes customers' comments and suggestions on the quality and usefulness of this document.

Your feedback is important, and helps us to best meet your needs as a user of our products. For example:

- Are the implementation steps correct and complete?
- Did you understand the context of the procedures?
- Did you find any errors in the information?
- Does the structure of the information help you with your tasks?
- Do you need different information or graphics? If so, where, and in what format?
- Are the examples correct? Do you need more examples?

If you find any errors or have any other suggestions for improvement, then please tell us your name, the name of the company who has licensed our products, the title and part number of the documentation and the chapter, section, and page number (if available).

> **Note:** Before sending us your comments, you might like to check that you have the latest version of the document and if any concerns are already addressed. To do this, access the Online Documentation available on the Oracle Technology Network Web site. It contains the most current Documentation Library plus all documents revised or released recently.

Send your comments to us using the electronic mail address: retail-doc_us@oracle.com

Please give your name, address, electronic mail address, and telephone number (optional).

If you need assistance with Oracle software, then please contact your support representative or Oracle Support Services.

If you require training or instruction in using Oracle software, then please contact your Oracle local office and inquire about our Oracle University offerings. A list of Oracle offices is available on our Web site at http://www.oracle.com.

Preface

The Oracle Retail Demand Forecasting User Guide for the RPAS Classic Client describes the application's user interface and how to navigate through it.

Audience

This document is intended for the users and administrators of Oracle Retail Demand Forecasting. This may include merchandisers, buyers, and business analysts.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Retail Demand Forecasting Release 16.0.3 documentation set:

- *Oracle Retail Demand Forecasting Installation Guide*
- *Oracle Retail Demand Forecasting Release Notes*
- *Oracle Retail Demand Forecasting User Guide for the RPAS Classic Client*
- *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*
- Oracle Retail Predictive Application Server documentation

The following documentation may also be needed when implementing RDF:

- *Oracle Retail Predictive Application Server Batch Script Architecture Implementation Guide*

Supplemental Documentation

The following document is available through My Oracle Support at the following URL:

https://support.oracle.com

Oracle Retail Demand Forecasting 16.0.3 Cumulative Fixed Issues (Note ID 2413398.1)

This document details the fixed issues and defects for all RDF, Curve, and Grade patch releases prior to and including the current release.

Customer Support

To contact Oracle Customer Support, access My Oracle Support at the following URL:

https://support.oracle.com

When contacting Customer Support, please provide the following:

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

Review Patch Documentation

When you install the application for the first time, you install either a base release (for example, 16.0) or a later patch release (for example, 16.0.3). If you are installing the base release, additional patch, and bundled hot fix releases, read the documentation for all releases that have occurred since the base release before you begin installation. Documentation for patch and bundled hot fix releases can contain critical information related to the base release, as well as information about code changes since the base release.

Improved Process for Oracle Retail Documentation Corrections

To more quickly address critical corrections to Oracle Retail documentation content, Oracle Retail documentation may be republished whenever a critical correction is needed. For critical corrections, the republication of an Oracle Retail document may at times not be attached to a numbered software release; instead, the Oracle Retail

document will simply be replaced on the Oracle Technology Network Web site, or, in the case of Data Models, to the applicable My Oracle Support Documentation container where they reside.

This process will prevent delays in making critical corrections available to customers. For the customer, it means that before you begin installation, you must verify that you have the most recent version of the Oracle Retail documentation set. Oracle Retail documentation is available on the Oracle Technology Network at the following URL:

http://www.oracle.com/technetwork/documentation/oracle-retail-100266.html

An updated version of the applicable Oracle Retail document is indicated by Oracle part number, as well as print date (month and year). An updated version uses the

same part number, with a higher-numbered suffix. For example, part number E123456-02 is an updated version of a document with part number E123456-01.

If a more recent version of a document is available, that version supersedes all previous versions.

Oracle Retail Documentation on the Oracle Technology Network

Documentation is packaged with each Oracle Retail product release. Oracle Retail product documentation is also available on the following Web site:

http://www.oracle.com/technology/documentation/oracle_retail.html

(Data Model documents are not available through Oracle Technology Network. These documents are packaged with released code, or you can obtain them through My Oracle Support.)

Documentation should be available on this Web site within a month after a product release.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1
Introduction

Oracle Retail Demand Forecasting (RDF) is a statistical and promotional forecasting solution. It uses state-of-the-art modeling techniques to produce high quality forecasts with minimal human intervention. Forecasts produced by the RDF 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 RDF can provide tremendous benefits to businesses.

A connection from RDF to Oracle Retail's Advanced Retail Planning and Optimization (ARPO) solutions is built directly into the business process by way of the automatic approvals of forecasts, which may then be fed directly to any ARPO solution. This process allows you to accept all or part of a generated sales forecast. Once that decision is made, the remaining business measures may be planned within an ARPO solution such as Merchandise Financial Planning, for example.

Forecasting Challenges and RDF 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. Oracle Retail'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, Croston's 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 and the complexity of the forecast model are evaluated in order to determine the most accurate forecast method. You simply select the AutoES forecast generation method and the system finds the best model.

Overcoming Data Sparsity Through Source Level Forecasting

It is a common misconception in forecasting that forecasts must be directly generated at the lowest levels (final-levels) of execution. 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 the sales data from a final-level up to a higher level (source-level) in the hierarchy in which demand patterns can be seen, and then generate a forecast at this source-level. After a forecast is generated at the source-level, the resulting data can be allocated (spread) back down to the lower-level, based on the lower-level's (final-level) relationship to the total. This relationship can then be determined through generating an additional forecast (interim forecast) at the final-level. RDF can be set up to dynamically generate a profile based on the interim forecasts. Also, a non-dynamic profile can be generated and approved in Curve to be used as this profile. It is this profile that determines how the source-level forecast is spread down to the final-level. For more information on Curve, refer to Chapter 12, "Curve".

Figure 1–1 Final Level versus Source Level Forecasting

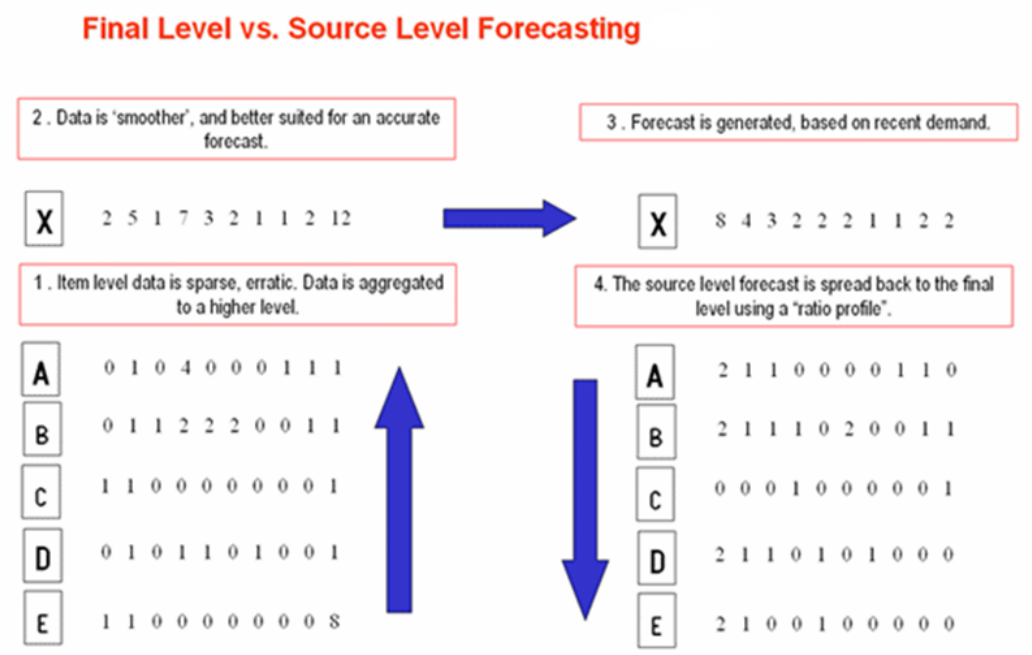

Some high-volume items may possess sufficient sales data for robust forecast calculations directly at the final-forecast level. In these cases, forecast data generated at an aggregate level and then spread down to lower-levels can be compared to the interim forecasts run directly at the final-level.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

Your RDF system may include multiple final-forecast levels. Forecast data must appear at a final-level for the data to be approved and exported to another system for execution.

Forecasting Demand for New Products and Locations

RDF also forecasts 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 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, refer to the Forecast Like-Item & Sister-Store Workbook section.

Managing Forecasting Results Through Automated Exception Reporting

The RDF end user may be responsible for managing the forecast results for thousands of items, at hundreds of stores, across many weeks at a time. The Oracle Retail Predictive Application Server (RPAS) provides users with an automated exception reporting process (called Alert Management) that indicates to you where a forecast value may lie higher than or lower than an established threshold, thereby reducing the level of interaction needed from you.

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.

For more information on the Alert Manager, refer to the *Oracle Retail Predictive Application Server User Guide*.

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 might have a significant impact on a product's sales history, as can fluctuating holidays such as Easter.

Using Promotional Forecasting (an optional, add-on module to RDF), promotional models of forecasting can be developed to take these and other factors into account when forecasts are generated. Promotional Forecasting 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 is similar.

This type of advanced forecasting identifies the behavioral relationship of the variable you want to forecast (sales) to both its own past and explanatory variables such as promotion and advertising.

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

Managing and Forecasting for Floating Events

For some retailers, an RDF implementation with causal forecasting is necessary. Whenever price discount, marketing information (TV, flyer) or space information (front isle) need to be factored in the forecast, causal forecasting is the right tool.

However, in RDF implementations without the causal capability, some floating events still exist. For instance, an increase in sales leading to Easter will happen even if there is no promotion specifically associated with Easter (chocolate bunnies sell more during Easter even though they are not specifically promoted). And because it does not happen in the same time period every year, the Easter spike will appear 'randomly' any time in March or April. The spike is baked in the seasonality of each item/store, making its prediction, timing and magnitude, inaccurate at best.

RDF is able to handle such situations, not only thru its causal capabilities, but also for baseline forecasting.

To do that, a set of events has to be defined. This usually happens at configuration time. Then, a user has to flag the periods where an event was active in the past (history) and in the future (forecast horizon). Once the dates are set up, the effect of the event is estimated based on past instances, and applied in the forecast horizon if the flag is active.

For the user interaction, that is the managing the periods when events were active, RDF provides the Floating Event Administration workbook. Its single worksheet lists events and you have the ability to flag periods when certain product/location combinations are active during an event.

53 Week Calendar

For the majority of retailers, the business is managed using a calendar (364 days organized into 13 week quarters) that periodically includes an extra 53rd week so that the year end stays in about the same time of the year. It is useful to have some control over how this 53rd week will be managed within the forecasting system's time dimension. Management of this issue causes customers the pain, time and cost of configuring their data every few years that this happens.

The problem described has two implications. The first case is when two years – each with 52 weeks – of historical sales are available, and the retailer needs to forecast for the following year, which has 53 weeks. The second case is when one of the years of historical sales has 52 weeks, and the other has 53 weeks.

The correction for the extra week happens as part of generate, in particular when the baseline is written out. This is necessary, such that all additional effects (promo, price change, demand transference) are layered on top of the baseline.

The information RDF needs to handle a 53rd week is the name of the measure that indicates which week is the extra week. The measure is loaded or populated through user input, and it is stored in the Forecast Administration parameter called **Extra Week Indicator Data Source**. On the same view (Advanced Final and Source Level Parameters) there is also the measure **Extra Week Interpret Method** that indicates how to calculate the forecast value for a week that was flagged as 53rd or extra week.

Forecasting Returns

RDF has the ability to calculate returns based on the unconstrained demand that it generates. The approach is based on the following assumptions:

- All returned items are re-sellable; in other words there is no damage or decomposition.
- Each week, the return is solely based on a previous week's sale, instead of weighting several weeks.

You need to enable the returns forecast generation, as well as input two parameters:

- The percentage of the merchandise that was returned
- The number of periods, t, from when the merchandise was sold until it was returned.

Given this information, the formula for the returns at time period t, is:

```
Returns(t) = Forecast(t- t) * percentage
```

In Figure 1–2, the time shift is specified as 1 and the return rate is 5%. Given the forecast (Fcst), the returns are calculated as:

```
Returns(t) = Forecast(t-1) * 5%
```

Note: The returns for the first week of forecast are based on the latest week of historical sales (HS), since no previous forecast value is available.

Figure 1–2 Calculating Returns

Logic	HS 107	HS 108	Forecast Start	Week 1	Week 2	Week 3	Week 4	Week 5
Fcst/Sls	50	50		50	80	80	100	60
Return = $Fcst_{shifted\_time}$*percentage	NA	NA		2.5	2.5	4	4	5

The parameters are set in the Forecast Administration Workbook and the Forecast Maintenance Workbook The returns forecast can be viewed in the Forecast Approval Workbook.

Oracle Retail Demand Forecasting Architecture

This section describes The Oracle Retail Predictive Application Server and RDF.

The Oracle Retail Predictive Application Server and RDF

The RDF application is a member of the Advanced Retail Planning and Optimization Suite (ARPO), including other solutions such as Merchandise Financial Planning, Item Planning, Category Management, and Advance Inventory Planning. The ARPO solutions share a common platform called the Oracle Retail Predictive Application Server (RPAS). RDF leverages the versatility, power, and speed of the RPAS engine and user-interface. Features such as the following characterize RPAS:

- Multidimensional databases and database components (dimensions, positions, hierarchies)
- Product, location, and calendar hierarchies
- Aggregation and spreading of sales 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
- Charting and graphing capabilities

More details about the use of these features can be found in the *Oracle Retail Predictive Application Server User Guide* and online help provided within your RDF solution.

Oracle Retail Demand Forecasting Workbook Groups

In addition to the standard RPAS Administration and Analysis workbook groups, there are several groups that are associated with the RDF solution which may include: Forecast, Promote, Curve, or any Planning and Optimization suite application (available modules are based upon licensing agreement).

Figure 1–3 Forecast Workbooks Tab

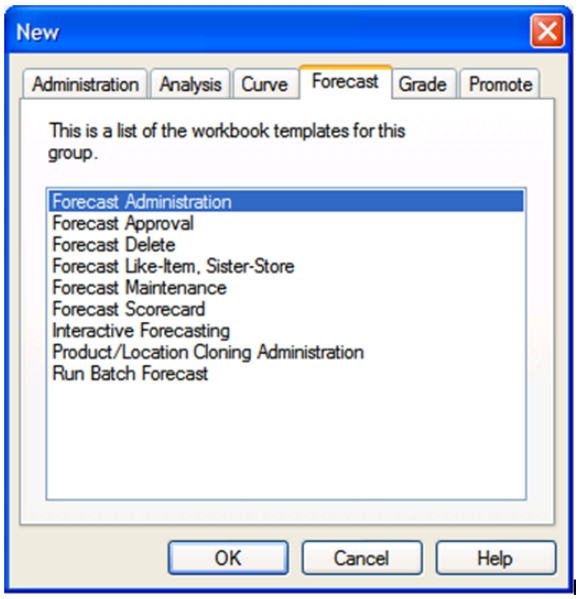

Forecast

The Forecast module refers to the primary RDF functionality and consists of the workbooks, measures, and forecasting algorithms that are needed to perform time-series forecasting. This includes the:

- Forecast Administration Workbook
- Forecast Maintenance Workbook
- Forecast Like-Item & Sister-Store Workbook
- Forecast Approval Workbook
- Forecast Batch Run Workbook
- Forecast Delete Workbook
- Interactive Forecasting Workbook
- Forecast Scorecard Workbook

The Forecast module also includes the batch forecasting routine and all of its component algorithms.

Promote (Promotional Forecasting)

The Promote module consists of the workbooks 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 Promotion Maintenance, Promotion Planner and Promotion Effectiveness workbooks.

Curve

The Curve module consists of the workbooks and batch algorithms that are necessary for the creation, approval, and application of profiles that may be used to spread

source-level forecasts down to final-levels as well to generate profiles, which may be used in any RPAS solution. The types of profiles typically used to support forecasting are: store Contribution, Product, and Daily profiles. These profiles may also be used to support Profile-Based Forecasting; however; Curve may be used to generate profiles that are used by other ARPO solutions for reasons other than forecasting. Profiles Types include Daily Seasonal, Lifecycle, Size, Hourly, and User-Defined profiles. For more information on the Curve workbooks and worksheets, refer to Chapter 12, "Curve.".

RDF Solution and Business Process Overview

Oracle Retail has designed a forecasting solution separate from replenishment, allocation or planning. In order to provide a single version of the truth, it is crucial to free up your time and supply the tools to focus on the analysis of forecast exceptions, historical data, and different modeling techniques. This empowers you to make better decisions, thus improving overall accuracy and confidence in the forecast downstream.

RDF and the Oracle Retail Enterprise

Within the Oracle Retail Enterprise, Oracle Retail Merchandising System (RMS) supplies RDF with Point-of-Sale (POS) and hierarchy data that is used to create a forecast. Once the forecast is approved, it is exported to RMS in order to calculate a recommended order quantity. The RDF libraries can be used in any RPAS solution to create forecasts to support merchandise, financial, collaborative, and price planning processes.

GA integration scripts export RDF output to solutions including RMS, RPO, RO, and AIP. For other RPAS based solutions, integration scripts need to be customized by the implementer to reflect each retailer's specific needs.

See Figure 1–5 for an overview of RDF and the Oracle Retail Enterprise.

Figure 1–4 RDF and the Oracle Retail Enterprise

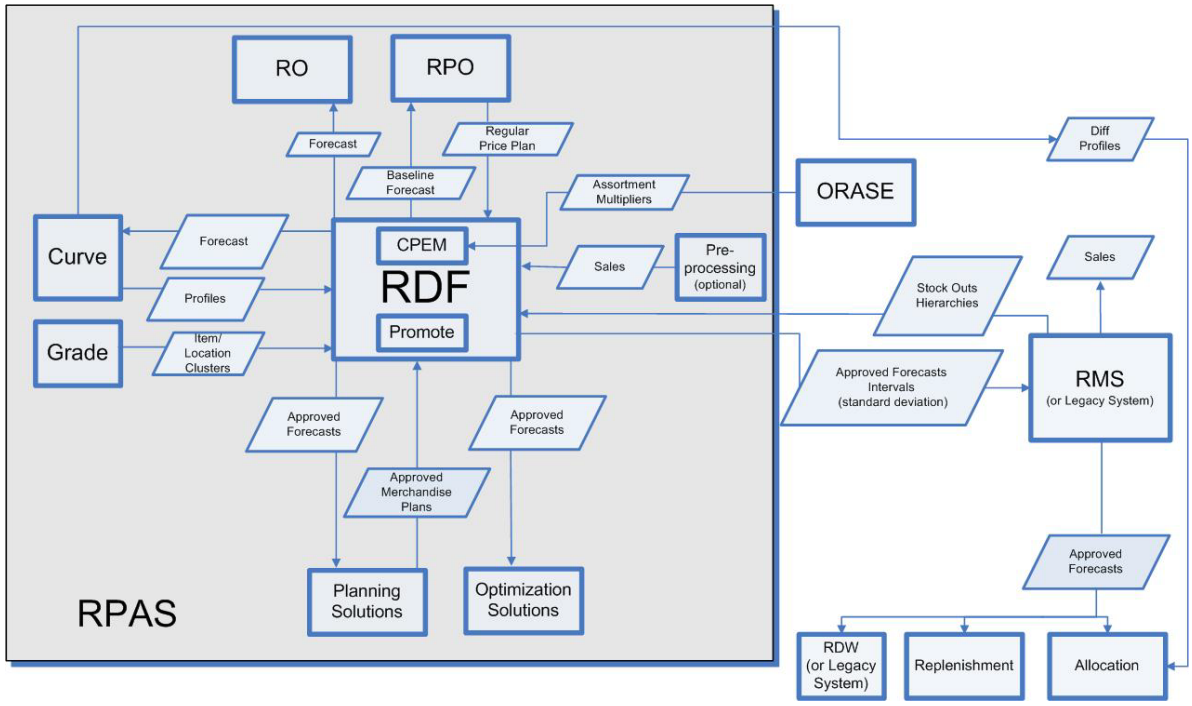

RDF Primary Workflow

There are a number of core super-user/end-user forecasting steps in the RDF workflow that are essential for producing accurate forecasts for the millions of item and location combinations that exist in a domain.

RDF Batch Flow Process

RDF's batch flow process follows six phases as shown in Figure 1–6.

Figure 1–5 Batch Flow Process in RDF

2

Oracle Retail Demand Forecasting Methods

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

Forecasting Techniques Used in RDF

RDF 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 uses 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 an Oracle Retail 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 Oracle Retail, uses a sales plan as the starting point that is adjusted based on observed data. This method fills a gap in standard time series forecasting when new, short-lifecycle or products with significant lifecycles are being forecast.

Prediction Intervals

Predictions 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 solutions 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, Oracle Retail 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 trade-off between the model's fit over the historic data and its complexity. A description of the competing models used within AutoES is described in "Exponential Smoothing (ES) Forecasting Methods" on page 2-9. 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 (item/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 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 is run in the future, these promotional forecasting techniques can better predict demand in the future.

Time Series (Statistical) Forecasting Methods

This section describes those techniques within RDF that generate forecasts directly from only a single time series. Generally the time series provided is past sales history for a given item/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 department/chain, and so on).

The following topics present fundamentals of the RDF 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 previously, and they do this by fitting quantitative models to statistical patterns from historical data. Put simply, the better the history of the variable being forecast, the stronger these statistical patterns are. Increased forecast accuracy depends on the strength of these patterns in relation to background irregularities.

RDF 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:

- Average (Simple Moving Average)
- Bayesian Information Criterion (BIC)
- Exponential Smoothing (ES) Forecasting Methods - Automatic Exponential Smoothing (AutoES)
  - Simple Exponential Smoothing
  - Croston's Method
  - Simple/Intermittent Exponential Smoothing
  - Holt Exponential Smoothing
  - Multiplicative Winters Exponential Smoothing

    Oracle Winters
  - Additive Winters Exponential Smoothing
  - Seasonal Exponential Smoothing (SeasonalES)
  - Seasonal Regression

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.

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.
- Can be generated with little historical data.
- Are very fast to generate.

Typically, moving average forecasts are generated at the final-forecast level (for example, item/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.

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 period (between history start date and forecast start date). 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 zero (0)). These changes tend to favor the

seasonal models to a slightly higher degree that improves the forecasts on retail data, especially for longer forecast horizons.

AutoES Flowchart

The following procedure outlines the processing routine steps that the system runs through to evaluate each time series set to forecast using the AutoES method. See Figure 2–1, "AutoES Flowchart".

1. Filter all leading zeros in the input data that is within the training window.
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.
3. Does the time series contain enough relevant data to generate a forecast? If yes, generate a forecast and statistics using the SimpleES method and move on to Step 4. If no, do not forecast and go to the next time series.
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 9.
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.
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.
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.
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.
9. Compare all candidate forecasts using BIC Criterion.
10. Return the corresponding forecast and statistics for the system-selected forecast method and move on to the next time series.

Figure 2–1 AutoES Flowchart

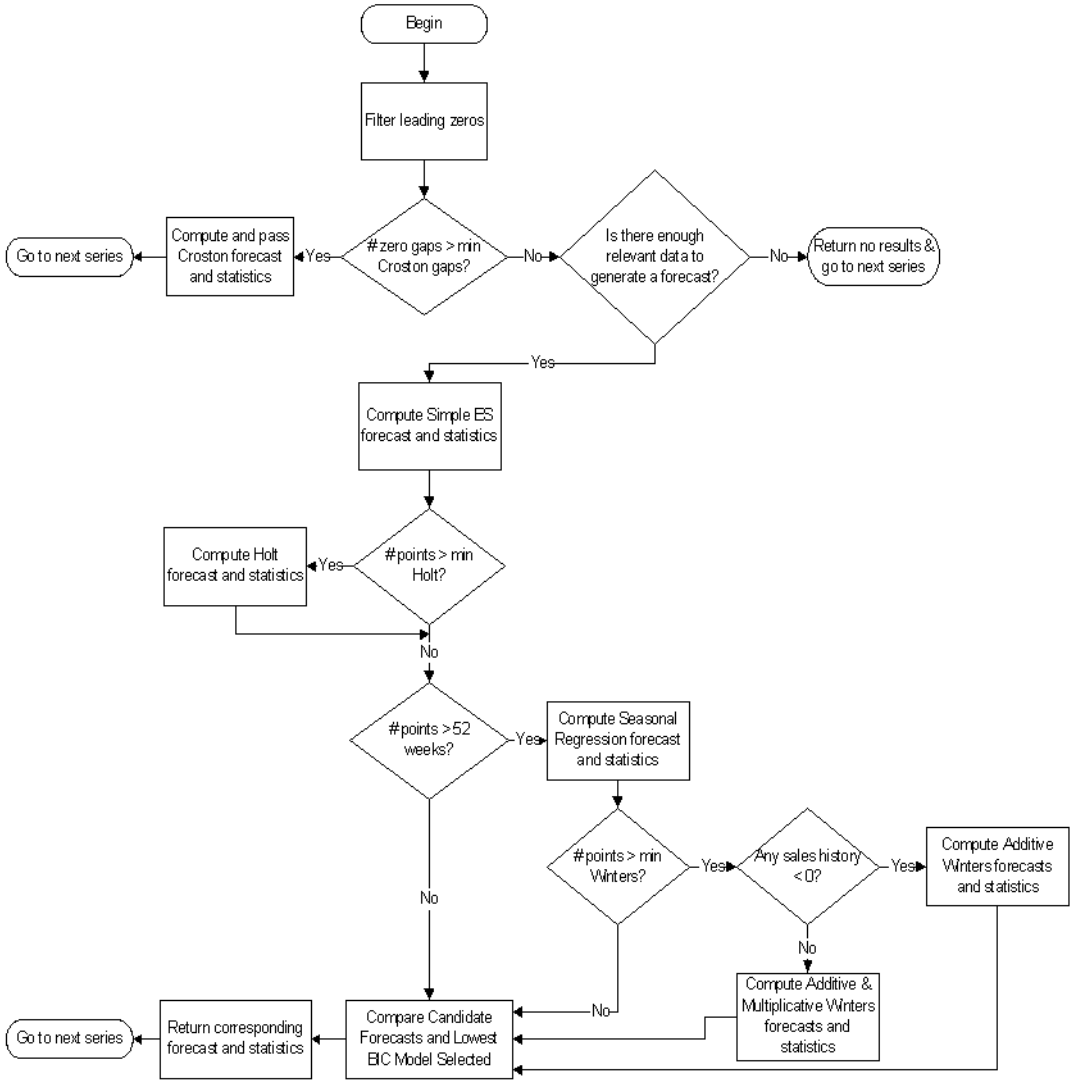

Automatic Forecast Level Selection (AutoSource)

This section describes how the Automatic Forecast Level Selection (AutoSource) 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 Figure 2–2, a final-forecast is generated by:

1. Aggregating up from the base level to the source-level
2. Generating a source-level forecast
3. Spreading the source-level forecast down to the final-level

For example, assume base-level sales data is at the item/store level, the final-forecast level is at the item/store level, and the candidate source generation level is at the style/store level. In this case, base-level sales data is aggregated from the item/store level up to the style/store level. A style/store forecast is generated, and the forecast data is spread back down to the item/store level. This forecast represents the final-forecast.

Figure 2–2 Forecast Level Selection Process

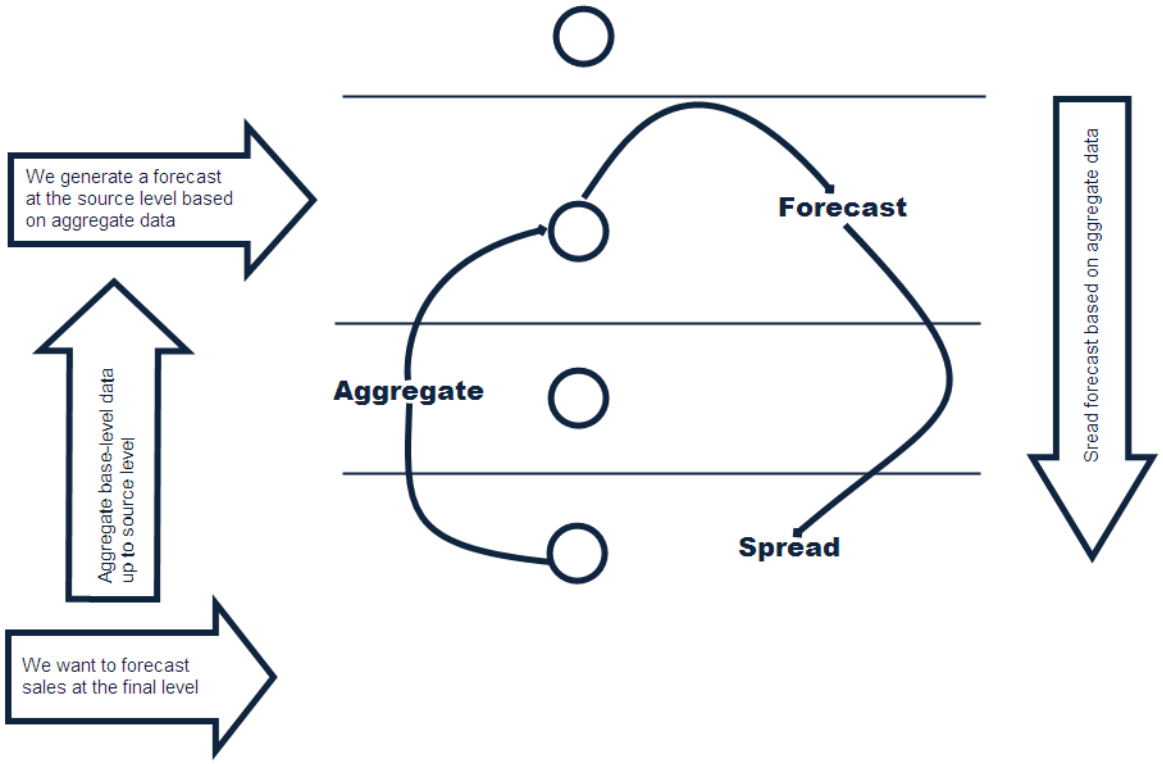

Determining the Best Source Level Forecast

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), the source generation level that generated that better PAE becomes the new best level.

Status and 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 meaning that 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:

- Generating forecasts at each unique aggregation level
- Using the train-test approach to evaluate the percent absolute error statistics for each

One or more of these sub-tasks is 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 (AutoSource) 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 eight 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 RDF appears in the Optimal Source Level measure on this worksheet. You can enable the use of this level by placing a check mark in the Pick Optimal Level measure for that product/location. The absence of a check mark in this measure causes the system to default to the Default Source Level or the Source Level Override value if this has been set by you.

Exponential Smoothing (ES) Forecasting Methods

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 Exponential Smoothing
- Croston's Method
- Simple/Intermittent Exponential Smoothing
- Holt Exponential Smoothing
- Multiplicative Winters Exponential Smoothing
    - Oracle Winters
- Additive Winters Exponential Smoothing
- Seasonal Exponential Smoothing (SeasonalES)
- 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 is applied to noisier data). The final selection between the resulting models is made according to a performance criterion that involves a trade-off 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, the AutoES forecasting process finds 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 generally projects better forecasts.

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.

Figure 2–3 is an example of a forecast in which data seems to be un-trended and un-seasonal; note the flat appearance of the forecast.

Figure 2–3 Simple Exponential Smoothing

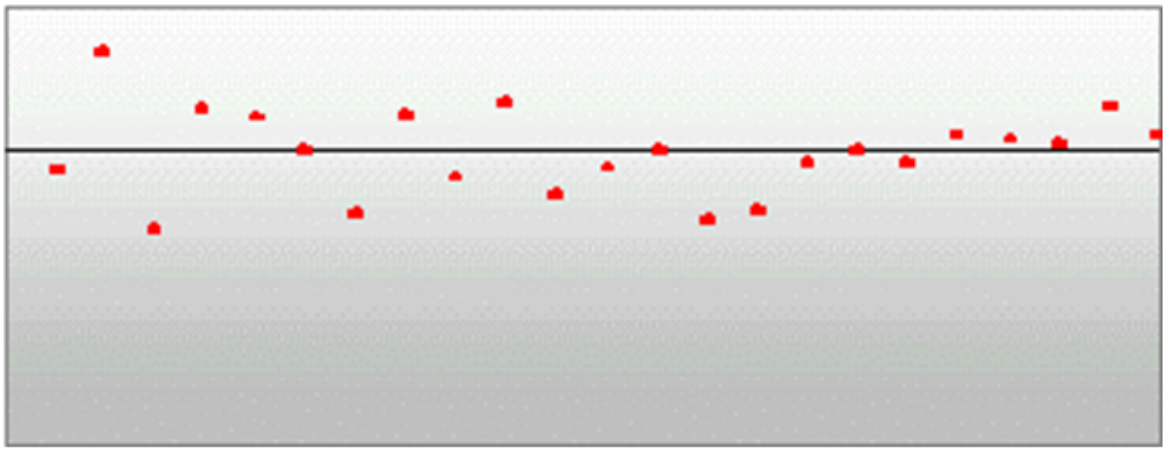

Croston's Method

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:

- Magnitude series
- 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.

Figure 2–4 shows a sales history of data where the demand for a given period is often zero.

Figure 2–4 Croston's Method

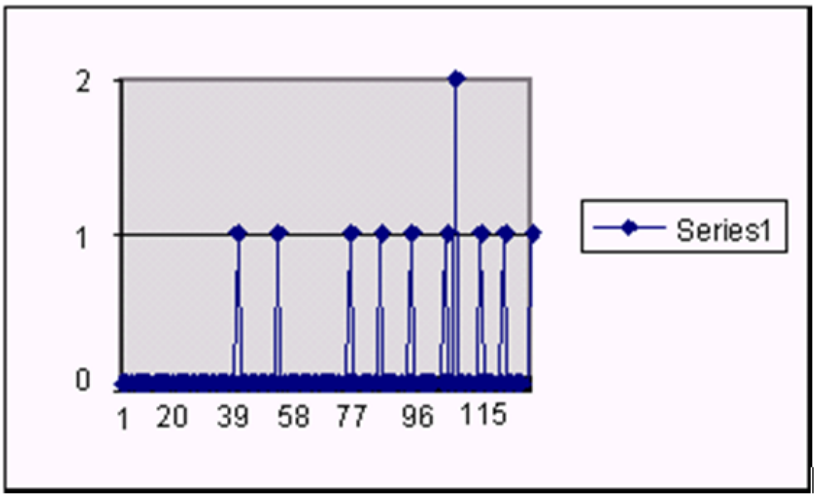

Simple/Intermittent Exponential Smoothing

This method is a combination of the SimpleES and Croston's (IntermittentES) methods. The SimpleES model is applied to the time series unless a large number of transitions from non-zero sales to zero data points are present. In this 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.

When this forecasting method is selected, the forecasts are seen as trending either up or down, as in Figure 2–5.

Figure 2–5 Holt Exponential Smoothing

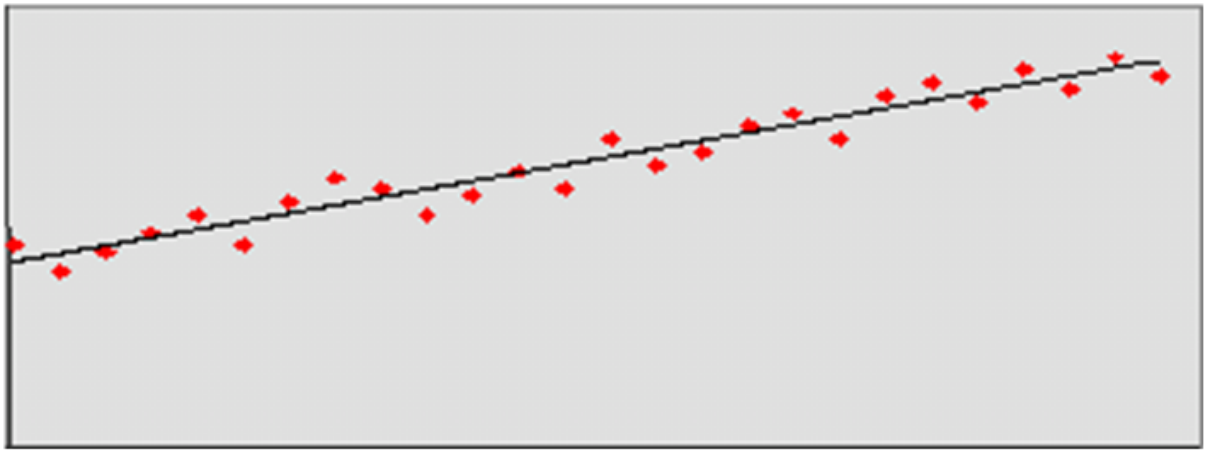

Multiplicative Winters Exponential Smoothing

Overall, a forecast point estimate is evaluated as:

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

a function of level, trend, seasonality, and trend dampening factor.

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

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

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

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

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

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

Oracle Winters calculates initial seasonal indices from a baseline Holt forecast. Seasonal indices, level, and trend are then updated in separate stages, using Winter's model as a basis for the updates.

Oracle Winters

Oracle Winters is the current seasonal forecasting approach, which uses a combination of Winters approach and decomposition. Decomposition allows level and trend to be optimized independently while maintaining a seasonal curve.

From sufficient data, RDF extracts seasonal indexes that are assumed to have multiplicative effects on the deseasonalized series.

Note: The component describing the deseasonalized values, which is multiplied by the seasonal index:

$$\hat{S}_t(h)$$

is the Holt model previously described.

In this case, three parameters are used to control smoothing of the components of level, trend, and seasonality.

When the Multiplicative Seasonal forecasting method is selected, the forecasts tend to look squiggly, as shown in Figure 2–6.

Figure 2–6 Multiplicative Winters Exponential Smoothing

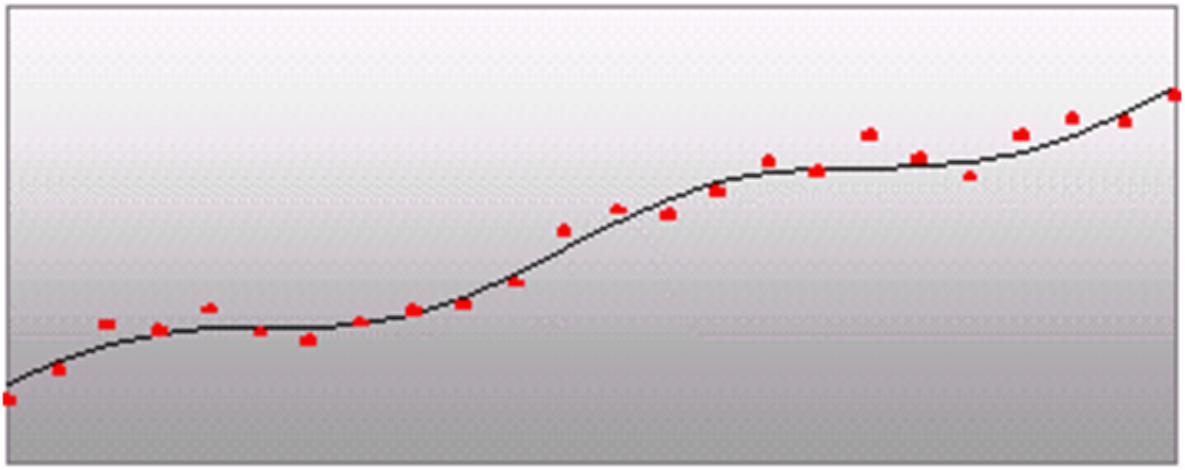

Additive Winters Exponential Smoothing

Overall, a forecast point estimate is evaluated as:

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

a function of level, trend, seasonality and trend dampening factor.

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

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

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

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

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

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

Oracle Winters is a Winters-based decomposition approach to update the level, trend, and Seasonal Indexes.

Refer to Multiplicative Winters Exponential Smoothing in this document for a description of each of the Forecasting Approaches.

Seasonal Exponential Smoothing (SeasonalES)

A combination of several seasonal methods. This method is generally used for known seasonal items or forecasting for long horizons. This method lets the Multiplicative Seasonal and Additive Seasonal models compete and picks the one with the better fit. If less than two years of data is available, a Seasonal Regression model is used. Even if there is too little data to create a seasonal forecast (in general less than 52 weeks, or a full sales cycle) the system still attempts to generate a seasonal forecast if at all possible. The sales of last year are time shifted by 52 weeks, and become the forecast for the corresponding periods. If the third week in 2020 is in the forecast horizon, and sales for the third week in 2019 are available, they will become the forecast. If not, the forecast will be generated by running SimpleES with the available (less than 52 weeks) history. For this particular case, the method picked is Seasonal Regression.

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 RDF (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 RDF 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 it 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 learns that shift and appropriately weigh 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 forecast 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 is selected if it best fits the data.

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), whenever a negative slope is detected, the regression is rerun with the intercecpt fixed to zero. This guarantees that a positive slope is calculated and thus a more logical forecast is 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, the method rejects itself out of hand and allows one of the other competing methods to provide the forecast.

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 is deseasonalized using the profile and then fed to Simple method.
3. RDF allows a max alpha between 0 and 1.
4. The Simple forecast is re-seasonalized using the profiles.

Profile-Based Method and New Items

The Profile-based forecasting method can be successfully used to forecast new items. In order to do that, we need to have a profile (which can be copied from an item 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 it 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.

Note: The Profile-based forecasting method is intended for a final forecasting level and not a source level.

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 lifecycle 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.

- 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 takes 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 (Bayesian Alpha), which you can set between zero and infinity.

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. The following are major problems in automatically developing these forecasts:

- 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 into a forecasting model. For example, the overall sales level of the product, how quickly the product takes off, how the product's sales is affected by planned promotions.

Using a Bayesian approach, a short life-cycle forecasting algorithm was 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, Oracle Retail assists clients in obtaining good forecasts for future demands for their products based upon historical sales data and available causal information. Depending on the information available, Oracle Retail's software supports various forms of exponential smoothing and regression-based forecasting. Frequently, clients already have some expectations of future demands in the form of sales plans. Depending on 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 was developed that merges a customer's 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 versus 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, and so on) produce adequate results. The problem arises when attempting to

forecast products with little or no history. In such instances, expert knowledge (generally in the form of sales plans) is required. 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? In answering that question (in a particular scenario), suppose that we have determined that 13 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 a hard-edge boundary existing; 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.

Bayesian Algorithm

Bayesian forecasting is primarily designed for product/location positions for which a plan exists. These product/locations can be, but are not limited to: items new in the assortment, fashion items, and so on. 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 assumes a plan exists and equals zero and acts accordingly.
3. The time period of interest for the Bayesian algorithm starts with the first non-zero value of the plan or the history start date (whichever is more recent), and ends at the end of the forecast horizon.

LoadPlan Forecasting Method

LoadPlan Forecasting Method copies the measure that was specified as Data Plan in the Forecast Administration Workbook into the Forecast measure. This method does not generate confidence and cumulative intervals when it is the final level method and no source level is specified. This method only generates confidence and cumulative intervals when a source level is specified and the source and final levels are the same.

Copy Forecasting Method

Copy Forecasting Method copies the measure that was specified as Forecast Data Source in the Forecast Administration Workbook into the Forecast measure. This method does not generate confidence and cumulative intervals.

Causal (Promotional) Forecasting Method

Causal, or promotional, forecasting requires four input streams:

- Time Series Data
- Historical Promotional Calendar
- Baseline Forecasts
- Future Promotional Calendar

Promote decomposes the problem of promotional forecasting into two sub-tasks:

- Estimating the effect that promotions have on demand

- Applying the effects on the baseline forecasts

To accomplish the first task, a stepwise regression sub-routine is used. 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.

A promotion variable can represent an individual promotion or a combination of overlapping promotions. An individual promotion represents the case where for a particular time period a single promotion is enabled. Overlapping promotions represents the case where there are two or more promotions that happen at the same time period at the same location for the same product. Causal Forecasting Method can calculate not only each individual promotion effect, but also the overlapping promotions effects.

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, knowing the effects and the baseline forecasts, we can generate a promotional forecast by applying the effects wherever the promotion or the overlapping promotion are active in the future.

The causal forecasting process has been simplified by first estimating the effects of promotions. Then the effects are applied on top of a baseline that is created externally from the causal forecasting process. For example, the baseline can be a loaded measure, or it can be generated in RDF using AutoES and source-level techniques.

It should be noted that just because promotional forecasting is selected, it does not necessarily imply that a promotional forecast results. In some instances, no promotional variables are found to be statistically significant. In these cases, the forecast ends up equivalent to standard time series forecasts.

If you want to force certain promotional variables into the model, this can be managed through forecasting maintenance parameters.

The Oracle Retail 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 is 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 items 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 is most likely attributed to one or the other event.

Causal Forecasting with Baselines

Rather than use a sales history that may not have sufficient or accurate data, users can load a baseline into the RDF Causal Engine instead. Baselines are often generated using data that is rolled up to a higher dimension than item/store, providing a greater depth of data and hence a less-noisy sales history. The sales data used to generate the baseline can be corrected for out-of-stock, promotional, and other short-term event information using Preprocessing. These baselines are then spread back to the item/store level and then loaded in the RDF Causal Engine. The Engine uses the

baseline along with the historic promotional data and future planned promotions data to create the system forecast, which is the baseline with the lifts, which were calculated from the promotional data, applied on top.

Figure 2–7 Baseline Process

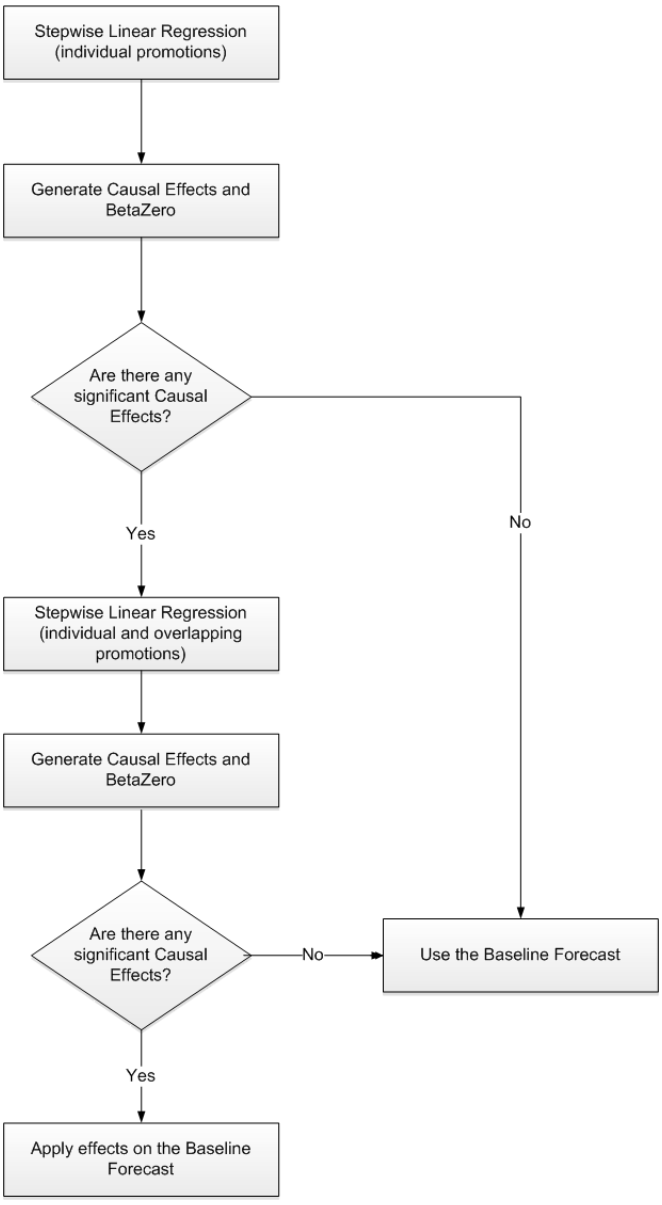

Causal Forecasting Using the Baseline Process

When a baseline is used, the AutoES binary code 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 an 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.
6. If the regression finds no significant promotional variables, the casual method is considered to have failed to fit. In this case, the forecast equals the baseline.
7. The fit at time *t*, *fit(t)*, is defined in terms of b_o, the intercecpt of the regression, b_1, the effect corresponding to promotional variable *i*, and $p_i(t)$, is the value of promo variable *I*, in time *t* as:

   $fit(t) = b_o + \hat{A}\ b_o * p_i(t)$
8. The forecast is obtained by re-causalizing the baseline. This is done by adding back the causal effects for the product/location/time positions, where the corresponding promo variables are on in the forecast region.
9. The binary writes the winning promotional variables effects back to the database.
10. The selected model is recorded in the database.
11. The binary records the forecast and the baseline in the database.

Causal Forecasting for Short-lifecycle Items

To forecast short-lifecycle promotional items, Causal deprices, depromotes, and smoothes the forecasting data source to generate the short-lifecycle forecast causal baseline.

In essence, this is the causal forecasting process where the generation of the baseline and the estimation of the promotional lifts are modified for items with short lifecycle.

First, the baseline is generated. Then the forecast data source is depriced, depromoted and smoothed. The resulting feed is aggregated and then spread down using rate of sales as a profile to create a lifecycle curve. The lifecycle curve is shifted and stretched or shrunk to fit the new season length. This curve represents the pre-season baseline forecast. In season, the pre-season forecast serves as a forecast plan to the Bayesian forecasting method. The Bayesian forecast is the causal baseline for short-lifecycle items.The next step is to calculate the promotional lifts. To do this, the promotional lifts are filtered from the historical sales and applied on top of the item's rate of sale. The resulting measure is regressed against the promotional variables to determine the promotion effects. Finally, the promotion effects are applied on top of the short-lifecycle causal baseline to generate the final-forecast.

Causal Forecasting at the Daily Level

The causal forecasting at the daily level is calculated by spreading the weekly causal forecast down to day. The spreading utilizes causal daily profiles, thus obtaining a causal forecast at the day granularity.

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 item/store 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 item/store, the effects are replicated down to item/store 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, the effects from the final-level is used.
4. Daily profiles are calculated using the Curve module. Since as much history as possible is used and is averaged over seven days, it's assumed that these profiles are de-causalized. The de-causalized daily profiles capture the day-of-week effect and should be quite stable.
5. Causal effects are applied to the daily profiles. The profiles are multiplied by the causal effects and then the profiles have to renormalize.

Example

For every item/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:

Mon	Tues	Wed	Thu	Fri	Sat	Sun	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:

Mon	Tues	Wed	Thu	Fri	Sat	Sun
		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 is held in a future week (w36), but only on Thursday:

Mon	Tues	Wed	Thu	Fri	Sat	Sun
			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.

- Update the week-to-day profile of w36 so that the weight of Thursday is doubled (the multiplicative factor is 2):

Mon	Tues	Wed	Thu	Fri	Sat	Sun	Week
10%	10%	10%	**20%**	20%	30%	10%	**110%**

- Normalize the profile for w36:

Mon	Tues	Wed	Thu	Fri	Sat	Sun	Week
9%	9%	9%	18%	18%	27%	9%	100%

Finally, spread the forecast of w36 using the normalized profile.

Final Considerations Regarding Causal Forecasting

In the Oracle Retail 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 a model is fit 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 it 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. 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 would otherwise not have the causal variables on at the same time. An alternate solution is whenever a causal effect cannot be computed because of lack of significant data. An averaged effect from another time-series in the same aggregation class is going to be used instead.

3

Setting Forecast Parameters

The Forecast workbook group allows you to perform functions related to statistical time series forecasting. This chapter provides information on defining and maintaining the parameters that govern the generation of forecasts in RDF.

Forecast Administration Workbook

This section describes the components of the Forecast Administration workbook.

Basic versus Advanced Tabs

Forecast Administration is the first workbook used in setting up RDF to generate forecasts. It provides access to forecast settings and parameters that govern the whole domain (database). These settings and parameters are divided into two areas, accessed through the Basic and Advanced tabs beneath main toolbar.

The Basic tab is used to establish a final-level forecast horizon, the commencement and frequency of forecast generation, and the specification of aggregation levels (Source Levels) and spreading (Profile) methods used to yield the final-level forecast results.

The Advanced tab is used to set default values for parameters affecting the algorithm and other forecasting techniques used to yield final-level and source-level forecasts, thus eliminating the need to define these parameters individually for each product and location. 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.

Final versus Source Level Forecasts

Often, forecast information is required for items at a very low level in the hierarchy. However, problems can arise 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-level or final-level dimension based on the lower-level's relationship to the total.

In order to spread this forecast 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 interim 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 other systems. The Source Level worksheet represents the default values for forecast parameters at the more robust aggregate (source) level.

Forecasting Methods Available in RDF

A forecasting system's main goal is to produce accurate predictions of future demand. The RDF 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. Some of these methods may not be visible in your solution based on configuration options set in the RPAS Configuration Tools. More detailed information on these forecasting algorithms is provided in Chapter 2, "Oracle Retail Demand Forecasting Methods."

Average

RDF uses a simple average model to generate forecasts.

Moving Average

RDF uses a simple moving average model to generate forecasts. You can specify a Moving Average Window length.

IntermittentES

RDF 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 transitions from non-zero sales to zero data points (that is, intermittent demand data).

The original time series is split into a Magnitude and Frequency series, and then the SimpleES 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.

TrendES

RDF 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 and then provides forecast point estimates by combining an estimated trend and the smoothed level at the end of the series. For instance, where the forecast engine cannot produce a forecast using the TrendES method, the Simple/IntermittentES method is used to evaluate the time series.

Causal

Causal is used for promotional forecasting and can only be selected if Promote is implemented. 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 and source-level forecasting are used 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. For instances where the forecasting engine cannot produce a forecast using the Causal method, the system evaluates the time series using the SeasonalES method.

To forecast short-lifecycle, promotional items, Causal deprices, depromotes, and smoothes the forecasting data source to generate the short-lifecycle forecast causal baseline. The promotion effects are calculated using the same Stepwise Regression code as Causal. However, the promotion lifts are first filtered from the sales and applied on top of the item's rate of sale. The resulting measure is regressed against the promotional variables to determine the promotion effects. Finally, the promotion effects are applied on top of the short-lifecycle baseline to generate the final-forecast.

No Forecast

No forecast is generated for the product/location combination.

Bayesian

Useful for short-lifecycle 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 lifecycle) 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. A Data Plan must be specified when using the Bayesian method. For instances where the Data Plan equals zero (0), the system evaluates the time series using the SeasonalES method.

Profile-based

RDF generates a forecast based on a seasonal profile that can be created in RPAS or legacy system. Profiles can also be copied from another profile and adjusted. Using historic data and the profile, the data is deseasonalized and then fed to the SimpleES method. The Simple forecast is then re-seasonalized using the profiles. A Seasonal Profile must be specified when using the Profile-Based method. For instances where the Seasonal Profile equals zero (0), the system evaluates the time series using the SeasonalES method.

Components

The Components forecast method multiplies pre-calculated baseline, promotional lifts, and regular price lifts to generate the forecast. These components are specified in Forecast Administration Workbook, Advanced Settings tab, Final Level Parameters worksheet. The following table lists the components and their referenced measures.

Component	Measure
Baseline	Components - Baseline
Regular Price Lifts	Components - Regular Price Lifts
Promotional Lifts	Components - Promotional Lifts
Interval	Components - Baseline Interval

Intervals are not generated and need to be provided. This method is useful when the forecast components are available since the forecast can be quickly generated.

Automatic Exponential Smoothing (AutoES)

RDF 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:

- SimpleES
- Croston's Method
- Simple/IntermittentES
- HoltES
- Multiplicative Seasonal
- Additive Seasonal
- SeasonalES
- Seasonal Regression

The final selection between the models is made according to a performance criterion (Bayesian Information Criterion) that involves a trade-off between the model's fit over the historic data and its complexity.

SimpleES

RDF uses a simple exponential smoothing model to generate forecasts. SimpleES 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. For more information, see Simple Exponential Smoothing.

Croston's Method

Croston's Method is used when the input series contains a large number of transitions from non-zero sales to zero data points (that is, intermittent demand data). For more information, see Croston's Method.

Simple/IntermittentES

Simple/IntermittentES is a competition between the SimpleES and IntermittentES methods. This method applies the SimpleES model unless a large number of transitions from non-zero sales to zero data points are present, in which case the Croston's model is applied. For more information, see Simple/Intermittent Exponential Smoothing

HoltES

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. For more information, see Holt Exponential Smoothing.

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. For more information, see Multiplicative Winters Exponential Smoothingwhich includes Oracle Winters.

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). For

more information, see Additive Winters Exponential Smoothing.

SeasonalES

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 too many zeros are present in the data, in which case the Additive Winters model of exponential smoothing is used. If less than two years of data is available, a Seasonal Regression model is used.

Even if there is too little data to create a seasonal forecast (in general less than 52 weeks, or a full sales cycle) the system still attempts to generate a seasonal forecast if at all possible. The sales of last year are time shifted by 52 weeks, and become the forecast for the corresponding periods. If the third week in 2020 is in the forecast horizon, and sales for the third week in 2019 are available, they will become the forecast. If not, the forecast will be generated by running SimpleES with the available (less than 52 weeks) history. For this particular case, the method picked is Seasonal Regression. For more information, see Seasonal Exponential Smoothing (SeasonalES).

Seasonal Regression

Seasonal Regression cannot be selected as a forecasting method, but is a candidate model that is used only when the SeasonalES 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. For more information, see Seasonal Regression.

Creating a Forecast Administration Workbook

Use the following procedure, to create a Forecast Administration Workbook.

1. Within the Master or Local Domain, select **New** from the File menu.
2. Select the Forecast tab to display a list of workbooks for statistical forecasting.
3. Select **Forecast Administration**.
4. Click **OK**.
5. 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.
6. Click **Finish** to open the workbook.

Worksheets for the Forecast Administration Workbook

The following sections describe the components of the Forecast Administration workbook.

The Forecast Administration Workbook has two tabs and includes the worksheets listed in Table 3–1.

Table 3–1 Forecast Administration Workbook Worksheets

Tab	Description	Worksheets
Basic Settings Workflow	Contains forecast administration settings.	Final Level Worksheet - Basic Settings Final and Source Level Parameters Worksheet - Basic Settings
Advanced Settings Workflow	Used to set parameters related to either the data that is stored in the system or the forecasting methods that are used at the final or source-levels. The parameters on this workflow tab are not as likely to be changed on a regular basis as the ones on the Basic Settings workflow tab.	Causal Parameters Worksheet - Final Level Causal Parameters - Source Levels Demand Transference Parameters Worksheet- Advanced Settings Final Level Parameters Worksheet - Advanced Settings Final and Source Level Worksheet - Advanced Settings Promo Affinity Parameters Worksheet Regular Price Parameters Worksheet- Advanced Settings

Final Level Worksheet - Basic Settings

The 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 is ultimately spread). Forecast approvals and data exports can only be performed on forecasts at a final-level. Figure 3–1 provides an example of a view of the Final Level Parameters worksheet in a master domain with three partitions/local domains, partitioned based on group.

Figure 3–1 Final Level Parameters Worksheet

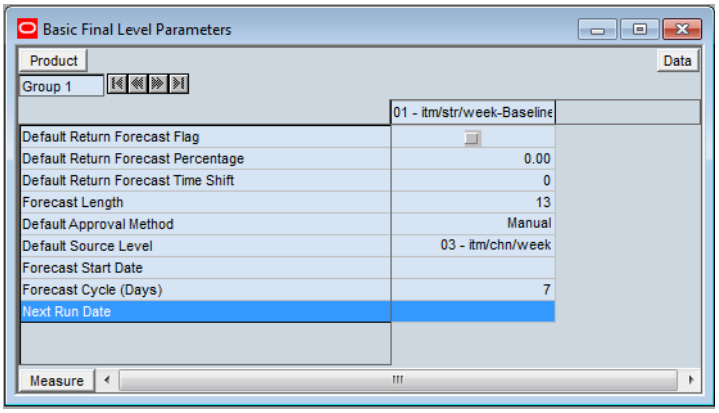

Measures: Final Level Worksheet - Basic Settings

The Final Level Worksheet - Basic Settings contains the following measures:

Default Approval Method

This field is a list from which you select the default automatic approval policy for forecast items. Valid values are:

Field	Description
Manual	The system-generated forecast is not automatically approved. Forecast values must be manually approved by accessing and amending the Forecast Approval Workbook.

Field	Description
Automatic	The system-generated quantity is automatically approved as is.
By Alert	This list of values may also include any Forecast Approval alerts that have been configured for use in the forecast approval process. Alerts are configured during the implementation and can be enabled to be used for Forecast Approval in the Enable Alert for Forecast Approval worksheet.. Refer to the *Oracle Retail Predictive Application Server Configuration Tools User Guide* for more information on the Alert Manager.The Alert Parameters workbook contains a list and descriptions of available alerts, and for which level (causal/baseline) that they are designed for.

Default Return Forecast Flag

If this Boolean measure is checked, it indicates that a returns forecast will be generated.

Default Return Forecast Percentage

This measure indicates the percentage of the merchandise sold in a time period that is expected to be returned to the store. A value of zero indicates that no returns will be calculated.

Default Return Forecast Time Shift

This measure indicates the number of time periods between when the merchandise was sold until it was returned. For instance, if the measure is set to one, the merchandise is expected to be returned one week after it was sold.

Default Source Level

The pick list of values displayed in this field allows you to change the forecast level that is used as the primary level to generate the source forecast. The source-levels are set up in the RPAS Configuration Tools. A value from the pick list is required in this field at the time of forecast generation.

Forecast Cycle

The Forecast Cycle is the amount of time (measured in days) that the system waits between each forecast generation. Once a scheduled forecast has been generated, this field is used to automatically update the Next Run Date field. A non-zero value is required in this field at the time of forecast generation.

Forecast Start Date

This is the starting date of the forecast. If no value is specified at the time of forecast generation, the system uses the data/time at which the batch is executed as the default value. If a value is specified in this field and it is used to successfully generate the batch forecast, this value is cleared.

Forecast Length

The Forecast Length is used with the Forecast Start Date to determine forecast horizon. The forecast length is based on the calendar dimension of the final-level. For example, if the forecast length is to be 10 weeks, the setting for a final-level at day is 70 (10 x 7 days).

Enable New Item Functionality

The New Item functionality is not automatically available for every forecast level. You can enable it for a certain level by checking this measure.

Causal Run Mode

Note: This measure is not displayed for baseline levels.

Allows you to select the default behavior of the causal engine. The are three options:

- Run Pooling Estimation—the promotion effects are estimated at the specified source levels. No forecast is generated.
- Run Forecast—the promotion effects are estimated at the final forecast level. Then, if a source level is specified as either default or override, the promotion effects from the final and specified source level are blended. Finally, the blended promotion effects are used to generate the forecast. As Pooling Effects only come from Source Levels, then those effects are used at the final level only if the specified source level is not the same as the final level.
- Run Both—the causal engine will performs all the actions described in the other options.

Next Run Date

The Next Run Date is the date on which the next batch forecast generation process automatically runs. RDF automatically triggers a set of batch processes to be run at a pre-determined time period. When a scheduled batch is run successfully, the Next Run Date automatically updates based on the Start Date value and the Forecast Cycle. No value is required in this field when the Forecast Batch Run Workbook wizard is used to generate the forecast or if the batch forecast is run from the back-end of the domains using the override True option. Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on forecast generation.

Final and Source Level Parameters Worksheet - Basic Settings

The Final and Source Level worksheet allows you to set the default parameters that are common to both the final and source-level forecasts.

Figure 3–2 Final and Source Level Parameters Worksheet

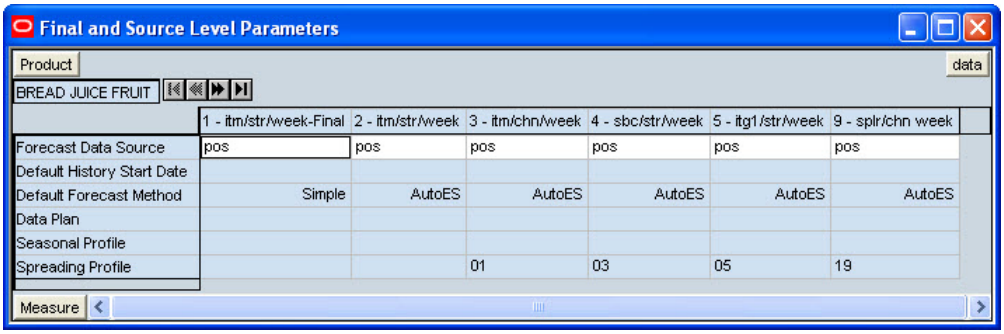

Measures: Final and Source Level Parameters Worksheet - Basic Settings

The Final and Source Level Parameters Worksheet - Basic Settings contains the following measures:

Forecast Data Source

This is a read-only value that displays the sales measure (the measure name) that is the data used for the generation of forecasts (for example, POS). The measure that is

displayed here is determined at configuration time in the RPAS Configuration Tools. Different Data Sources can be specified for the Final and Source Levels.

Default History Start Date

This field indicates to the system the point in the historical sales data at which to use in the forecast generation process. If no date is indicated, the system defaults to the first date in your calendar. 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, 2011 and an item/location does not have sales history until February 1, 2011, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value. Different History Start Dates can be specified for the Final and Source Levels.

Data Plan

Used in conjunction with the Bayesian forecast method, Data Plan is used to input the measure name of a sales plan that should be associated with the final-level forecast. Sales plans, when available, provide details of the anticipated shape and scale of an item's selling pattern. If the Data Plan is required, this field should include the measure name associated with the Data Plan.

Note: This measure is not displayed for causal levels.

Default Forecast Method

The Default Forecast Method is a complete list of available forecast methods from which you can select the primary forecast method that is used to generate the forecast. A summary of methods is provided in Forecasting Methods Available in RDF. The chapter, Oracle Retail Demand Forecasting Methods, covers each method in greater detail. It is important to note that Causal should not be selected unless the forecast level was set as a Causal level during the configuration. Refer to the *Oracle Retail Demand Forecasting Configuration Guide* for more information on configurations using the Causal forecast method.

Note: If this a Causal forecast level, then the only available methods are Causal and No Forecast. Alternately, when the level is non-causal, the Causal method is not available.

Seasonal Profile

Used in conjunction with the Profile-Based forecasting method, this is the measure name of the seasonal profile that is used to generate the forecast at either the source or final-level. Seasonal profiles, when available, provide details of the anticipated seasonality (shape) of an item's selling pattern. The seasonal profile can be generated or loaded, depending on your configuration. The original value of this measure is set during the configuration of the RDF solution.

Note: This measure is not displayed for causal levels.

Spreading Profile

Used for source-level forecasting, the value of this measure indicates the profile level that is used to determine how the source-level forecast is spread down to the final-level. No value is needed to be entered at the final-level. For dynamically

generated profiles, this value is the number associated with the final profile level (for example, 01). Note that profiles 1 through 9 have a zero (0) preceding them in Curve—this is different than the forecasting level numbers.

For profiles that must be approved, this is the measure associated with the final profile level. This measure is defined as apvp +level (for example, apvp01 for the approved profile for level 01 in Curve).

Note: This measure is not displayed for causal levels.

Source Level Forecast Data Source

You can specify a measure to serve as a forecast data source at the source-level. If the measure is left blank for a certain source-level, the source measure from the corresponding final-level is used.

Source Level History Start Date

This is the starting date for historical sales data at the source-level. For example, if your system start date is January 1, 2003, but you only want to use historical sales data from the beginning of 2011, you need to set your History Start Date to January 1, 2011. 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, RDF 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, 2011, and an item/location does not have sales history until February 1, 2011, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

Note: History start dates for final and source-levels can be different.

Causal Parameters Worksheet - Final Level

The Causal Parameters worksheet allows you to set the parameters that support promotional forecasting. The parameters that support overlapping promotions are also specified in this view. Overlapping promotions are promotions happen at the same time period at the same location for the same product. This worksheet only includes the causal forecast levels for the final-level selected during the wizard process. This worksheet is not visible if Promote/Promotional Forecasting is not implemented in your RDF environment.

Figure 3–3 Causal Parameters Worksheet - Final Level

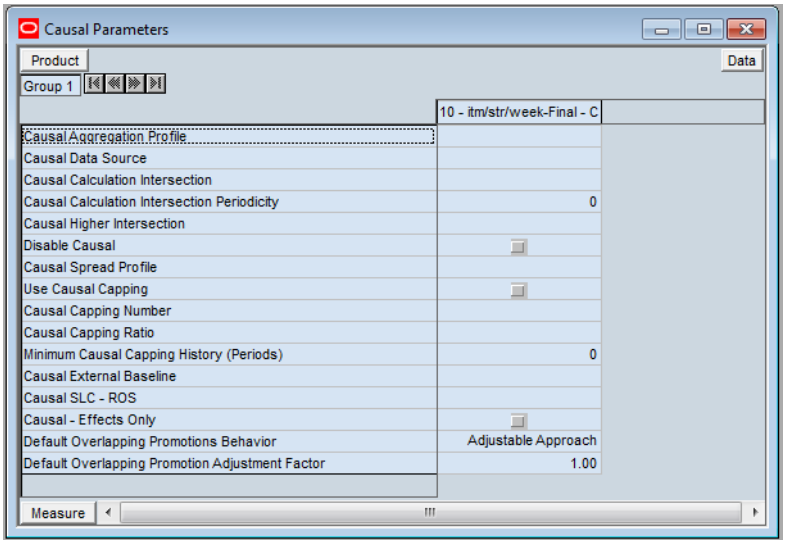

Measures: Causal Parameters Worksheet - Final Level

The Causal Parameters Worksheet - Final Level contains the following measures:

Causal Aggregation Profile

Used only for Daily Causal Forecasting, the Causal Aggregation Profile is measure name of the profile used to aggregate promotions defined at day up to the week. The value entered in this field is the measure name of profile and it must be valid. If this profile is generated within Curve, the format of the measure name is apvp+level (for example, apvp01). Note that the only aggregation of promotion variables being performed here is along the Calendar hierarchy. RDF does not support aggregation of promotion variables along other hierarchies such as product and location hierarchies.

Causal Baseline Minimum Winters History Length

Used only in Causal forecasting, this parameter is the minimum number of historical time periods that are required for Winters to be used to produce the Causal baseline. The default setting is 104 periods, but it may be set to lower number, for example, 52 periods.

Causal Calculation Intersection

Used only for Daily Causal Forecasting, the Causal Calculation Intersection is the intersection at which the causal forecast is run. The format needs to match the hierarchy dimension names set in the RPAS Configuration Tools (such as itemstr_ week). Each dimension must have only four characters. The order of the dimension does not matter. There is no validation of correct format of this intersection.

Causal Capping Number

The value of this parameter is a measure based on the same product and location intersection as the forecast level. The maximum value of the referenced measure times the Causal Capping Ratio is the maximum value to use for calculating the causal forecast for time series that meet the Causal Capping conditions.

Note: For the Causal Daily forecast, the causal capping number is used to cap the week causal forecast, then the weekly level peak is spread down to the daily level.

If Use Causal Capping is set to *True*, the history for the time series is greater than or equal to the Minimum Causal Capping History, and the preliminary forecast is greater than or equal to the value in the Causal Capping Number measure times the Causal Capping Ratio, then the forecast is recalculated to be the value in the Causal Capping Number measure multiplied by the value in the Causal Capping Ratio measure.

Causal Capping Ratio

The value of this parameter is a measure based on the same product and location intersection as the forecast level. This measure contains the ratio that is used to calculate the forecasts for time series that meet the Causal Capping conditions. If Use Causal Capping is set to *True*, the history for the time series is greater than or equal to the value in the Minimum Causal Capping History measure, and the preliminary forecast is greater than or equal to the value in the Causal Capping Number measure times the Causal Capping Ratio, then the forecast is recalculated to be the Causal Capping Number multiplied by the value in the Causal Capping Ratio measure.

Causal Data Source

Used only for Daily Causal Forecasting, the Causal Data Source is an optional setting that contains the measure name of the sales data to be used if the data for causal forecasting is different than the Data Source specified at the Final level. If needed, this field should contain the measure name of the source data measure (for example, DPOS).

Causal Baseline

Enter the measure name that you want to serve as the baseline for the causal forecast at this level. The measure can be loaded or it can be generated in a different forecast run.

Causal Higher Intersection

An optional setting for Causal Forecasting, this intersection is the aggregate level to model promotions if the causal intersection cannot produce a meaningful causal effect. This intersection applies to promotions that have a Promotion Type set to Override From Higher Level (set in the Promo Effect Maintenance Workbook). The format of this intersection needs to match the hierarchy dimension names set in the RPAS Configuration Tools-such as sclsrgn_ (Subclass/Region), and it must not contain the calendar dimension. Each dimension must have only four characters. The order of the dimension does not matter. There is no validation of correct format of this intersection.

Causal Spread Profile

Used only for Daily Causal Forecasting, the Causal Spread Profile is the measure name of the profile used to spread the causal baseline forecast from the Causal Calculation Intersection to the Final Level. If this profile is generated in Curve, this measure value is apvp+level (for example, apvp01).

Default Overlapping Promotion Adjustment Factor

The Overlapping Promotion Behavior measure is applied only if Adjustable Approach is selected for Default Overlapping Promotions Behavior.

This Default Overlapping Promotion Adjustment Factor specifies at a high level how the individually calculated promotions interact with each other when they are overlapping in the forecast horizon. This parameter serves as a global setting, but can be overridden at lower levels. The default value is 1.

- A value greater than 1 means the promotion effects will be compressed when applied in the model, instead of linearly summing up to get the total promotion effect. The larger the value is, the larger the compression effect will be, meaning the smaller the total effect will be.
- A value between 0 and 1 means the promotion effects will be amplified when applied in the model, instead of linearly summing up to get the total promotion effect. The smaller the value is, the larger the amplification effect will be, meaning the larger the total effect will be.
- If a value less than or equal to 0 is put in the cell, the calculation engine will find the best adjustment factor to fit the history data. This factor is then used to combine overlapping promotions in the forecast horizon. With few exceptions, the value should be anywhere from 1 to 5. This factor is then used to combine overlapping promotions in the forecast horizon

> **Caution:** The process to find the best adjustment factor can be time consuming and introduce performance issues.

Default Overlapping Promotions Behavior

The Overlapping Promotion Behavior measure allows you to specify if overlapping promotions function should be activated during the batch run. If so, the causal algorithm will create a new promotion for each unique combination of promotions in the history. The new promotions are used in the stepwise regression, and they are used in calculating the forecast if they are found significant. The new promotion combinations are temporary. This parameter serves as a global setting, but can be overridden at lower levels.

The Default Overlapping Promotions Behavior allows you to specify at a high level if overlapping promotion calculate engine should be called during the batch run. The option is the Adjustable Approach:

- **Adjustable Approach**

  This approach means that overlapping promotion function will not be activated. Each promotion's effect is calculated individually in the stepwise regression, and RDF will apply the Adjustable Factor (link) when overlapping happens in the forecast horizon.

  This approach should be used when the promotion history is not well established and less reliable.

It should be noted that just because the overlapping promotion function is activated, it does not necessarily imply an overlapping promotional forecast result. In some instances, it is possible that no overlapping promotions are found to be statistically significant. The indicator of the significant overlapping promotions' existence can be viewed within the Forecast Approval Workbook.

> **Note:** It is no longer true that only Boolean type promotions and Real type promotions with the value of one (1) can be considered the candidates of overlapping promotions. Every promotion can be used for overlapping.

Minimum Causal Capping History (Periods)

If Use Causal Capping is set to *True*, this parameter is used to set minimum number of historical time periods required before the system considers a time series for causal capping. If left empty, the algorithm currently uses 52 weeks of history as the default for this parameter.

Default Blending Parameter

This parameter sets the weights for combining the Final and Source Level promotion effects, when calculating the blended effect. The range of the parameter is 0 to 1. A value closer to 1 will yield a blended effect closer to the Source Final Level effect. A value closer to zero yields an effect closer to the Final Level effect. The value can be overridden in the Forecast Maintenance workbook.

Note: More detailed information on the Causal forecasting algorithm is provided in Forecasting Methods Available in RDF.

Use Causal Capping

Place a check in this parameter (set to True) if capping is to be applied to the causal forecast. Also required for causal capping are the following parameters:

- Causal Capping Number
- Causal Capping Ratio
- Minimum Causal Capping History (Periods)

Note: More detailed information on the Causal forecasting algorithm is provided in Forecasting Methods Available in RDF.

Causal Parameters - Source Levels

This worksheet allows you to select actions specific for causal source levels. At the source level, promotion and demand data is not simply aggregated. The final level data for all time series in a source level dimension is analyzed at the same time, resulting in potentially huge amounts of data.

Measures: Causal Parameters - Source Levels

The Causal Parameters - Source Levels contains the following measures:

Number of Extra Data Points

This measure determines the number of periods before and after that the causal engine analysis when determining the causal effects. For instance if the number is 1, and we have one promoted period in history, then the input to the causal engine is the demand for the period prior to the promoted period, the promo period demand and the demand for the period past the promoted period, for a total of 3 data points. If the value of the measure is 2, then the total number of data points is 5. This measure is very important because it influences the effect estimation, both in accuracy and processing time. If the number is too high, too many data points are processed by the engine, resulting in poor performance.

Run Estimate Flag

This measure indicates if the effect estimation should be performed for a certain source level, the next time the estimation is run.

Last Estimate Date

This measure displays the date when the effect estimation was run for a certain source level. This is useful when deciding if a rerun is necessary. For instance if there are three source levels, and two have a Estimate Date of last week, and the third a date from last month. The user may want to enable estimation only for the source level with the earlier date. There was not much activity such that a re-estimation is necessary for the other two source levels.

Demand Transference Parameters Worksheet- Advanced Settings

The Demand Transference Parameters Worksheet- Advanced Settings allows you to set parameters to decide if and how demand transference is incorporated in the forecast.

Measures: Demand Transference Parameters Worksheet- Advanced Settings

The Demand Transference Parameters Worksheet- Advanced Settings contains the following measures:

Default Apply Demand Transference Effects

A check in this field (set to True) indicates that RDF will incorporate demand transference in the forecast. Otherwise, by clearing the box (set to False), no demand transference is incorporated in the forecast generation process. The Demand transference function will be enabled if Demand Transference Enabled is specified during implementation, and Default Apply Demand Transference is set to True. This measure serves as a global setting, but can be overridden at lower levels.

RDF loads assortment multipliers from ORME (Oracle Retail Modeling Engine). The demand transference effects are derived from these assortment multipliers, at a prod/loc/clnd intersection. Before RDF incorporates the effects in the forecast, the value of the assortment multipliers can be checked and capped.

- If the loaded value is smaller than the Maximum Assortment Multiplier, but greater than the Minimum Assortment Multiplier, the loaded value is applied for the demand transference effect calculation.

Default Demand Transference Effect Decay Factor

This measure is applied if and only if the demand transference function is enabled and applied. This parameter serves as a global setting, but can be overridden at lower levels.

This parameter determines how the magnitude of the demand transference effects is applied over time. An assortment change usually triggers a transference in demand. However, the effect can be seen only for a few periods following the assortment change. In addition, an adjustable damping mechanism ensures that the demand transference effect trends down as time goes by. The decaying factor is a number between zero (0) and one (1).

RDF loads assortment multipliers from ORME (Oracle Retail Modeling Engine). The demand transference effects are derived from these assortment multipliers, at a prod/loc/clnd intersection. Before RDF incorporates the effects in the forecast, the value of the assortment multipliers can be checked and capped.

The closer the decaying factor is to zero (0), the faster the demand transference effect decays.

Default Demand Transference Effects Applying Period

This measure is applied only if the demand transference function is enabled and applied. This parameter serves as a global setting, but can be overridden at lower levels.

This integer measure specifies the number of periods for which the demand transference effect is applied. An assortment change usually triggers a transference in demand. However, the effect is usually seen only for a few periods following the assortment change. The number of periods is specified as an integer in this cell.

Figure 3–4 Demand Transference Decaying Effect

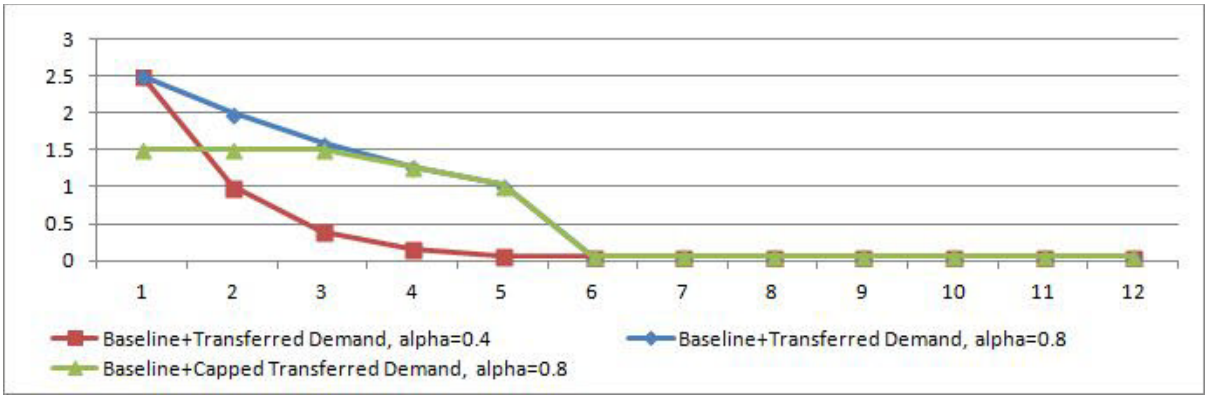

Default Maximum Assortment Multipliers

This measure is applied if and only if demand transference function is enabled and applied. This parameter serves as a global setting, but can be overridden at lower levels. It is used to set the maximum threshold of the assortment multipliers.

- If the loaded value is greater than the Maximum Assortment Multiplier, the Maximum Assortment Multiplier is applied for the demand transference effect calculation.
- If the loaded value is smaller than the Maximum Assortment Multiplier but greater than the Minimum Assortment Multiplier (link), the loaded value is applied for the demand transference effect calculation.

Default Minimum Assortment Multipliers

This measure is applied if and only if demand transference function is enabled and applied. This parameter serves as a global setting, but can be overridden at lower levels. It is used to set the minimum threshold of the assortment multipliers.

- If the loaded value is smaller than the Minimum Assortment Multiplier, the Minimum Assortment Multiplier is applied for the demand transference effect calculation.
- If the loaded value is greater than the Minimum Assortment Multiplier but smaller than the Maximum Assortment Multiplier (link), the loaded value is applied for the demand transference effect calculation.

Demand Transference Enabled

This read-only measure specifies if demand transference is configured for the final level. The decision to enable demand transference is made in the Configuration Tools.

Note: The content of the measure is visible if you roll up to All Products on the product hierarchy. At a lower intersection the cell is unavailable.

Final Level Parameters Worksheet - Advanced Settings

The Final Level worksheet allows you to set the advanced parameters for the final-level forecasts. Figure 3–5 provides an example of a view of this worksheet in a master domain with three partitions/Local Domains, partitioned based on group.

Figure 3–5 Final Level Parameters Worksheet

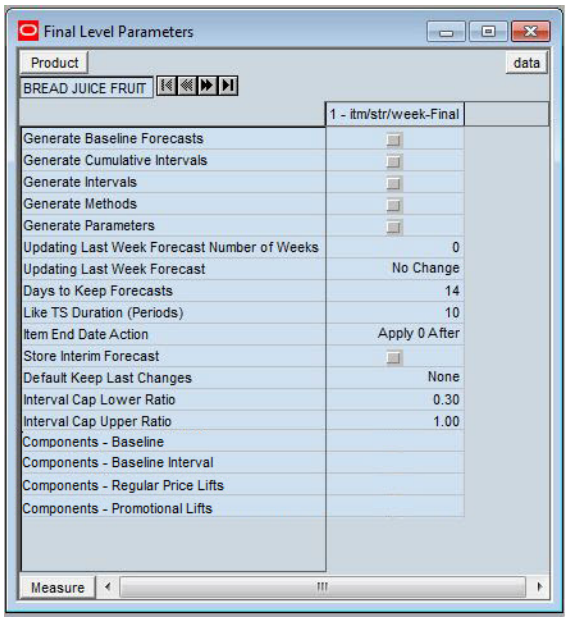

Measures: Final Level Parameters Worksheet - Advanced Settings

The Final Level Parameters worksheet - Advanced Settings contains the following measures:

Components - Baseline

This measure stores the name of the measure that serves as baseline in the Components forecasting measure. If no measure is specified, the Components forecasting method assesses the value as one (1).

Note: This measure is not displayed for causal levels.

Components - Baseline Interval

This measure stores the name of the measure that serves as baseline interval in the Components forecasting measure.

Note: This measure is not displayed for causal levels.

Components - Regular Price Lifts

This measure stores the name of the measure that serves as regular price lifts in the Components forecasting measure. If no measure is specified, the Components forecasting method assesses the value as one (1).

Note: This measure is not displayed for causal levels.

Components - Promotional Lifts

This measure stores the name of the measure that serves as promotional lifts in the Components forecasting measure. If no measure is specified, the Components forecasting method assesses the value as one (1).

Note: This measure is not displayed for causal levels.

Days to Keep Forecasts

This field is used to set the number of days that the system stores forecasts based on the date/time the forecast is generated. The date/time of forecast generation is also referred to as birth date of the forecast. A forecast is deleted from the system if the birth date plus the number of days since the birth date is greater than the value set in the Days to Keep Forecast parameter. This process occurs when either the Forecast Batch Run Workbook wizard is used to generate the forecast or when PreGenerateForecast is executed.

When you start the Forecast Batch Run Workbook, one of the first steps that RDF takes is to remove forecasts older than the value given in the Days to Keep Forecasts parameter. RDF uses the value of that parameter in your local domain. However, because forecasts are registered globally, the Forecast Delete has a global effect. That is, if you set the parameter to some low value, for example, one day all forecasts in all local domains older than one day are removed. This can adversely impact other users of the system if done without consideration.

When the PreGenerate utility is used, typically for automated batch production of forecasts, RDF uses the value for the Days to Keep Forecasts parameter in the first local domain. Again, if this is set to a value that is not expected by the other users, it can cause unintended disruption.

Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on PreGenerateForecast.

Default Apply Interaction Factor

When generating the forecast, there are three possible lifts added on top of the baseline to get the final result. They are promotion lift, regular price change lift, and demand transference lift. The Default Apply Interaction Factor measure allows you to specify at a high level if these individually calculated lifts should interact with each other during the forecast generation. This parameter serves as a global setting, but can be overridden at lower levels.

A check in this field (set to True) indicates if you want RDF to combine the promotion, regular price, and demand transference lifts. The impact of the combination can be adjusted in the Default Interaction Factor measure.

Otherwise, by clearing the box (set to False) these effects will be linearly added to the baseline to generate the forecast. By default, the value is set to False. The calculated lifts can be viewed in the Forecast Approval Workbook.

Default Interaction Factor

The Interaction Factor is applied only if the Default Apply Interaction Factor measure is set to True.

This Default Interaction Factor specifies at a high level how the individually calculated promotion lift, regular price change lift and demand transference lift interact with each other during the forecast generation process. This parameter serves as a global setting, but can be overridden at lower levels. The default value is 2.

- A value equal to 1 means that each lift will be summed up to calculate the total lift.
- A value greater than 1 means the lifts will be compressed when applied in the model, instead of linearly summing up to get the total lift. The larger the value is, the larger the compression effect will be, meaning the smaller the total lift will be.
- A value between 0 and 1 means the lifts will be amplified when applied in the model, instead of linearly summing up to get the total lift. The smaller the value is, the larger the amplification effect will be, meaning the larger the total lift will be.

Note that if Apply Interaction Factor is set to True and the Interaction Factor is set to a value other than 1, user should expect to see the combined effects for Demand Transference, Promotion, and Regular Price in Forecast Approval Workbook.

Default Keep Last Changes

This field is a list from which you select the default change policy for forecast items. Valid values are:

Field	Description
Keep Last Changes (None)	There are no changes that are introduced into the adjusted forecast. The adjusted forecast is equal to the system forecast.
Keep Last Changes (Total)	Considers both the Last System Forecast and the Last Approved Forecast in determining change policy, but only uses the Last Approved Forecast to calculate the adjustment value. For each item/week combination that is forecast, RDF automatically introduces the same quantity that was approved in the last approved forecast into the change only if that quantity differed from that in the last system forecast. If the quantities are the same, RDF introduces the current system-generated forecast into the adjusted forecast.
Keep Last Changes (Diff)	This method allows you to keep the difference between previous adjusted forecast and system forecast if there is any adjustment existing. For each product/location/week in the forecast horizon, when last approved system forecast is different from last approved forecast, it means that you have made an adjustment previously. The last approved forecast minus the last approved system forecast plus the current system forecast is used to populated the current adjusted forecast. If the last approved system forecast is the same as last approved forecast, it means that you have not made any adjustment previously and the current adjusted forecast will be populated with current system forecast.

Field	Description
Keep Last Changes (Ratio)	This method allows you to keep the ratio of last adjusted forecast versus the last system forecast if there is any adjustment existing. For each product/location/week in the forecast horizon, when the last approved system forecast is different from last approved forecast, it means that you have made an adjustment previously. The ratio of last approved forecast to last approved system forecast multiplied by the current system forecast is used to populated the current adjusted forecast. If the last approved system forecast is the same as last approved forecast, it means that you have not made any adjustment previously and the current adjusted forecast will be populated with current system forecast.

Generate Cumulative Interval

A check in this field (set to True) specifies whether you want RDF to generate cumulative intervals (this is similar to cumulative standard deviations) during the forecast generation process. Cumulative Intervals are a running total of Intervals and are typically required when RDF is integrated with the Oracle Retail Merchandising System. 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 the Forecast Approval Workbook.

Generate Intervals

A check in this field (set to True) indicates that intervals (similar to Standard Deviations) should be stored as part of the batch forecast process. Intervals can be displayed in the Forecast Approval Workbook. If you do not need intervals, excess processing time and disk space may be eliminated by clearing the check box. For many forecasting methods, intervals are calculated as standard deviation, but for Simple, Holt, and Winters the calculation is more complex. Intervals are not exported.

Generate Methods

A check in this field (set to True) indicates that when an ES forecast method is used, the chosen forecast method for each fitted time series should be stored. The chosen method can be displayed in the Forecast Approval Workbook.

Generate Parameters

A check in this field (set to True) indicates that the alpha, level, and trend parameters for each fitted time series should be stored. These parameters can be displayed in the Forecast Approval Workbook.

Interval Cap Lower Ratio

The value entered in this field multiplied by the forecast represents the lower bound of the confidence interval for a given period in the forecast horizon.

Interval Cap Upper Ratio

The value entered in this filed multiplied by the forecast represents the upper bound of the confidence interval for a given period in the forecast horizon.

Item End Date Action

This parameter allows the option for items with end dates within the horizon to have zero demand applied to time series before or after the interim forecast is calculated. The two options are:

Option	Description
Apply 0 After Spreading	This is the default value. Spreading ratios are calculated for time series with no consideration made to the end date of an item. It is after the source forecast is spread to the final-level when zero (0) is applied to the System Forecast.
Apply 0 to Interim	For items that have an end date within the forecast horizon; zero (0) is applied to the Interim Forecast before the spreading ratios are calculated. This ensures that no units are allocated to the final-level for time series that have ended.

Store Interim Forecast

A check should be placed in this field (set to True) if the interim forecast is stored. The Interim Forecast is the forecast generated at the Final Level. This forecast is used as the Source Data within Curve to generate the profile (spreading ratios) for spreading the source-level forecast to the final-level. The interim forecast should only be stored if it is necessary for any analysis purposes.

Note: This measure is not displayed for causal levels.

Updating Last Week Forecast

This field is a list from which you can select the method for updating the Approved Forecast for the last specified number of weeks of the forecast horizon. This option is valid only if the Approval Method Override (set in the Forecast Administration Workbook) is set to Manual or Approve by alert, and the alert was rejected.

Method	Description
No Change	When using this method, the last week in the forecast horizon does not have an Approved Forecast value.
Replicate	When using this method the last week in the forecast horizon is forecast using the Approved Forecast for the week prior to this time period. To determine the appropriate forecast time period, the value is subtracted from the Forecast Length.
Use Forecast	When using this method, the System Forecast for the last weeks in the forecast horizon is approved.

Final and Source Level Worksheet - Advanced Settings

The Final and Source Level worksheet allows you to set the advanced parameters that are common to both the final and source-level forecasts.

Note: This worksheet is not available for causal levels.

Figure 3–6 Final and Source Level Parameters Worksheet

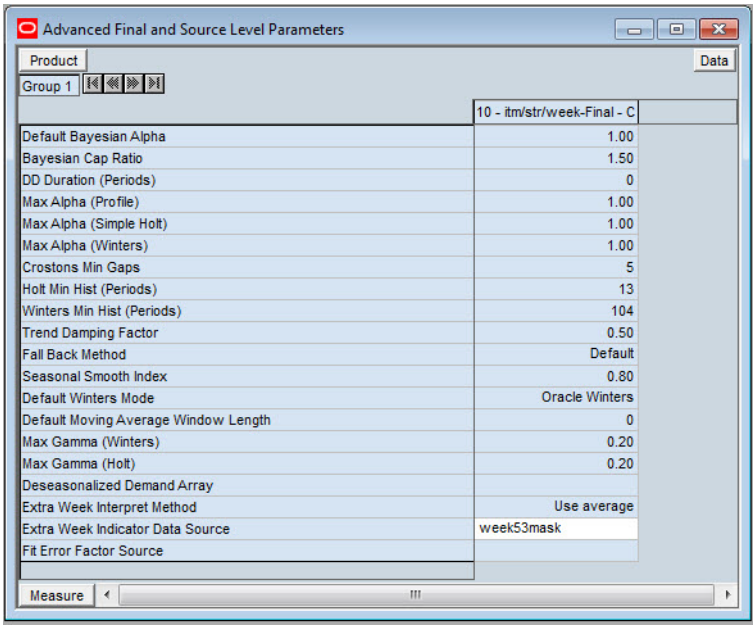

Measures: Source Level Worksheet - Advanced Settings

The Source Level Worksheet - Advanced Settings contains the following measures:

Bayesian Alpha (range (0, infinity))

When using the Bayesian 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 one (or infinity) weights the sales plan more in creating the forecast, whereas alpha closer to zero (0) weights the known history more. The default is one (1).

Bayesian Cap Ratio

The Bayesian Cap ratio is used to cap the resulting Bayesian forecast if it deviates significantly from the sale plan. The Bayesian Cap ratio is used as follows:

Example 3–1

```
If forecast I > Bayesian Cap ratio * Max value of Past Sales AND
forecast I >  Bayesian Cap ratio * Max value of Past Plan AND
forecast I >  Bayesian Cap ratio * Plan I
then
forecast I  = Plan I
```

It defaults to a value of 1.5.

Crostons Min Gaps

The Crostons Min Gaps is the default minimum number of transitions from non-zero sales to zero sales. Thus, if Croston’s Min Gap is set to five, then the method may fit if you have five or more transitions from non-zero sales to zero sales. If there are not enough gaps between sales in a given product's sales history, the Croston's model is

not considered a valid candidate model. The system default is five minimum gaps between intermittent sales. The value must be set based on the calendar dimension of the level.

For example, if the value is to be 5 weeks, the setting for a final-level at day is 35 (5x7days) and a source-level at week is 5.

DD Duration (weeks)

Used with Profile Based forecast method, the DD Duration is starting from the first week of populated history, the number of weeks of seasonal profiles required after which the system stops using the DD (Deseasonalized Demand) approach and defaults to the normal Profile-Based method. The value must be set based on the calendar dimension of the level. For example, if the value is to be 10 weeks, the setting for a final-level at day is 70 (10x7days) and a source-level at week is 10.

Default Moving Average Window Length

Used with Moving Average forecast method, this is the Default number of data points in history used in the calculation of Moving Average. This parameter can be overwritten at item/location from the Forecast Maintenance Workbook.

Fallback Method

Set this parameter only if the Fallback Method is to vary from the default Fallback Methods used by the selected forecasting algorithm. If the method selected as the Default Forecast Method or Forecast Method Override does not succeed for a time series, this method is used to calculate the forecast and the default Fallback Methods in the forecasting process is skipped entirely.

The default Fallback Methods are listed in Table 3–2.

Table 3–2 Fallback Methods

Fallback Method	Steps
If either the Causal, Bayesian, or Profile-Based are selected as the Default Forecast Method or Forecast Method Override and the method does not fit the data.	1. RDF attempts to fit SeasonalES 2. RDF attempts to fit TrendES 3. RDF attempts to fit Simple/IntermittentES
If the SeasonalES is selected as the Default Forecast Method or Forecast Method Override and neither Multiplicative Seasonal or Additive Seasonal fits the data.	1. RDF attempts to fit TrendES 2. RDF attempts to fit Simple/IntermittentES
If either the Multiplicative Seasonal or Additive Seasonal are selected as the Default Forecast Method or Forecast Method Override and the method does not fit the data.	1. RDF attempts to fit TrendES 2. RDF attempts to fit Simple/IntermittentES
If the TrendES is selected as the Default Forecast Method or Forecast Method Override and the method does not fit the data.	1. RDF attempts to fit Simple/IntermittentES

Holt Min Hist (Periods)

Used with the AutoES forecast method, Holt Min Hist is the minimum number of periods of historical data necessary for the system to consider Holt (TrendES) as a potential forecasting method. RDF 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, Holt is not be considered as a valid option. The system default is 13 periods. The value must be set based on the calendar dimension of the level. For example, if the value is to be 13 weeks, the setting for a final-level at day is 91 (13x7days) and a source-level at week is 13.

Max Alpha (Profile) (range [0.001 to 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. This field displays the maximum value (that is, cap value) of alpha allowed in the model-fitting process. An alpha cap value closer to one (1) allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to zero (0) only allows less reactive models. The default is one (1).

The value for the optimized alpha parameter will be in the range 0.001 to Max Alpha.

Max Alpha (Simple, Holt) (range [0.001 to 1])

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

The value for the optimized alpha parameter will be in the range 0.001 to Max Alpha.

Max Alpha (Winters) (range [0.001 to 1])

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

The value for the optimized alpha parameter will be in the range 0.001 to Max Alpha.

Max Gamma (Holt) (range [0,1])

In the Holt (TrendES) model-fitting procedure, gamma (a model parameter capturing the trend) is determined by optimizing the fit over the time series. This field displays the maximum value (cap value) of gamma allowed in the model-fitting process.

The value for the optimized gamma parameter will be in the range 0.001 to Max Gamma.

Max Gamma (Winters) (range [0,1])

In the Winters (SeasonalES) model-fitting procedure, gamma (a model parameter capturing the trend) is determined by optimizing the fit over the time series. This field displays the maximum value (cap value) of gamma allowed in the model-fitting process.

The value for the optimized gamma parameter will be in the range 0.001 to Max Gamma.

Seasonal Smooth Index

This parameter is used in the calculation of seasonal index. The current default value used within forecasting is 0.80. Changes to this parameter impacts the value of seasonal index directly and impact the level indirectly. When seasonal smooth index is set to one (1), seasonal index is closer to the seasonal index of last year sales. When seasonal smooth index is set to zero (0), seasonal index is set to the initial seasonal indexes calculated from history. This parameter is used for Oracle Winters.

Trend Damping Factor (range [0,1])

This parameter determines how reactive the forecast is to trending data. A value close to zero (0) is a high damping, while a value if one (1) implies no damping. The default is 0.5.

Winters Min Hist (Periods)

Used with the AutoES forecast method, the value in this field is the minimum number of periods of historical data necessary for Winters to be considered as a potential forecast method. If not enough years of data are available for a given time series, Winters is not used. The system default is two years of required history. The value must be set based on the calendar dimension of the level. For example, if the value is to be 104 weeks/2 years, the setting for a final-level at day is 728 (104 weeks x 7 days) and a source-level at week is 104.

It is recommended that you choose a Forecasting Approach that best suits the nature of your business. The default forecasting approach is Oracle Winters. For additional information on the forecasting approaches, refer to the Forecasting Methods Available in RDF section.

Note: If patching this change into a domain, in order to view this measure in the Forecast Administration Workbook, you must add it to the Final and Source Level Parameters Worksheet - Basic Settings by selecting it from the Show/Hide dialog within the RPAS Client.

Deseasonalized Demand Array

Enter the name of the measure that serves as the deseasonalized data source in the profile-based forecasting method. If left blank, the system uses the forecast data source in combination with a seasonal profile to generate the deseasonalized demand.

Enable AutoSource

This measure allows you to select the forecast levels that are competing when running the AutoSource functionality. For example, if you do not want the final forecast level included, then leave all source levels enabled, and make sure the final level is not selected.

Extra Week Indicator Data Source

This measure stores the name of the measure that indicates which week is the 53rd, or extra week. To set this measure, use the Extra Week Indicator found in the Extra Week Administration Workbook.

Extra Week Interpret Method

This measure indicates how the forecasting value should be calculated if for a week that was flagged extra, or 53rd. This is necessary when an extra week is added within the forecast horizon, because initially no value is calculated for that week.

The options to calculate the value are:

- Use average: the value for the extra week is calculated as the average of the values of the weeks immediately prior and after the flagged week.
- Use before: the value for the extra week is a replication of the value in the previous week
- Use after: the value for the extra week is equal to the value of the following week

Fit Error Factor Source

The measure stores the measure name that error factor of sales vs forecast. This measure needs to be at the intersection of the final forecast level, less the calendar dimension, example, item/store.

The Error Factor is necessary for the calculation of the Confidence Intervals.

For the majority of the forecasting methods, the error factor is calculated during forecast generation, but for the following methods the calculation is not possible and needs to be input:

- Copy
- LoadPlan
- Component

Day of Week Profile Worksheet

This measure displays the default values of the day of week profile that can be used to aggregate and/or spread promotions between the two dimensions, day and week.

During the causal run it is often necessary to aggregate promotions from day to week. For example, the effect estimation is happening at the week level, but promotion indicators are stored at the day level, because events may run Thursday to Wednesday as opposed to Sunday to Saturday. If no profile is specified, a ratio of 1/7 is used. A more granular profile can be specified at a product/location intersection.

If a granular intersection is not necessary, but the 1/7 value is too rough, then the user can set a default at the local domain (product group) level. This default is stored in this measure.

Enable Alert for Forecast Approval

In this worksheet, you can select the batch alerts that can be used for forecast approval.

However, not all alerts are relevant when used to approve a certain type of forecast. For instance, for the baseline, the following are common choices:

- Forecast vs last approved forecast – also available as a real time alert
- Forecast vs sales last year – also available as a real time alert
- Forecast vs recent sales

Upon making the selection of the alerts to be used for approving the baseline forecast, click **Calculate** to update the approval list to reflect your choices:

For the causal level, relevant approval alerts are:

- Causal peaks – also available as a real time alert
- Forecast vs last approved forecast – also available as a real time alert

Upon making the selection of the alerts to be used for approving the causal forecast, click **Calculate** to update the approval list to reflect your choices:

Promo Affinity Parameters Worksheet

The Promo Affinity Parameters Worksheet allows you to set parameters to decide if cross promotional elasticities are incorporated in the forecast.

Figure 3–7 Promo Affinity Parameters Worksheet

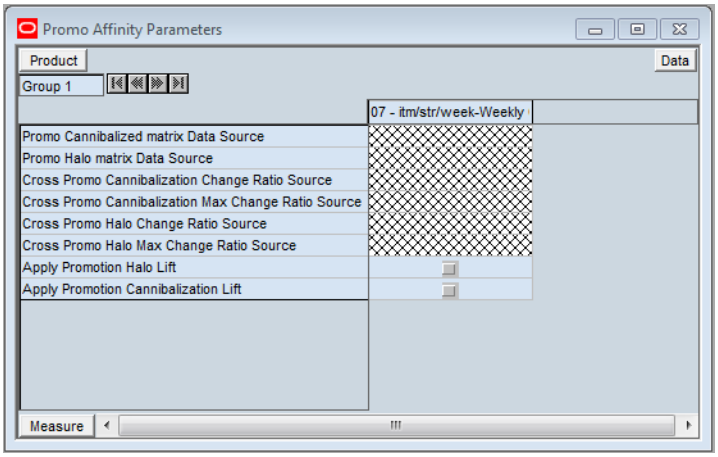

Measures: Promo Affinity Parameters Worksheet

The Promo Affinity Parameters contains the following measures:

Apply Promotion Cannibalization Lift

Select this parameter (set to True) if the application of Cannibalization lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower-levels.

Apply Promotion Halo Lift

Select this parameter (set to True) if the application of Halo lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower-levels.

Cross Promo Cannibalization Change Ratio Source

This parameter contains the name of the measure that determines the percentage of the promotional lift that is going to cannibalize related items.

> **Note:** This measure is specified at an intersection higher than item/store. The content of the measure is visible if you roll up to All Products on the product hierarchy. At a lower intersection the cell appears as a hash mark.

Cross Promo Cannibalization Max Change Ratio Source

This parameter contains the name of the measure that determines an item's maximum allowed drop in sales due to cannibalization. For instance if the sales of an item for a given period are 20 units, and the maximum allowed percentage is 20%, the drop in sales due to cannibalization for the period can not exceed 4 units.

> **Note:** This measure is specified at an intersection higher than item/store. The content of the measure is visible if you roll up to All Products on the product hierarchy. At a lower intersection the cell appears as a hash mark.

Cross Promo Halo Change Ratio Source

This parameter contains the name of the measure that determines the percentage of the promotional lift that is going to increase demand of complimentary items due to halo effect.

> **Note:** This measure is specified at an intersection higher than item/store. The content of the measure is visible if you roll up to All Products on the product hierarchy. At a lower intersection the cell appears as a hash mark.

Cross Promo Halo Max Change Ratio Source

This parameter contains the name of the measure that determines an item's maximum allowed increase in sales due to halo. For instance if the sales of an item for a given period are 15 units, and the maximum allowed percentage is 20%, the increase in sales due to halo for the period can not exceed 3 units.

> **Note:** This measure is specified at an intersection higher than item/store. The content of the measure is visible if you roll up to All Products on the product hierarchy. At a lower intersection the cell appears as a hash mark.

Promo Cannibalized Matrix Data Source

This parameter contains the name of the measure that is used to spread the cannibalization lift caused by a promoted item to the related items.

> **Note:** This measure is specified at an intersection higher than item/store. The content of the measure is visible if you roll up to All Products on the product hierarchy. At a lower intersection the cell appears as a hash mark.

Promo Halo Matrix Data Source

This parameter contains the name of the measure that is used to estimate the Halo effects between items. The measure name is implementation-specific.

> **Note:** The content of the measure is visible if you roll up to All Products on the product hierarchy. At a lower intersection the cell is unavailable.

Regular Price Parameters Worksheet- Advanced Settings

The Regular Price Parameters worksheet allows you to set the parameters that allow you to incorporate the effects of regular price changes within the forecast.

Figure 3–8 Regular Price Parameters Worksheet

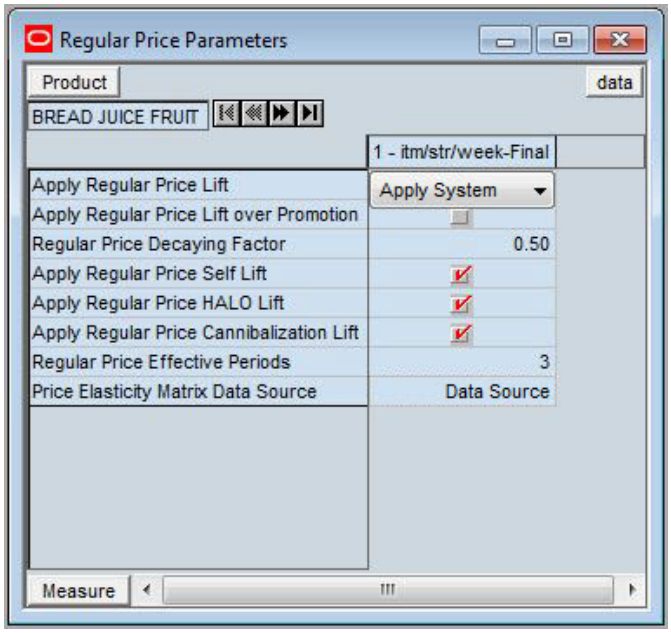

Measures: Regular Price Parameters Worksheet- Advanced Settings

The Regular Price Parameters Worksheet- Advanced Settings contains the following measures:

Apply Regular Price Lift

This parameter serves as a global setting, but can be overridden at lower-levels.

Select a value from the list for this parameter:

- **Do Not Apply**— no regular price lifts are incorporated in the forecast
- **Apply System** — the effects calculated in RDF are incorporated in the forecast
- **Apply Override** — the lifts calculated in RDF are ignored and the override lifts are incorporated

There are three components to the Regular Price Lift:

- Self Lift — determined by the Regular Price Self Elasticity
- Cannibalization — determined by the Regular Price Cannibalization Cross Elasticity
- Halo — determined by the Regular Price Halo Cross Elasticity

Apply Regular Price Lift over Promotion

Is only effective if Apply Regular Price Lift is set to Apply System or Apply Override. Place a check in this parameter (set to True) if the application of regular price lifts is enabled for time periods that are promoted. If the parameter is not checked, regular price lifts will not be applied for periods with promotions. This parameter serves as a global setting, but can be overridden at lower-levels.

Regular Price Decaying Factor

This parameter determines how the magnitude of the price effects is applied over time.

A regular price change usually triggers a change in demand. However, the effect can be seen only for a few periods following the price change. An adjustable damping mechanism ensures that the regular price effect trends down as time goes by. The decaying factor is a number between zero (0) and one (1).

The closer the decaying factor is to zero (0), the faster the regular price effect decays.

Figure 3–9 Decaying Factor Trend

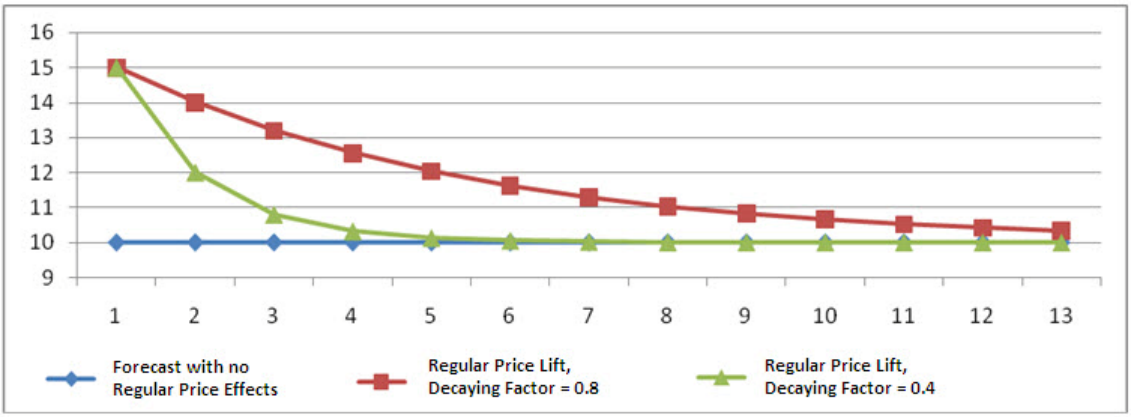

Apply Regular Price Self Lift

Is only effective if Apply Regular Price Lift is set to Apply System or Apply Override. Place a check in this parameter (set to True) if the application of Regular Price Self Lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower-levels.

Apply Regular Price Halo Lift

Is only effective if Apply Regular Price Lift is set to Apply System or Apply Override. Place a check in this parameter (set to True) if the application of Regular Price Halo Lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower-levels.

Apply Regular Price Cannibalization Lift

Is only effective if Apply Regular Price Lift is set to Apply System or Apply Override. Place a check in this parameter (set to True) if the application of Regular Price Cannibalization Lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower-levels.

Regular Price Effective Periods

Specifies the number of periods for which the regular price effect is applied.

Price Elasticity Matrix Data Source

Specifies the name of the measures that stores the regular price elasticities, such as gammas (rdfgamma). The measure name is implementation specific.

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 SeasonalES 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, and make the appropriate change to the forecast method.

The Forecast Maintenance task is split into these workflows:

- Forecast Maintenance Workbook: Basic Final Level Parameters
- Forecast Maintenance Workbook: Advanced Final Level Parameters
- Forecast Maintenance Workbook: Basic Source Level Parameters Worksheet

Creating a Forecast Maintenance Workbook

Use the following procedure to create a Forecast Maintenance workbook:

1. Within the Master or Local Domain, select **New** from the **File** menu.
2. Select the **Forecast** tab to display a list of workbooks for statistical forecasting. Select Forecast Maintenance.
3. Click **OK**. The Forecast Maintenance wizard opens and prompts you to select the level of the final-forecast. Depending on your organization's setup, you may be offered a choice of several final-forecast levels.
4. Select the final-forecast level to be viewed in the workbook.
5. Click **Next**.
6. Select the locations to include in the workbook.
7. Click **Next**.
8. Select the products to include in the workbook.
9. Click **Next**.
10. Select any additional measures (that is, measures not standard in the Forecast Maintenance workbook) that you would like included. The measure options available in this window are set in the RPAS Security Administration workbook/Workbook Template Measure Rights worksheet
11. Click **Finish** to display the workbook.

Forecast Maintenance Workbook: Basic Final Level Parameters

This worksheet includes the parameters that are overridden most frequently.

The Final Level Worksheet and Source Level worksheet allows for certain parameters set at a global level (in Forecast Administration) to vary at different item/locations..

Basic Final Level Parameters Worksheet Measures

The Final Level Worksheets for baseline and causal contain the following measures:

Approval Method Override

Set only at the final-level, the Approval Method Override is a list from which you select the approval policy for individual product/location combinations. No value is in this field if the system default set in the Forecast Administration Workbook is to be

used. Valid values are:

Field	Description
Manual	The System Forecast and Adjusted Forecast are not automatically approved. Forecast values must be manually approved.
Automatic	The Adjusted Forecast is automatically approved as is.
By Alert <name of the alert>	This list of values may also include any Forecast Approval alerts that have been configured for use in the forecast approval process. Alerts are configured during the implementation. Refer to the *Oracle Retail Predictive Application Server Configuration Tools User Guid*e for more information on the Alert Manager and the *Oracle Retail Demand Forecasting Implementation Guide* for more information on configuring Forecast Approval Alerts.

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

Forecast Method Override

Set at both final and source-levels, the Forecast Method Override is a list from which you can select a different forecast method than the Default Forecast Method set in the Forecast Administration Workbook. *No Override* appears in this field if the system default set in the Forecast Administration Workbook is to be used. Valid options depend on your system setup.

Keep Last Changes Override

Set only at final-levels, Keep Last Changes Override field may be used to override the default setting at a product/location intersection. *None* appears in this field if there is no override applied to the intersection.

Source Level Override

Set only at final-levels, the Source Level Override is the level at which the aggregate, more robust baseline 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. If it is a causal level, the source level refers to the pooling level where aggregate promotion effects are estimated.*No Override* appears in this field if the system default set in the Forecast Administration Workbook is to be used.

Bayesian Alpha Override

Note: This measure is not displayed for causal levels.

This is the override of the Bayesian Alpha parameter that is specified globally in the Forecast Administration Workbook. The override can be specified at every forecasting level, that is, final and source-levels.

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 one (1).

Moving Average Window Length Override

Note: This measure is not displayed for causal levels.

Used with Moving Average forecast method, this is the number of data points in history used in the calculation of Moving Average. This parameter, when set to a non-zero value, overrides the value in the Forecast Administration Workbook.

Optimal Source Levels

Displayed only at final-levels, a value is populated in this field if AutoSource has been run on the final-level. The AutoSource executable evaluates all levels associated to a final-level and returns the source-level that yields the optimal forecast results. For more information on AutoSource, refer to the *Oracle Retail Demand Forecasting Implementation Guide*.

Note: For item/stores that are new or highly seasonal, AutoSource may not return the best recommendation since new items may not have an adequate sales history length and highly seasonal items may only sell for a short period during the year. For these items, you should not set the AutoSource recommendation as default at the final-level. Only use AutoSource recommendations for item/stores that have an adequate sales history. This measure is not displayed for causal levels.

Pick Optimal Level

Set only at final-levels, a check in this field (set to *True*) indicates that the batch forecast should use the Optimal Source Level selected by AutoSource. For more information on AutoSource, refer to the *Oracle Retail Demand Forecasting Implementation Guide*.

Note: This measure is not displayed for causal levels.

Return Forecast Flag Override

If this Boolean measure is checked, it indicates that a returns forecast will be generated for the desired product/location combinations.

Return Forecast Percentage Override

This measure overrides the default value of the percentage of the merchandise sold in a time period that is expected to be returned to the store. A value of zero indicates that no returns will be calculated.

Return Forecast Time Shift Override

This measure overrides the default value of the number of time periods between when the merchandise was sold until it was returned. For instance, if the measure is set to one, the merchandise is expected to be returned one week after it was sold.

Source Level Override

Set only at final-levels, the Source Level Override is 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. *No Override* appears in this field if the system default set in the Forecast Administration Workbook

is to be used.

Blending Parameter Override

Note: This measure is not displayed for baseline levels.

The blending parameter is used to combine promotion effects calculated at the pooled as well as item/store levels. This is the override of the value assigned in the Forecast Administration Workbook. The default setting for this parameter is No Override.

Forecast Maintenance Workbook: Advanced Final Level Parameters

The Advanced Settings workflow activity is used to override the dates that are used in the forecast generation process as well as historical start dates for any intersection at the final-level that varies from the default settings in the Forecast Administration Workbook.

Advanced Final Level Parameters Worksheet Measures

The Advanced Final Parameter worksheet contains the following measures:

Apply Promotional Lift Override

Note: This measure is not displayed for baseline levels.

Select this parameter (set to *True*) if you want RDF to ignore the system calculated lifts, and apply the overrides.

Overlapping Promotion Adjustment Factor Override

Note: This measure is not displayed for baseline levels.

The Overlapping Promotion Adjustment Factor Override measure is applied if Adjustable Approach is selected for Overlapping Promotion Behavior for the item/location level. This is an override of the global setting.

The Overlapping Promotion Adjustment Factor Override specifies at an item/location how the individually calculated promotion effects are combined with each other when the promotions are overlapped in the forecast horizon. The default value is 1.

Forecast End Date Override

This parameter represents the last point in time for which the forecasting engine forecasts for a particular intersection. Should this parameter be set to a date less than the Forecast Start Date plus the Forecast Length (in Forecast Administration), the engine forecasts zero (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. No value is in this measure if the system default set in the Forecast Administration Workbook is to be used.

Note: Changes to this measure can be seen in the Forecast Maintenance Workbook. The most recent commit (in either task) is the value used by the system.

Forecast Start Date Override

This parameter represents the date to start forecasting at a particular intersection. 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 change the Forecast Start Date once it is no longer in the future. No value is in this measure if the system default set in the Forecast Administration Workbook is to be used.

Note: Changes to this measure can be seen in the Forecast Maintenance Workbook. The most recent commit (between either task) is the value used by the system.

History Start Date Override

This parameter represents the first point in time from which the Forecasting Engine begins training and modeling (that is, if there are two years of history, but you only want to use one year, you set the start date to a year ago). This parameter overrides the History Start Date set in the Forecast Administration Workbook to the desired item/location intersection. For example, if you have a large spike in the first three weeks of sales for an item on sale, you can set the Historical Start Date to one week past that period, and those first few weeks are not 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, 2003, and an item/location does not have sales history until February 1, 2003, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

If this parameter is set into the future, there would be no forecast, as the history training window is read as zero.

Note: Changes to this measure can be seen in the Forecast Maintenance Workbook. The most recent commit (between either task) is the value used by the system.

Forecast Length Override

This parameter decides how many periods ahead RDF should forecast for a certain item/store. It is the override of the Forecast Length in the Forecast Administration workbook. Such a feature is desired, because not all merchandise require the same forecast horizon. For instance, the majority of items require a seven week long forecast to be used in replenishment. So the default value is set to seven weeks. However, a few items which need longer term planning and have long lead times, may need a forecast that is 52 weeks long. For those the length can be overwritten with the appropriate value.

Forecast Maintenance Workbook: Basic Source Level Parameters Worksheet

Measures: Basic Source Level Parameters Worksheet

The Basic Source Level Parameters worksheet contains the following measures:

History Start Date Override

This parameter represents the first point in time from which the Forecasting Engine begins training and modeling (that is, if there are two years of history, but you only want to use one year, you set the start date to a year ago). This parameter overrides the History Start Date set in the Forecast Administration Workbook to the desired item/location intersection. For example, if you have a large spike in the first three weeks of sales for an item on sale, you can set the Historical Start Date to one week past that period, and those first few weeks are not 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, 2003, and an item/location does not have sales history until February 1, 2003, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

If this parameter is set into the future, there would be no forecast, as the history training window is read as zero.

Note: Changes to this measure can be seen in the Forecast Maintenance Workbook. The most recent commit (between either task) is the value used by the system.

Forecast Method Override

Set at both final and source-levels, the Forecast Method Override is a list from which you can select a different forecast method than the Default Forecast Method set in the Forecast Administration Workbook. No Override appears in this field if the system default set in the Forecast Administration Workbook is to be used. Valid options depend on your system setup.

Bayesian Alpha Override

Note: This measure is not displayed for causal levels.

This is the override of the Bayesian Alpha parameter that is specified globally in the Forecast Administration Workbook. The override can be specified at every forecasting level, that is, final and source-levels.

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 one (1).

Moving Average Window Length Override

Note: This measure is not displayed for causal levels.

Used with Moving Average forecast method, this is the number of data points in history used in the calculation of Moving Average. This parameter, when set to a non-zero value, overrides the value in the Forecast Administration Workbook.

Forecast Like-Item & Sister-Store Workbook

The Forecast Like-Item & Sister-Store workbook provides the ability to model a new product's demand after an existing product. Forecasts can thus be generated for the new product based on the selected history or the forecast of the existing product plus an Adjustment Ratio. Likewise, the sales history or the forecast of existing store locations can be used as the forecast foundation for new locations.

Note: This workbook is only valid at final-levels. Therefore, this workbook may include hierarchy dimensions that are at higher positions than item or store (for example, subclass or region).

Creating a Forecast Like-Item & Sister-Store Workbook

1. Within the Local Domain, select **New** from the **File** menu.
2. Select the **Forecast** tab to display a list of workbooks. Select **Forecast Like-Item & Sister-Store**.
3. Click **OK**. The wizard opens and prompts you to select the level of the final-forecast. Depending on your organization's setup, you may be offered a choice of several final-forecast levels.
4. Select the final-forecast level to be viewed in the workbook.
5. Click **Next**.
6. Select the locations to include in the workbook.
7. Click **Next**.
8. Select the products to include in the workbook.
9. Click **Next**.
10. Select any additional measures (that is, measures not standard in the Forecast Like-Item & Sister-Store workbook) to be included. The measure options available in this window are set in the RPAS Security Administration workbook/Workbook Template Measure Rights worksheet.
11. Click **Finish** to display the workbook.

Worksheets for the Forecast Like-Item & Sister-Store Workbook

The following sections describe the components of the Forecast Like-Item & Sister-Store workbook.

The Forecast Like-Item & Sister-Store workbook includes these worksheets:

- Like-Item Worksheet
- Sister-Store Worksheet
- Final Advanced Parameter Worksheet

Like-Item Worksheet

The Like-Item worksheet is used to forecast a new item by modeling it after an existing item.

Figure 3–10 Final Like-Item Worksheet

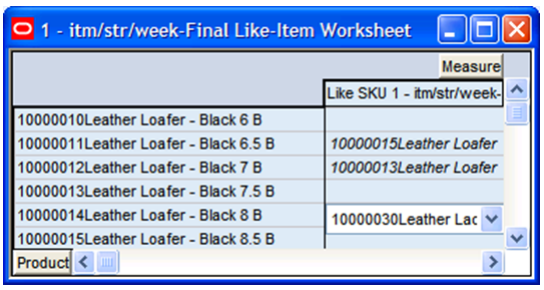

Measures: Like-Item Worksheet

The Like-Item worksheet contains the following measure:

Like SKU

The Like SKU field displays the items selected during the wizard process. In the example shown in the Final Like-Item 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 is used for the new item based on the parameters selecting in the Advance Parameter worksheet. In the previous example, Leather Loafer - Black 8 B is a new item, and Leather Lace-up Boat Shoes 8 B is being selected as the existing item that is used to Leather Loafer - Black 8 B model's forecast.

Note: When working with both the Like SKU and Sister-Store worksheets, making a selection in one of the worksheets requires Calculate to be run before the pick-list options are available in the other worksheet.

Sister-Store Worksheet

The Sister-Store worksheet is used to forecast demand for a new store by modeling it after an existing store.

Figure 3–11 Final Sister-Store Worksheet

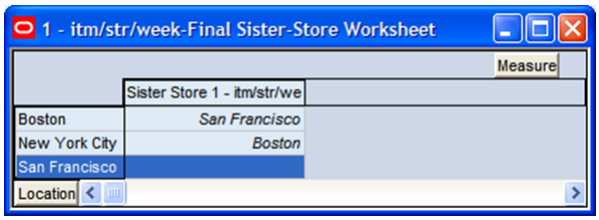

Measures: Sister-Store Worksheet

The Sister-Store worksheet contains the following measures:

Sister-Store

Displays the locations selected during the wizard process. In the example shown in Figure 3–10, 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 is used for the new store based on the parameters selecting in the Advance Parameter worksheet.

In Figure 3–11, the Boston is the new store and San Francisco is being selected as the store to be used for modeling Boston's forecast.

> **Note:** When working with both the Like SKU and Sister-Store worksheets, if a selection is made in one of the worksheets, Calculate must be run before the pick-list options are available in the other worksheet.

Final Advanced Parameter Worksheet

The Final Advanced Parameter worksheet is used to manage item/locations in which the Forecast Start Date, Forecast End Date, or History Start Date varies from the default settings in the Forecast Administration Workbook and to set the Adjustment Ratio for the new item or location being forecast.

Figure 3–12 Final Advanced Parameter Worksheet

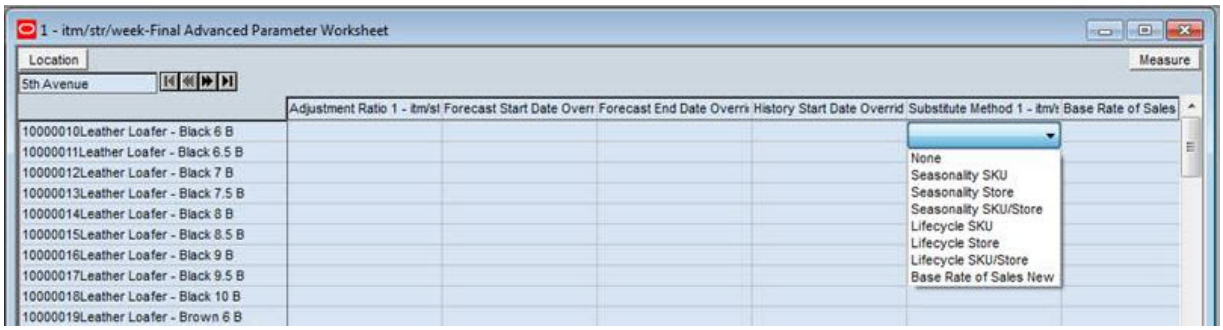

Measures: Final Advanced Parameter Worksheet

The Final Advanced Parameter worksheet contains the following measures:

Adjustment Ratio

You may enter an Adjustment Ratio to apply to the forecast for the new product/location combination displayed. This is a real number between [zero (0), infinity]. The default (NA) value is 1.00, (in other words 100%) which translates to no adjustment.

> **Note:** If demand for a new item is expected to be 30% greater than its Like-item, the Adjustment Ratio would be set to 1.30. If demand for a new item is expected to be 30% less than its Like-item, the Adjustment percent would set to 0.70.

Base Rate of Sales

This measure represents the average sales of a new item/store combination. It is used when specifying 'Base Rate of Sales New SKU' as the Substitution Method to create a forecast for the new item/store combination. The measure can be generated in another application and loaded into RDF, or it can be manually entered by a user.

Forecast Start Date Override

This parameter represents the date to start forecasting for an item/location combination. This parameter can be set in the future if using Like-item or Sister-Store

functionality, and, upon reaching that time, the forecast is 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 the Forecast Start Date for this intersection needs to be edited once it is no longer in the future. For Like item or Sister-Store, the Forecast Start Date and the History Start Date should be set to the same date.

No value is in this measure if the system default set in the Forecast Administration Workbook is to be used.

Note: This measure can also be set in the Forecast Maintenance Workbook. Changes to this measure can be seen in the Forecast Maintenance Workbook and the Forecast Like-Item & Sister-Store Workbook. The most recent commit (between either of the workbooks) is the value used by the system.

Forecast End Date Override

This parameter represents the last point in time for which the Forecasting Engine forecasts for an item/location combination. Should this parameter be set to a date less than the Forecast Start Date plus the Forecast Length (in the Forecast Administration Workbook), the engine forecasts 0 past this date. If Forecast End Date is more than Forecasting Start Date plus Forecasting Length, no forecast is generated outside Forecasting Start Date plus Forecasting Length window. 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. No value is in this measure if the system default set in the Forecast Administration Workbook is to be used.

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

Note: This measure can also be set in the Forecast Maintenance Workbook. Changes to this measure can be seen in the Forecast Maintenance Workbook and the Forecast Like-Item & Sister-Store Workbook. The most recent commit (between either of the workbooks) is the value used by the system.

History Start Date Override

This parameter represents the first point in time from which the Forecasting Engine begins training and modeling (that is, if there are two years of history, but only one year is required, set the start date to a year ago). This parameter overrides at the item/store level from the global settings in the Forecast Administration Workbook. This can be used to level out past sales. For example, if there is a large spike in the first three weeks of sales for an item was on sale, set the Historical Start Date to one week past that period. Those first few weeks are not 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, 2011 and an item/location does not have sales history until February 1, 2011, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

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 any of the Lifecycle Methods (refer to Substitute Methods); 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 begins using as its sales.

Note: This measure can also be set in the Forecast Maintenance Workbook. Changes to this measure can be seen in the Forecast Maintenance Workbook and the Forecast Like-Item & Sister-Store Workbook. The most recent commit (between either of the workbooks) is the value used by the system.

Substitute Methods

Displays a list from which you can select the substitute method. When a Substitute Method is used to forecast, the method set for an intersection is cleared once the Default Forecast Start Date is greater than the Forecast Start Date Override plus the Like TS Duration for the intersection. Valid options are listed in Table 3–3.

Table 3–3 Substitute Methods for Forecasting

Option	Description
None	There is no substitution for this product/location combination. This is the default value.
Seasonality SKU	You provide a Like-item that has a similar seasonality pattern that sells at the same store. The new product's forecast is the Like-item's demand forecast with the applied adjustment. The forecast is set to zero (0) for all dates before the new product's start date.
Seasonality 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 is the demand forecast of the same product at the Sister-Store with the applied adjustment. The forecast is set to zero (0) for all dates before the new store's open date.
Seasonality SKU/ 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 is the demand forecast of the Like-item at the Sister-Store with the applied adjustment. The forecast is set to zero (0) for all dates before the new product's start date after the new store opens.
Lifecycle SKU	You provide a Like-item that had a similar lifecycle pattern that sells at the same store. The new product's forecast is 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.

Table 3–3 (Cont.) Substitute Methods for Forecasting

Option	Description
Lifecycle 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 is 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 SKU/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 is 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.
Base Rate of Sales New SKU	You provide a base rate of sales for a new item/store combination. The new product's forecast is a combination of the seasonality of the forecast at the corresponding source level and the base rate of sales. Specifically, the formula to calculate the forecast is: *Forecast at time t = base rate of sales multiplied by a seasonal profile at time t* The seasonal profile is derived from the source level forecast. *Forecast at time t = source level forecast at time t divided by the average of the source level forecast times the base rate of sales*. The base rate of sale is a convenient way to generate (seasonal) forecasts for a new item. The base rate of sales does not require a Like SKU or Sister Store to generate a forecast. It requires a source level forecast. For example, source level should be an aggregation of product level like subclass or class, and an aggregation of location, like region. It also requires a value for how much the item is selling on average per period. That value works as the interim forecast used to spread the source level forecast down to final level. Because this item is new, the system knows that as you have specified a history start date. The base rate of sale forecast is generated until the TS duration is reached. The TS duration is the difference between forecast start date and history start date. Usually the default history start date is empty, which means it is the beginning of the calendar. If the history start date override is also empty, the forecast start date less the beginning of the calendar (>=2 years) is probably much more than the TS duration, so the item is not considered new. In this case base rate of sales or any new item method, will not work.

Required Steps for Forecasting Using Each of the Like SKU/Sister-Store 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 RDF stops using the substitution method and starts using the system forecast generated by the forecast engine.

Procedures

The following procedures outline the steps required for using each of the methods described in Table 3–3, " Substitute Methods for Forecasting":

Seasonality/SKU

Introduces a new item at an existing store (Like Item with a similar forecast):

1. From the Like-Item worksheet, select a Like Item from the list across from the new item.

2. From the Advanced Parameter worksheet, Set the Forecast Start Date for the new item at an existing store.
3. From the Advanced Parameter worksheet, Set the History Start Date for the new item at the existing store to the same date as the Forecast Start Date.
4. From the Advance Parameter worksheet, Set the Adjustment percent (optional) for the new item at the existing store.

Seasonality/STR

Introduces an existing item at a new store (Sister-Store with similar forecast):

1. From the Sister-Store worksheet, Select a Sister-Store from the list across from the new store.
2. From the Advanced Parameter worksheet, Set the Forecast Start Date for the existing item at the new store.
3. From the Advanced Parameter worksheet, Set the History Start Date for the existing item at the new store to the same date as the Forecast Start Date.
4. From the Advance Parameter worksheet, Set the Adjustment percent (optional) for the existing item at the new store.

Seasonality/SKU_STR

Introduces a new item at a new store (Like-item and Sister-Store with a similar forecast):

1. From the Like-Item worksheet, Select a Like-item from the list across from the new item.
2. From the Sister-Store worksheet, Select a Sister-Store from the list across from the new store.
3. From the Advanced Parameter worksheet, Set the Forecast Start Date at the intersection of the new item and the new store.
4. From the Advanced Parameter worksheet, Set the History Start at the intersection of the new item and new store.
5. From the Advance Parameter worksheet, Set the Adjustment percent (optional) at the intersection of the new item and new store.

Lifecycle/SKU

Introduces a new item at an existing store (Like item's sales history to be used as the forecast for the new item):

1. From the Like-Item worksheet, Select a Like-item from the drop-down list across from the new item.
2. From the Advanced Parameter worksheet, Set the Forecast Start Date for the new item at the existing store.
3. From the Advanced Parameter worksheet, Set the History Start Date for the new item at the existing store to the same date as the Forecast Start Date.
4. From the Advanced Parameter worksheet, Set the History Start Date for the Like item at the existing store to the point in its sales history that maps to the new item's forecast.
5. From the Advance Parameter worksheet, Set the Adjustment percent (optional) for the new item at the existing store.

Lifecycle/STR

Introduces an existing item at new store (Sister-Store's sales history to be used as the forecast for the new store):

1. From the Sister-Store worksheet, Select a Sister-Store from the list across from the new store.
2. From the Advanced Parameter worksheet, Set the Forecast Start Date for the existing item at the new store.
3. From the Advanced Parameter worksheet, Set the History Start Date for the existing item at the new store to the same date as the Forecast Start Date.
4. From the Advanced Parameter worksheet, Set the History Start Date at the intersection of the Sister-Store and existing item to the date in its sales history that maps to the new store's forecast.
5. From the Advance Parameter worksheet, Set the Adjustment percent (optional) for the existing item at the new store.

Lifecycle/SKU_STR

Introduces a new item at a new store (Like item's and Sister-Store's sales history to be used as the forecast for a new item at a new store:

1. From the Like-Item worksheet, Select a Like-item from the list across from the new item.
2. From the Sister-Store worksheet, Select a Sister-Store from the list across from the new store.
3. From the Advanced Parameter worksheet, Set the Forecast Start Date at the intersection of the new item and new store.
4. From the Advanced Parameter worksheet, Set the History Start Date at the intersection of the new item and new store to the same date as the Forecast Start Date.
5. From the Advanced Parameter worksheet, Set the History Start Date at the intersection of the Like item and Sister-Store to the date in its sales history that maps to the new item and new store's forecast.
6. From the Advance Parameter worksheet, Set the Adjustment percent (optional) at the intersection of the new item and new store.

Product/Location Cloning Administration Workbook

The Product/Location Cloning Administration workbook allows users to specify parameters to clone or copy history for new items or locations from existing items and locations. Since RDF needs history for forecasting, this functionality allows users to specify clone items and clone stores for these new Product/Locations.

Cloning is different from Like-item/Sister-Store functionality. Like-item/Sister-Store functionality copies forecasts or history of the Like-item or Sister-Store as the forecasts for the new item/store. However, in cases where the two items have different promotion calendars, it is necessary to forecast based on history data and promotion calendar. Item/Location Cloning functionality facilitates this process.

Also, by copying history from the parent product, users can start forecasting almost immediately after the launch of the new product.

Product/Location Cloning Administration can be performed from the Master or Local domain. The actual cloning or copying of history can be performed as part of batch run.

Note: Clone is available as a special expression that can be configured to be called as a batch run to perform the actual copying of history.

When the Cloning administration workbook is built in the Master domain, if clone assignments are made such that a product is assigned a clone that belongs to a different local domain, then at the time of batch run RDF does not clone that specific product and continue processing after logging a message.

Note: This workbook is only valid at final-levels. Therefore, this workbook may include hierarchy dimensions that are at higher positions than item or store (for example, subclass or region).

Using Product and Location Cloning

This section describes examples of cloning as illustrated in Figure 3–13.

An owner of multiple stores is opening a new store (Store N) and wants to replicate some of the items (#1-5) from an established store (Store A) to the new store (Store N).

Location cloning occurs when the owner clones the same items (#1-2) from Store A to Store N.

Product cloning occurs when the owner clones a similar item (#3) to (#5) from within Store N.

Product and location cloning occurs when the owner clones a similar item from Store A (#5) to Store N (#4).

Figure 3–13 Cloning Examples

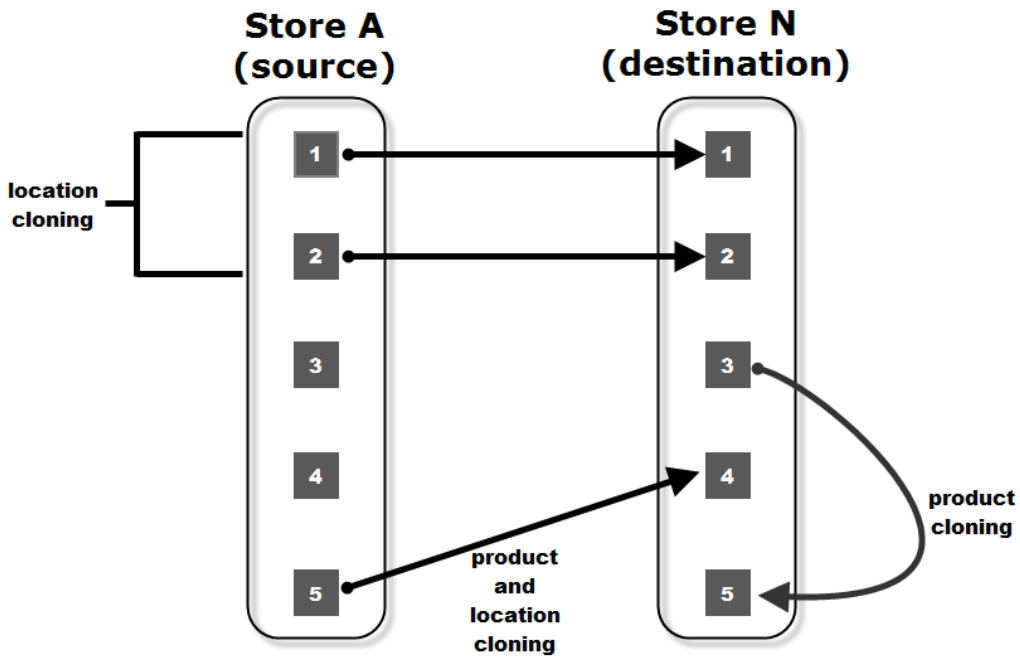

Worksheets for the Product/Location Cloning Administration Workbook

The following sections describe the components of the Product/Location Cloning Administration workbook.

The Product/Location Cloning Administration workbook has four tabs and includes the worksheets listed in Table 3–4.

Table 3–4 Tabs and Worksheets for the Product/Location Cloning Administration Workbook

Tab	Worksheets
Cloned History Adjustment	Adjustment Parameters Worksheet
Product Cloning	Product Cloning Worksheet
Location Cloning	Location Cloning Worksheet
Clone Mask	Clone Mask Worksheet

Adjustment Parameters Worksheet

The Adjustment Parameters worksheet allows users to specify parameters that determine how much the cloned history is adjusted to match the level of the actual sales.

Figure 3–14 Adjustment Parameters Worksheet

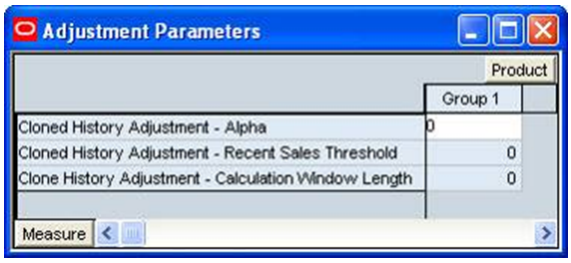

Measures: Adjustment Parameters Worksheet

The Adjustment Parameters worksheet contains the following measures:

Cloned History Adjustment - Alpha (Range 0-1)

This parameter is used to calculate the ratio that is applied to the cloned history. If this parameter is close to zero (0), then the cloned sales will be scaled so that they are aligned with the sales level. If this parameter is close to one (1), then little scaling of the cloned history will occur.

Cloned History Adjustment - Recent Sales Threshold

This parameter represents a specific number of periods. If it is less than the number of periods from the first sale date to the current date, then the cloned history is adjusted. If it is greater than the number of periods from the first sale date to the current date, then no adjustment is made to the cloned history.

Cloned History Adjustment - Calculation Window Length

This parameter is the number of periods that are considered when adjusting the cloned history.

Figure 3–15 Adjustment Ratio Formula

$$cloning\ adjustment\ ratio = \left(\frac{average\ of\ actual\ sales}{average\ of\ cloned\ sales\ over\ past\ M\ periods}\right)^{(1-alpha)}$$

Figure 3–16 Cloned History Formula

$$cloned\ history = adjustment\ ratio * unadjusted\ cloned\ history$$

Product Cloning Worksheet

The Product Cloning worksheet allows users to make Clone Product assignments for new Products. From this worksheet, users can specify up to three parent Products for a given Product that is being cloned. Also, it allows users to specify different parent Products for different Locations. For example, CS Spring Water 1 sells like Diet Soda single in North America, but sells like Regular Soda in Asia. The Product Cloning worksheet thus is at Product and Location intersection.

Users have the ability to specify associated contributions for each of the Clone Products and have the ability to specify an overall Adjustment Ratio.

Figure 3–17 Product Cloning Worksheet

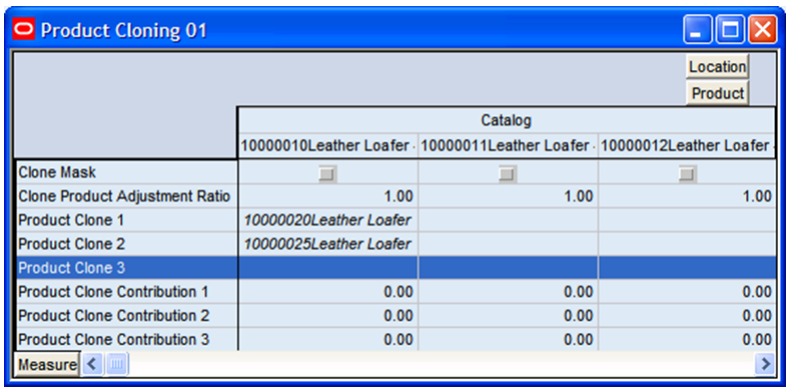

Measures: Product Cloning Worksheet

The Product Cloning worksheet contains the following measures:

Product Clones

Product Clone1, Product Clone 2, Product Clone 3

Parent Products for the new Product being cloned. It is optional to specify more than one product clone. The selection or assignment of product clone is made by clicking this measure, and a Single Hier Select wizard opens. Assignments are made by selecting a product and clicking OK. Refer to the *Oracle Retail Predictive Application Server User Guide* for more details on the Single Hier Select wizard.

Figure 3–18 Product Clone Selection

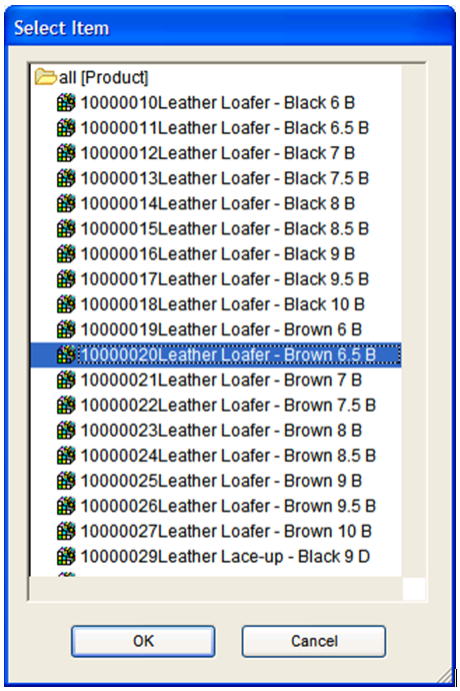

Product Clone Contributions

Product Clone Contribution1, Product Clone Contribution2, Product Clone Contribution3

Relative contributions of all clone products. If user does not specify contributions, but makes selections for product clones and commits the changes, at the time of commit the contributions are calculated to have equal weights for all specified product clones. If user specifies Product clone contributions, they get normalized such that they add up to one at the time of commit. Product clone contributions with no associated product clone values get set to zero (0) at the time of commit.

Clone Product Adjustment Ratio

You may enter an Adjustment Ratio to apply to the cloned history of the new product. This is a real number between [zero (0), infinity]. The default (NA) value is 1.00 (in other words 100%), which translates to no adjustment.

Example 3–2 Clone Product Adjustment Ratio

If demand for a new item is expected to be 30% greater than its clone item, the Adjustment Ratio would be set to 1.30. If demand for a new item is expected to be 30% less than its clone item, the Adjustment percent would set to 0.70.

Location Cloning Worksheet

Location Cloning worksheet allows users to make clone Location assignments for new Locations. Location Cloning worksheet allows users to specify up to three Clone Locations for a given Location that is being cloned. Also, it allows users to specify different parent Locations for different products. For example, the Beverly Hills store sells like the Miami store for swimwear and sells like the New York store for electronics. The Location Cloning worksheet thus is at Location and Product intersection.

Users have the ability to specify associated contributions for each of the Parent Locations and have the ability to specify an overall Adjustment Ratio.

Figure 3–19 Location Cloning Worksheet

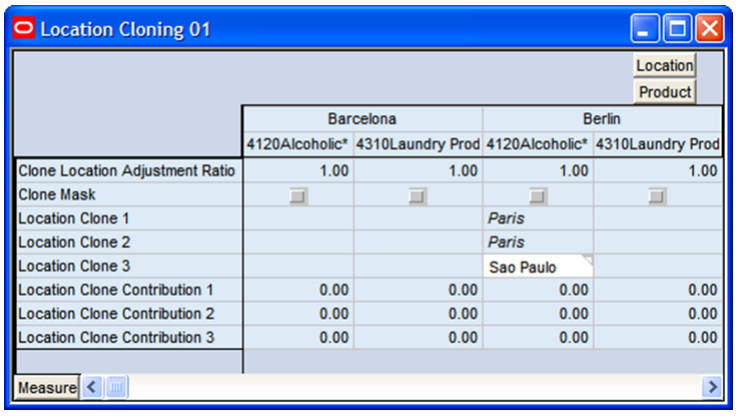

Measures: Location Cloning Worksheet

The Location Cloning worksheet contains the following measures:

Location Clones

Location Clone1, Location Clone 2, Location Clone 3

Parent Locations for the new Location being cloned. It is optional to specify more than one parent clone. The selection or assignment of Location Clone is made by clicking this measure, and a Single Hier Select wizard opens. Assignments are made by selecting a Location and clicking **OK**. Refer to the *Oracle Retail Predictive Application Server User Guide* for more details on the Single Hier Select wizard.

Figure 3–20 Select Location Clone

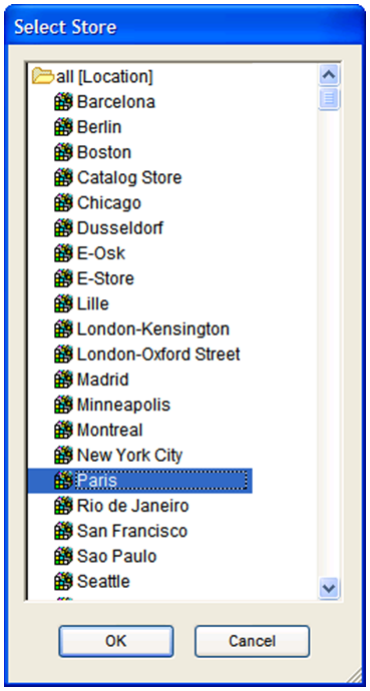

Location Clone Contributions

Location Clone Contribution1, Location Clone Contribution2, Location Clone Contribution3

Relative contributions of all Clone Locations. If user does not specify contributions, but makes selections for Location Clones and commits the changes, at the time of commit the contributions are calculated to have equal weights for all specified Location Clones. If user specifies contributions, at the time of commit, Location clone contributions get normalized so that they add up to one (1). Location Clone contributions with no associated Location Clone values get set to zero (0) at the time of commit.

Clone Location Adjustment Ratio

You may enter an Adjustment Ratio to apply to the cloned history of the new location. This is a real number between [zero (0), infinity). The default (NA) value is 1.00 (in other words 100%), which translates to no adjustment.

Example 3–3 Clone Location Adjustment Ratio

If demand for a new store is expected to be 30% greater than its clone store, the Adjustment Ratio would be set to 1.30. If demand for a new store is expected to be 30% less than its clone store, the Adjustment percent would set to 0.70.

Note: Adjustment Ratios and Clone Contributions specified in the Product Cloning and Location Cloning worksheets are used together while evaluating the result for a given Item/store.

Example 3–4 Adjustment Ratios and Clone Contributions

Item1	
Item2	**Item3**
20%	80%
Adjustment Ratio = 1	

STR1:	
STR2	**STR3**
50%	50%
Adjustment Ratio = 0.5 The contributions are calculated as: Item2/STR2 at 5% (=20% x 0.5x50%); Item2/STR3 at 5%;Item3/STR2 at 20%(=80%x0.5x50%);Item3/STR3 at 20%	

Clone Mask Worksheet

This worksheet allows users to enable cloning for product/locations by setting the Clone Mask measure to True.

Note: Clone Mask measure can also be brought in using the Show/Hide option in the Product Cloning and Location worksheets. The clone mask can also be set based on business rules by configuring RPAS rules.

Figure 3–21 Clone Mask Worksheet

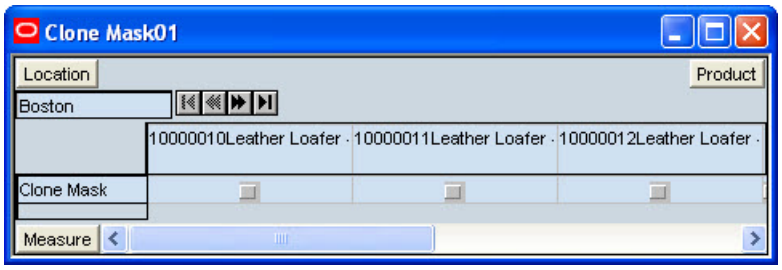

Measures: Clone Mask Worksheet

The Clone Mask worksheet contains the following measures:

Clone Mask

This measure specifies if cloning should be performed for the item/store combination.

4

Generating and Approving a Forecast

Once a user has completed setting all global and individual forecast parameters a forecast must be generated and approved. The Forecast workbook group provides the workbooks necessary to support these tasks.

Run a Batch Forecast

The forecast generation process creates demand forecasts for all product/location combinations that are set to forecast within the forecast horizon window. Forecasts are typically run automatically as scheduled batch jobs. RDF 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 supported in part through 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 support the automatic scheduling of batch forecasting jobs in the RDF solution. Refer to the Forecast Administration Workbook description for further information.

Forecast Batch Run Workbook

The Forecast Batch Run Workbook wizard allows you to manually execute the forecast generation process at a time other than the regularly scheduled batch job. If a global domain environment is implemented, forecasts generated in the Local domain can be viewed in the Master domain; however this forecast is isolated to the data in the Local domain. The execution of PreGenerateForecast at the Master domain, then passing the output of the process to Generate from the back-end of each Local domain allows for the Local domains to share a birth date, thus supporting a worksheet to forecast data across Local domains in the Master domain. Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on PreGenerateForecast and Generate.

Running a Batch Forecast Manually

The following sections describe the procedures to manually run a batch forecast.

Before You Begin

Prior to using the Forecast Batch Run Workbook wizard, at minimum the following tasks must be performed:

1. Create or access a Forecast Administration Workbook.

2. In the Forecast Start Date measure, enter the starting date of the forecast horizon; otherwise the system defaults to the system date (today).
3. Set a Default Forecasting Method for the Final and Source Level.
4. Set a Default Source Level for the Final Level.
5. Set the Forecast Length if there is no value already in this measure.
6. Set the Spreading Profile if the Source Levels are at aggregate dimensions higher than the Final Level.
7. Commit any changes by selecting from the File menu, then select Commit Now.
8. Close the workbook.

Procedure

Follow these steps to manually run a batch forecast:

1. Within the Local domain, select **New** from the File menu.
2. Select the Forecast tab to display a list of workbooks for statistical forecasting. Highlight Forecast Batch Run and click **OK**.
3. The Forecast Batch Run wizard opens and prompts you to select the Final Levels to forecast. Select **Next** or **Finish**.

The Run Batch wizard automatically executes PreGenerateForecast and Generate within the Local Domain. If **Next** is selected from the last wizard window, the wizard does not advance to the completion message until the forecast has been generated. Depending on the amount of product/locations to be forecast and the forecast horizon, it may take a several minutes before the system advances to the final window. When the forecast generation is completed, the wizard displays a window that notifies you of the forecast generation ID.

After a forecast is generated, the Forecast Start Date field is cleared. This ensures that the same forecast is not generated again on the same date. The Next Run Date field is also updated based on the birth date of the forecast plus the Forecast Cycle. Both fields can be viewed in the Forecast Administration Workbook.

Delete Forecasts

The Days to Keep Forecasts parameter (set in the Forecast Administration Workbook) supports the automatic deletion of old forecasts when the Run Batch wizard or PreGenerateForecast is executed. Occasionally, a user may need to manually delete a forecast.

Some reasons might include:

- A forecast was run with the wrong source-levels selected.
- The horizon was not properly set.
- Old forecasts need to be deleted to save space on the server.

The Delete Forecasts wizard guides you through the process of deleting unwanted forecasts from the system. Deletions of forecasts are permanent.

Forecast Delete Workbook

Note: The deletion of forecasts from RDF is permanent.

If the Forecast Delete wizard is used in a global domain environment. Deleting a forecast in the Master domain deletes the selected birth date within all domains. Deleting a forecast in the Local domain deletes the selected birth date only within the Local domain.

1. From the Master or Local domain, select **New** from the File menu.
2. Select the Forecast tab to display a list of workbooks for demand forecasting.
3. Select **Forecast Delete**.
4. Click **OK**.
5. Select the forecast generation dates of the forecast to delete.

Figure 4–1 Forecast Delete Wizard

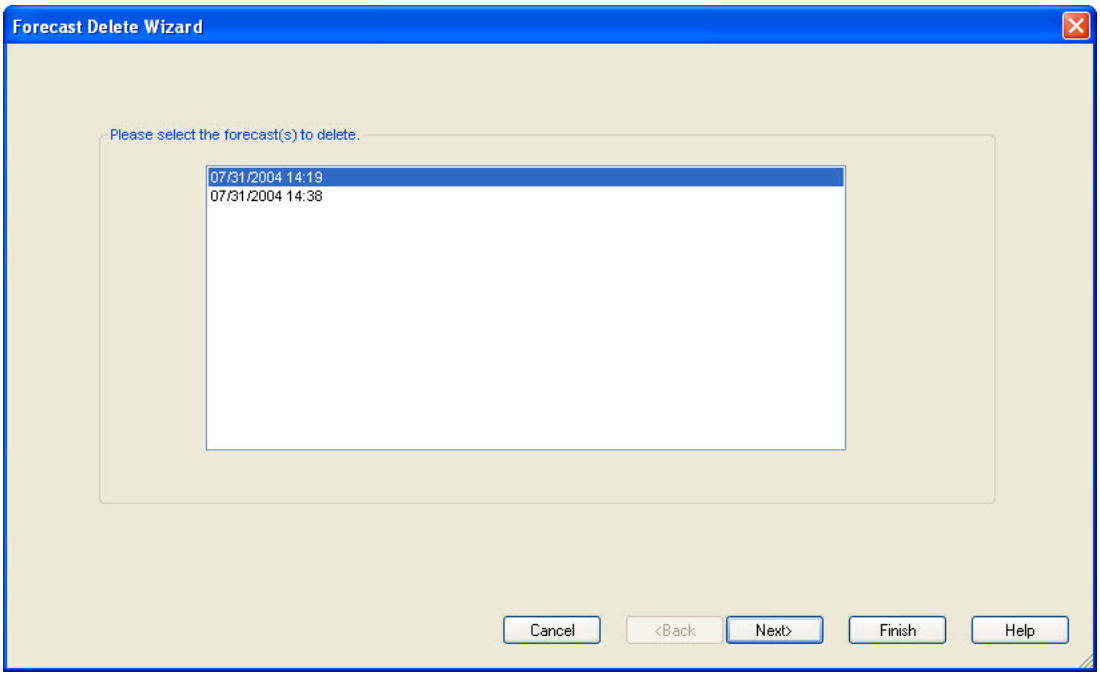

6. Click **Next**.
7. Select **Yes** to verify forecast deletion, or **No** to cancel deletion and exit the Forecast Delete wizard.
8. Click **Finish** to process the request. If **Yes** was selected on the final wizard window, the forecast is deleted.

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 forecast data can be exported to other processes, such as replenishment programs. The Forecast Approval Workbook allows you to view, analyze, adjust, and approve forecast results.

Some system forecasts may be set to be automatically approved by the system. The default approval method for items in a forecast is set in the Forecast Administration Workbook, and these policies can be amended for individual product/location combinations in the Forecast Maintenance Workbook. Any forecasts not set to

Automatic Approval may require evaluation, adjustments, and ultimately approval before subsequent processes are run.

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 are made on the appropriate worksheets in the Forecast Approval Workbook.

The Forecast Approval Workbook includes these worksheets:

Table 4–1 Worksheets for the Forecast Approval Workbook

Worksheet	Description
Advanced Final Level Worksheet	This worksheet displays all the regular price and promotional forecasting components.
Final Level Worksheet	This worksheet allows you to review final-level system-forecast quantities and make revisions to them if needed.
Source Level Worksheet	This worksheet displays the system-generated source-level forecast and allows you to compare this data with final-level forecast values.
Approval Worksheet	This worksheet allows you to specify the manual approval policy of forecasts by product and location.
Final System Parameters Worksheet	This option is only available if Generate System Parameters or Generate Methods is turned on in the Forecast Administration Workbook.
Source System Parameters Worksheets	This option is only available if Generate System Parameters or Generate Methods is turned on in the Forecast Administration Workbook, and a Source Level was designated in Forecast Administration.
Valid Forecast Run Worksheet	This worksheet allows you to review the partition dimensions in which the generated forecast was run.

When the Forecast Approval Workbook is displayed, you may review the system-generated forecast and measures for any levels included in the workbook and make adjustments to forecast values at the final-level. Forecast values are overwritten in the Adjusted Forecast measure on 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.

Opening or Creating the Forecast Approval Workbook

Select **Open** from the **File** menu to bypass the Forecast Approval wizard and open an existing Forecast Approval Workbook or perform the following:

1. Within the Master domain or Local domain, select **New** from the File menu.
2. Select the Forecast tab to display a list of workbooks for statistical forecasting.
3. Select Forecast Approval. Click **OK**.
4. 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**.

5. Select the forecast levels to include in the Forecast Approval Workbook. Select as many forecast levels as necessary for comparison. Click **Next**.

   Note: In the Source Level wizard, the source levels for a Causal Final Level are not displayed, only the Causal Final Level is displayed.

6. Select the birth date of the forecast you wish to approve.
   - Select **Use the most recently generated forecast** to build a workbook containing the most recent forecast values.

     Note: Use the most recently generated forecast - it must be selected if the task supports an AutoTask build.

   - Click **Select from a list of forecast** to select from a list of previously generated forecasts stored in the system. Note that if this option is selected, the system automatically highlights the first forecast in the list, and if no further selections are made by you, this first forecast is used.
   - If no selection is made by you, the birth date defaults to the latest generated generation ID. Click **Next**.
7. Select the specific locations you want to view. Click **Next**.
8. Select the merchandise you want to view. Click **Next**.
9. Select the first date of history to include in the workbook. You may choose to either set the **Forecast Start Date minus the number of periods** or select the first date of history from the displayed list. It is recommended that the **Forecast Start Date minus the number of periods** is set if the workbook is generated using AutoWorkbook build. Click **Next**.

   Note: If you make no selection, the system uses the forecast start date.

10. Select the last date in the forecast horizon to include in the workbook. You may either choose to **Include the following number of time periods** of the forecast horizon or select the last date to include of the horizon from the displayed list. It is recommended that **Include the following number of time periods of the forecast horizon** is set if the workbook is generated using AutoWorkbook build. The value entered must be within the forecast horizon. Click **Next**.
11. Select the last date in the calendar to include in the workbook. You may choose to either set the **Forecast end date plus the following number of time periods** or select the last date to include of the post-horizon calendar from the displayed list. It is recommended that **Forecast end date plus the following number of time periods** is set if the workbook is generated using AutoWorkbook build. Click **Next**.

    Note: If you make no selection, the system uses the last date in the forecast horizon chosen by you in the previous wizard window.

12. Place check marks next to any additional registered measures you would like to view in your workbook. The valid values of these measures may only be viewed if **Generate Intervals, Generate Cumulative Intervals, Generate Methods, Generate Parameters, Generate Baselines**, or **Store Interim Forecasts** were selected in the Forecast Administration Workbook. Click **Next**.

> **Note:** If any of these parameters are included in the Forecast Approval Workbook using AutoWorkbook build and are later deselected for generation in the Forecast Administration Workbook, AutoWorkbook must be set again.

13. From the list provided, select any additional measures beyond the default measures in the workbook that you would like to view. Click **Finish** to build and open the workbook.

Worksheets for the Forecast Approval Workbook

The following sections describe the components of the Forecast Approval Workbook.

The Forecast Approval Workbook includes these worksheets:

- Advanced Final Level Worksheet
- Final Level Worksheet
- Source Level Worksheet
- Approval Worksheet
- Final or Source System Parameters Worksheet
- Valid Forecast Run Worksheet

Advanced Final Level Worksheet

This worksheet displays all the regular price and promotional forecasting components. These include Self Lifts, as well as Halo and Cannibalization Effects for Regular Price and Promotions. This information is read-only.

Figure 4–2 Advanced Final Level Worksheet

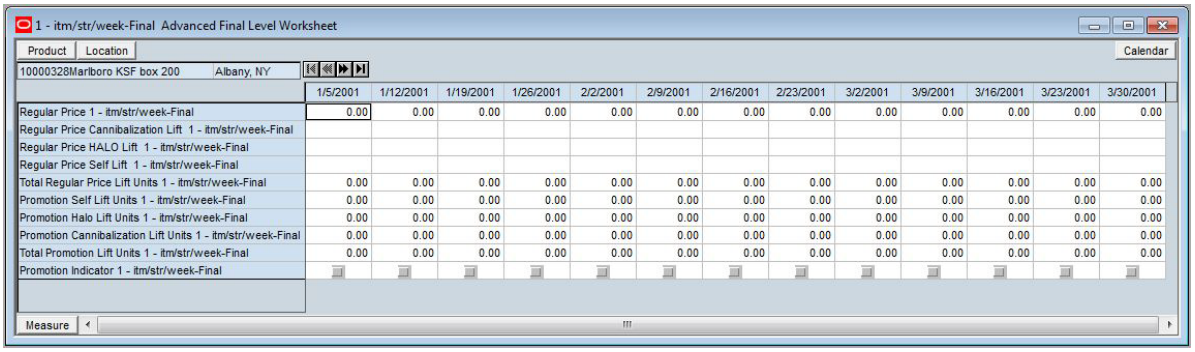

Measures: Advanced Final Level Worksheet

The Advanced Final Level worksheet contains the following measures:

Promotion Cannibalization Lift Units

This measure displays the negative effect, measured in units, that the promotion of certain items has on this item. The Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1.When the Interaction function is enabled, you should expect to see the displayed lift unit adjusted by the Interaction Factor calculation.

Promotion Halo Lift Units

This measure displays the positive effect, measured in units, that the promotion of certain items has on this item.If the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1, the displayed lift is the result of the combination of various lift types.When the Interaction function is enabled, you should expect to see the displayed lift unit adjusted by the Interaction Factor calculation.

Promotion Indicator

If this measure is checked (True), it indicates that a given item/location has any promotion active for a time period

Promotion Self Lift Units

This measure displays the effect that a promotion has on the sales units of an item.If the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1, the displayed lift is the result of the combination of various lift types.When the Interaction function is enabled, you should expect to see the displayed lift unit adjusted by the Interaction Factor calculation.

Regular Price

This measure displays the price for which the item has sold at the given store. This is a time-phased measure, and you can see how a change in price drives the regular price effects.

Regular Price Cannibalization Lift

This measure displays the negative effect, measured in units, that a price change of certain items has on this item. If the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1, the displayed lift is the result of the combination of various lift types.When the Interaction function is enabled, you should expect to see the displayed lift unit adjusted by the Interaction Factor calculation.

Regular Price Halo Lift

This measure displays the positive effect, measured in units, that a price change of certain items has on this item. If the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1, the displayed lift is the result of the combination of various lift types.When the Interaction function is enabled, you should expect to see the displayed lift unit adjusted by the Interaction Factor calculation.

Regular Price Self Lift

This measure displays the effect that a price change has on the sales units of an item. If the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1, the displayed lift is the result of the combination of various lift types.When the Interaction function is enabled, you should expect to see the displayed lift unit adjusted by the Interaction Factor calculation.

Total Promotion Lift Units

This measure displays the total effect that promotions can have on an item's sales units. It is the combination of Promotion Self Lift Units, Promotion Halo Lift Units and Promotion Cannibalization Lift Units. The default way to combine the promotional lifts is to sum them up. However, if the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than one, the Total Promotion Lift Units is not the straight summation of the three promotion lift measures.

Note: If the measure, Apply Promotional Lift Override, is selected in the Advanced Final Parameter Worksheet, then RDF ignores the system calculated lifts and applies the Aggregated Promotion Lift Override. It only applies the override in the week that has a promotion.

Total Regular Price Lift Units

This measure displays the total effect that regular price changes can have on an item's sales units. It is the sum of Regular Price Self Lift, Regular Price Halo Lift, and Regular Price Cannibalization Lift.

Final Level Worksheet

The Final Level worksheet allows you to review the forecast quantities and make adjustments to forecasts if needed. The primary 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, approve forecasts that have been user-adjusted or require manual approval.

Figure 4–3 Final Level Worksheet

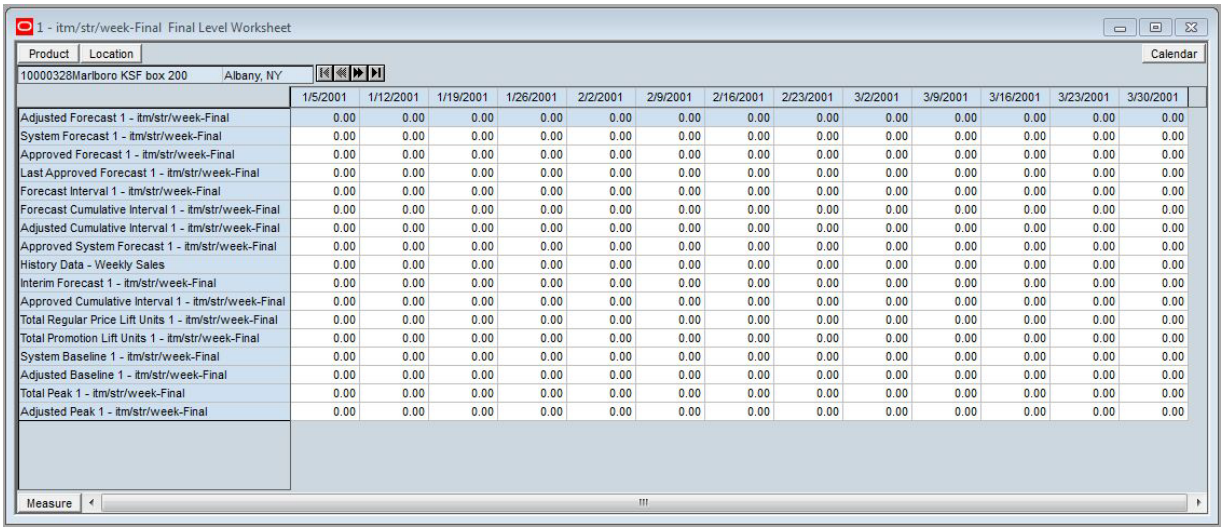

Measures: Final Level Worksheet

The Final Level worksheet contains the following measures:

Adjusted Cumulative Interval

When changes are made to the Adjusted Forecast, the value of the Forecast Cumulative Interval is recalculated in this measure. The values in this measure are read-only. To view and store this measure:

- Generate Cumulative Intervals must be selected in the Forecast Administration Workbook.
- Cumulative Intervals must be selected to be viewed in the Forecast Approval wizard.

Adjusted Forecast

The value in this field initially defaults to the System Forecast if the forecast is automatically approved by the system or through a Forecast Approval Alert. Otherwise, the value in the Adjusted Forecast is different than the System Forecast if:

- Adjusted Forecast has been updated by you.
- Keep Last Changes is set to Total, Difference or Ratio.
- Update Last Weeks Forecast is set to Replicate.

Changes to the Adjusted Forecast are not automatically approved Make all of your changes to the forecast in the Final Level view. To approve, go to the Approval view. and run the Approve Forecast custom menu. This updates your values into the Approved Forecast. measure with the adjusted values. and additionally update the date of the adjustment and the name of the user who made the adjustment within the Approval view.

Note: Changes to the Adjusted Forecast for periods outside of the forecast horizon are not committed.

Edits to any non-committed values in the Forecast Approval Workbook are overwritten when data is refreshed.

Approved Cumulative Interval

The Approved Cumulative Interval is the cumulative interval that was approved at the time of the workbook build. The values contained in this measure are read-only. If changes are made to the Adjusted Forecast, the Approved Cumulative Interval reflects the recalculated values when the workbook is committed and data is refreshed.

Approved Forecast

The Approved Forecast is the forecast quantity that was approved at the time of the task build. The values contained in this measure are read-only. If changes are made to the Adjusted Forecast, then go to the Approval view. and run the Approve Forecast custom menu. This updates your values into the Approved Forecast. measure with the adjusted values.

Approved System Forecast

The Approved System Forecast is populated with the last System Forecast approved for a time series:

1. Approval Method set to Automatic Approval
2. Approval Method set to a Forecast Approval Alert and the alert is not triggered
3. Update Last Weeks Forecast is set to Use Forecast

The values in this measure are read-only.

Demand Transference Lift Units

This measure displays the total lift caused by assortment changes. This is a time-phased measure, and you can see how assortment changes drive the demand transference lift in unit.

If the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1, the displayed lift is the result of the combination of various lift types.When the Interaction function is enabled, you should expect to see the displayed lift unit adjusted by the Interaction Factor calculation.

Forecast Cumulative Interval

Cumulative Intervals are used in safety stock calculation within allocation and replenishment systems. This value is similar to a running total of the Forecast Interval and is read-only. To see this measure:

- Generate Cumulative Intervals must be selected in the Forecast Administration Workbook.
- Cumulative Intervals must be selected to be viewed in the Forecast Approval wizard.

Forecast Interval

The Forecast Interval is calculated on the particular forecast region as capped standard deviation for some methods. It takes into consideration the system forecast for capping as well.

The forecast interval for any given week i is capped to a maximum of:

Forecast(*i*) * upperCapRatio

This cap is applied once the interval exceeds the system forecast, and it is applied for every value of the interval after that. The upperCapRatio (Interval Cap Upper Ratio) is a parameter set in the Forecast Administration Workbook.

The forecast interval is capped to a minimum of Forecast(i) * lowerCapRatio whenever the interval is less than this quantity. The lowerCapRatio (Interval Cap Lower Ratio) is another parameter set in the Forecast Administration Workbook.

When forecasts are approved, the corresponding Intervals and Cumulative Interval measures are approved. To see this measure:

- Generate Intervals must be selected in the Forecast Administration Workbook.
- Intervals must be selected to be viewed in the Forecast Approval wizard.

Forecast vs. Last Approved Forecast Workbook Alert

This measure is the target measure of a real time alert. It compares current forecast with the forecast generated last period, and if the difference is more than a threshold, the alert is triggered which is visible by the different format of the cells.

Forecast vs. Last Year Sales Workbook Alert

This measure is the target measure of a real time alert. It compares forecast with recent/ past sales, and if the difference is more than a threshold, the alert is triggered which is visible by the different format of the system forecast cells. This measure is only relevant for a baseline level.

Causal Peaks Workbook Alert

This measure is the target measure of a real time alert. It compares current forecast with the maximum historical demand. If the difference is positive, the alert is triggered which is visible by the different format of the system forecast cells. This measure is only relevant for a causal level.

History Data

History Data (a read-only measure) is the sales data used to generate the forecast. This allows you to compare Actuals to forecast values. When the workbook is created, the Data Source measure is copied into History Data.

Interaction Multiplier

The measure displays the interaction factor's impact when the Default Apply Interaction Factor is set to True and the Default Interaction Factor is set to a value other than 1. The value in this field equals to the combined total lift divided by the linearly added up lifts

If Apply Interaction Factor is set to False, or the Interaction Factor is set to 1, the value in this field will equal to 1.

Interim Forecast

The Interim Forecast is the forecast generated at the final-level that is used as the Data Source in Curve to produce the Spreading Profile. This profile determines how the Source Forecast is spread down to the Final Forecast level. The values in this measure are read-only. To see this measure:

- Generate Interim Forecast must be selected in the Forecast Administration Workbook.
- Interim Forecast must be selected to be viewed in the Forecast Approval wizard.

Last Approved Forecast

The Last Approved Forecast is the approved system forecast value when an approval occurred on a previous batch forecast for the time series. The values contained in this measure are read-only.

Return Forecast

If the Return Forecast option is enabled, this measure displays the units per week that are returned to the store.

System Baseline

This metric is an input to causal forecasting. It is set or specified in the Causal Baseline measure in the Forecast Administration workbook. It can be interfaced in the solution, or, more likely, it is also generated in RDF on another forecast level.It is generated using past sales data that contains no promotions (that is, normal demand given no causal effects).

System Forecast

The System Forecast displays the system-generated forecast for the time series. The values contained in this field are read-only. This measure includes baseline as well as promotional lifts and regular price lifts, if applicable.

Total Peak

This measure stores the total lift that an item can have on top of its baseline demand. The total lift is the sum of the Total Promotion Lift Units and other effects that are available to be applied on top of the baseline.

Total Promotion Lift Units

This measure displays the total effect that promotions can have on an item's sales units. A period can have one event active. For example, an item is advertised in a flyer and then the application of the promotion on the baseline is straightforward. However, for a different period the item can be advertised in a flyer, in the store, and have a price discount. Then the total promotion lift is not the straight summation of the three promotion lift measure.

Source Level Worksheet

The Source Level worksheet displays the system-generated source-level forecast. Final level forecast values in the Final Level worksheet can be viewed alongside and compared with their corresponding source-level forecasts.

Figure 4–4 Source Level Worksheet

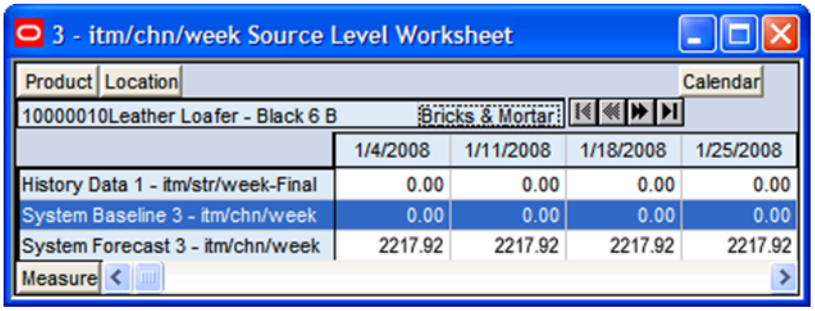

Measures: Source Level Worksheet

The Source Level worksheet contains the following measures:

History Data

History Data (a read-only measure) is the sales data used to generate the forecast. This allows you to compare Actuals to forecast values. When the workbook is created, the Data Source measure is copied into History Data.

System Baseline

The System Baseline is a forecast generated on past sales data that contains no promotions (that is, normal demand given no causal effects). To see a generated Baseline Forecast:

- Promote must be implemented.
- Generate Baseline must be selected in the Forecast Administration Workbook.
- The System Baseline must be selected to be viewed in the Forecast Approval wizard.

System Forecast

The System Forecast displays the system-generated forecast for the time series. The values contained in this field are read-only.

Approval Worksheet

The Forecast Approval worksheet allows for non-adjusted System Forecast to be approved. This worksheet can also be used to view the approval date of forecast values and to display the name of the user that 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.

Figure 4–5 Approval Worksheet

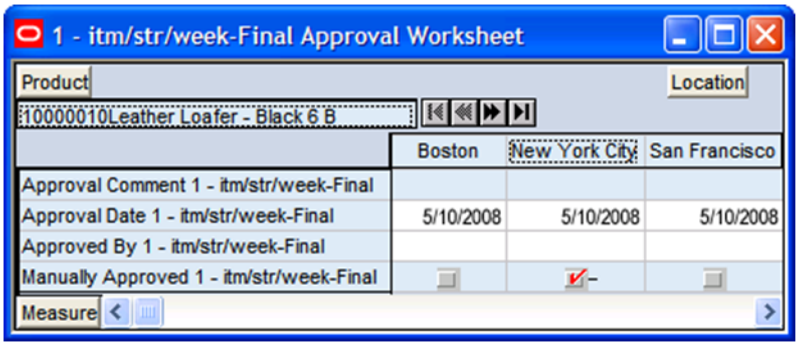

Measures: Approval Worksheet

The Approval worksheet contains the following measures:

Approval Comment

Approval Comment is a field in which notes may be entered regarding the forecast values or any pertinent information for specified product/location combinations.

Approval Date

Approval Date is a read-only field that displays the date that the forecast quantity is approved either automatically during the batch forecast process or when changes are made to the Adjusted Forecast. This information is necessary for RDF to carry out any subsequent processes, such as replenishment procedures.

Approved By

Approved By is a read-only field that displays the name of the user to approve forecasts for an item/location. This field may be populated with Sys if the system was set to automatically approve forecasts during the batch forecast process.

Manually Approved

Manually Approved is a Boolean flag. This field may be checked if you want to accept the System Forecast quantity for a time series that has yet to be approved. The flag is also activated when a change is made to the Adjusted Forecast and *Calculate* occurs. When this flag is activated, the time series for an item/location are approved and both the Approval Date and Approved By fields are updated.

Final or Source System Parameters Worksheet

The Final or Source Parameters worksheets allow you to worksheet relevant information for the forecast methods used during forecast generation. The information is different for forecast levels where baseline forecast is generated versus levels where causal is run. These parameters are only available to be viewed if Generate Parameters or Generate Methods are available in the Advanced settings of the Forecast Administration Workbook.

System Parameters Worksheet for Baseline Forecast For baseline levels the worksheet shows the alpha, level, and trend parameters for each fitted time series.

Measures: Final or Source System Parameters Worksheet for Baseline Forecast The Final or Source Parameters for Baseline Forecast worksheet contains the following measures.

To see these measures:

- **Generate Parameters** must be selected in the Forecast Administration Workbook.
- **System Generated Parameters** must be selected to be viewed in the Forecast Approval wizard.

Fit Error Factor

This is the relative error calculated at source level and transformed to become the error factor for the final level.

The transformation is a two step process. First, the source level error is replicated at the final level intersection. Then it is multiplied with a profile - still at the final level intersection, to add some variability for each product/location combination.

The value is multiplied with the forecast to create the confidence intervals.

Forecast Method Picked

This is the system-calculated alpha value (which is an internal optimization parameter that 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 Generated Alpha

This is the system-calculated alpha value (which is an internal optimization parameter that 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 Generated Level

This is 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 Generated Gamma

This is the system-calculated gamma value 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 Generated Trend

This is 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.

Final or Source System Parameters for Causal Forecast Worksheet This worksheet shows the system parameters for a causal forecast.

Measures: Final or Source System Parameters Worksheet for Causal Forecast The Final or Source Parameters for Causal Forecast worksheet contains the following measures.

Since no forecast is generated at the causal pooling levels, the information on this worksheet is relevant only for the final forecast level

The Final Parameters for Causal Forecast worksheet contains the following measures.

To see these measures:

- **Generate Parameters** must be selected in the Forecast Administration Workbook.
- **System Generated Parameters** must be selected to be viewed in the Forecast Approval wizard.

Fit Error Factor

This is the relative error calculated between the historical demand and the backcast. The backcast is calculated as the regression intercept, as a proxy for the baseline and the causal information applied on top of it.

Overlapping Promotion Adjustment Factor

This measure displays the Overlapping Promotion Adjustment Factor value that was applied during the forecast generation at the item/location level.

The Overlapping Promotion Adjustment Factor specifies how the individually calculated promotion effects are combined with each other when the promotions are overlapped in the forecast horizon.

Blending Parameter Merged

This parameter displays the weight for combining the Final and Source Level promotion effects, when calculating the blended effect. The range of the parameter is 0 to 1. A value closer to 1 will yield a blended effect closer to the Source Final Level effect. A value closer to zero yields an effect closer to the Final Level effect. The measure displays the merged default and override values set in the Forecast Administration and Forecast Maintenance workbooks, respectively.

Source Level

This measure displays the pooling level assigned to the item/location. It is the merge between the default and override values set in the Forecast Administration and Forecast Maintenance workbooks, respectively.

Forecast Approval Workbook Navigation

The following sections demonstrate how the exception management framework can be used to navigate the Forecast Approval workbook.

The amount of forecasts that are typically produced every week by retailers is huge. Manually reviewing and approving so many numbers is not realistic. On the other hand, automatically approving all numbers is not good practice either. To address this fact, RPAS offers the exception management framework, where rules can be set up,

and the forecast can be checked against those rules. If the rules are violated, the cells are flagged, and the user can review and take action. If not, the forecasts are approved, and the user can concentrate on the forecasts that need attention.

RDF has a set of alerts configured, that can be used for the forecast approval process. Some of these rules, are implemented both as batch as well as workbook alerts. The batch alerts are calculated in batch and are typically used to narrow down the information that is brought in a workbook. They will include the whole forecast horizon of the product/locations that are alerted. Once the workbook is built, the user can navigate it by continuing to use the batch alert, or she can switch to workbook alerts. The workbook alerts are more convenient once the workbook is built, because they point the user to the exact cell that needs attention, not just the product/location.

Enable Alerts for Approval The first step is to enable some alerts to be used for approval. This happens in the Forecast Administration workbook, and should be an infrequent task.

Not all alerts make sense for all types of forecasts. For instance, for a baseline forecast, comparing forecasts vs. sales last year probably makes sense. However, due to differences in the promotional calendar year over year, it may not make sense to compare these quantities for the causal forecast. After selecting one or more alerts to be used in the approval process, choose your approval option.

In this case the user selected to compare the forecast vs the last approved forecast. The reasoning behind this is that the user has spent a fair amount of time to make sure the forecast is in good shape the last time forecast was generated. For the next run she wants to be alerted when the new forecasts differ significantly from the last approved values. If they are close, the new numbers are automatically approved.

At this point the setup is complete, and the forecast generation batch can be kicked off. Among other tasks, the batch generates the forecasts and runs the forecast approval rules.

When the user logs in after the batch is run, she will see the list of alerts and the number of time series that are flagged in each local domain.

She can pick an alert and select to build a new workbook. The forecast approval alerts are most relevant in building the Forecast Approval workbook, but they can be included in any workbook. The workbook builds using the same wizard options as if an alert was not included. The only difference is that the last step is to pick a sheet for the alert measure. A good choice is the System Parameters worksheet.

Valid Forecast Run Worksheet

In a global domain environment, the Valid Forecast Run worksheet allows you to identify which Local domains share the forecast birth date selected in the wizard when viewed in the Master domain. If in a Local domain, only the single position of the partition dimension is displayed.

Figure 4–6 Valid Forecast Run Worksheet

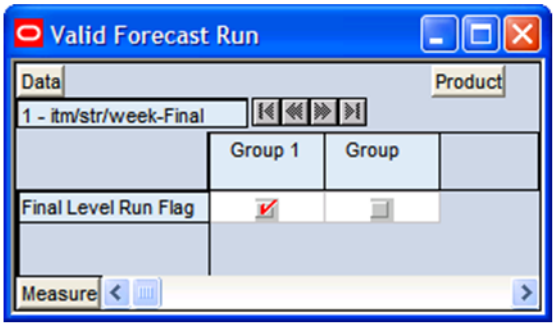

Measures: Valid Forecast Run Worksheet

The measure displayed on this worksheet is viewed at final-level/partition intersection. The Valid Forecast Run worksheet includes the following standard measure:

Final Level Run Flag

The Final Level Run Flag is a read-only Boolean measure. A check (set to True) displayed in this field indicates which positions from the partition dimension had a successful forecast run for the birth date selected in the Forecast Approval wizard.

Note: Edits to the Adjusted Forecast can only be committed for partitions that have an activated Final Level Run Flag.

Approving Forecasts Through Alerts (Exception Management)

RDF provides you with ability to manually approve every product/location forecast value. However, due to the extremely large volume of product/location forecast values that is generated for each forecast cycle, RPAS provides additional functionality that enhances your ability to evaluate and approve a forecast.

When configuring an RDF solution, specific Forecast Approval Alerts and alert parameters may be defined and applied during the batch process, and they may also be inserted into the Forecast Approval Workbook. Alerts report product/location forecast values that exceed expected thresholds. If a specific product/location forecast value exceeds a threshold, the alert measure is flagged as on and visible to you as a check mark in the appropriate product/location intersection.

The complexities of retail operations can result in the need to define many alerts; each designed to watch for a specific scenario that would require a user's evaluation of a forecast value.

For more information on how alert measures and rules are defined and registered, refer to the *Oracle Retail Predictive Application Server Configuration Tools User Guide* and *Oracle Retail Demand Forecasting Implementation Guide*.

Product Attributes and Dynamic Hierarchies

This section shows how dynamic attributes can be used to conveniently navigate the Forecast Approval workbook.

RDF already makes use of the exception management framework to help with the forecast review and approval process. The attributes can go one step further. Imagine an attribute has a 'replenish' attribute with values of yes or no. A user may want to make sure first that items that are replenished have robust and accurate forecasts. Reviewing forecasts of items that are not replenished may be a lower priority. Or imagine a retailer carries different brands, which are captured as item attributes. If a brand was newly introduced, the forecast analyst may want to first review the item forecasts of the newly introduced brand. Or if price tier is available as item attribute, the analyst may want to review high priced items first, to make sure margin is not lost.

Setup Product Tab

This tab allows you to review item attributes and setup dynamic hierarchies.

The available worksheets are:

- Review Product Attributes Worksheet
- Setup Product Attributes Worksheet

Review Product Attributes Worksheet

This worksheet displays the item attributes and the product attributes labels. This information is for review only, and can not be edited.

Setup Product Attributes Worksheet

This worksheet displays the item attributes and the product attributes labels. This information is for review only, and can not be edited.

Once the attributes are assigned, the Refresh Product Rollup custom menu needs to run for the changes to take effect.

The place where you can take advantage of the dynamic hierarchies is the Final Level Worksheet, where forecasts are reviewed, adjusted and approved. A typical layout – without making use of dynamic hierarchies.

Using Dynamic Hierarchies

If we want to use dynamic hierarchies, we need to enable which attribute we want to use. Remember, you can set up to three attributes.

Figure 4–7 Attribute Setup

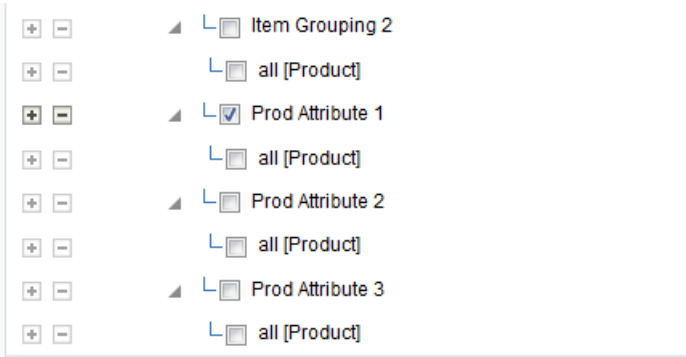

After setting up brand as the attribute, you can filter down to items that need attention first. Note how the private label and ***Columbian*** items are collapsed, and the ***House*** items are expanded.

5

Preprocessing Administration Workbook

This chapter describes how the preprocessing functionality is implemented in RDF using the Preprocess Admin Workbook.

The Preprocess Admin Workbook contains these tabs:

- Preprocess Admin Tab
- Preprocess Parameters Tab

> **Note:** Similar to the Preprocess Admin Workbook is the Preprocess Review and What-if Workbook. For functionality differences, see Preprocess Workbook Functionality.

About Preprocessing

Preprocessing is an optional module that is used to correct historical data prior to forecast generation when history does not represent general demand patterns. It is meant to automatically make adjustments to the raw POS (Point Of Sales) data so the next demand forecasts do not replicate undesired patterns.

Data Preprocessing is commonly used to:

- Correct for lost sales due to stock-outs
- Cleanse data for effects of promotions and short-term price changes (optional)
- Correct for outliers – unusually high or low values introduced by human error or special events (hurricane that left a store closed for a week)
- Scrub data manually to fake history and override user history
- Adjust demand for the occasional 53rd calendar week
- Manage demand created during events and holidays that do not occur in the same period every year, for example, Back to School.

Preprocessing runs after the data has been loaded from the host system and prior to forecast generation. Use the Preprocess Admin Workbook to produce and review data. It is common for an environment to require preprocessing to run multiple times to properly smooth the history.

Preprocessing Data in the RDF Workflow

Preprocessing offers a variety of Algorithm methods to support the business requirements. The main reason for preprocessing is to transform the raw sales data into a measure that gets as close as possible to unconstrained demand.

The preprocessing step is most often implemented in batch, by invoking a preprocessing special expression. The special expression takes several measures as input. For instance, one needs to specify the measure to be corrected, the desired algorithm, and the number of periods to be considered. However, you can go into more detail, and specify several filter window lengths, or exponential smoothing parameters.

Preprocess Administration Workbook, Tabs, and Worksheets

The following table lists the worksheets for the Preprocess Administration workbook.

Workbook	Tabs	Worksheets
Preprocess Admin Workbook	Preprocess Admin Tab	Preprocess Admin Worksheet
	Preprocess Parameters Tab	Preprocess Panel for Baseline Data Source Worksheet Preprocess Panel for Causal Data Source Worksheet Preprocess Method Parameters Worksheet Preprocess Method Parameters Override Worksheet

Preprocess Admin Workbook

Use the following procedure to create a Preprocess Administration workbook:

1. Within the Master or Local Domain, select **New** from the **File** menu.
2. Select the **Preprocess** tab to display a list of workbooks. Select Preprocess Administration.
3. Click **OK**. The Preprocess Administration wizard opens.
4. Select the locations to include in the workbook.
5. Click **Next**.
6. Select the products to include in the workbook.
7. Click **Finish** to display the workbook.

Worksheets for the Preprocess Administration Workbook

The following sections describe the components of the Preprocess Administration workbook.

The Preprocess Administration workbook has three tabs and includes the worksheets listed in Table 5–1.

Table 5–1 Preprocess Administration Workbook Worksheets

Tab	Description	Worksheets
Preprocess Admin Tab	Allows you to define the scope of the preprocessing run, as well as filter out item/locations where preprocessing does not make sense because of lack of enough historical sales.	Preprocess Admin Worksheet
Preprocess Parameters Tab	Makes available the preprocessing parameters for four rounds of preprocessing runs, necessary to calculate the data source for baseline forecasting and promotional forecasting.	Preprocess Panel for Baseline Data Source Worksheet Preprocess Panel for Causal Data Source Worksheet Preprocess Method Parameters Worksheet Preprocess Method Parameters Override Worksheet

Preprocess Admin Tab

This tab contains the Preprocess Admin worksheet.

Preprocess Admin Worksheet

The Preprocess Admin worksheet allows you to define the scope of the preprocessing run, as well as filter out item/locations where preprocessing does not make sense because of lack of enough historical sales.

Figure 5–1 Preprocess Admin Worksheet

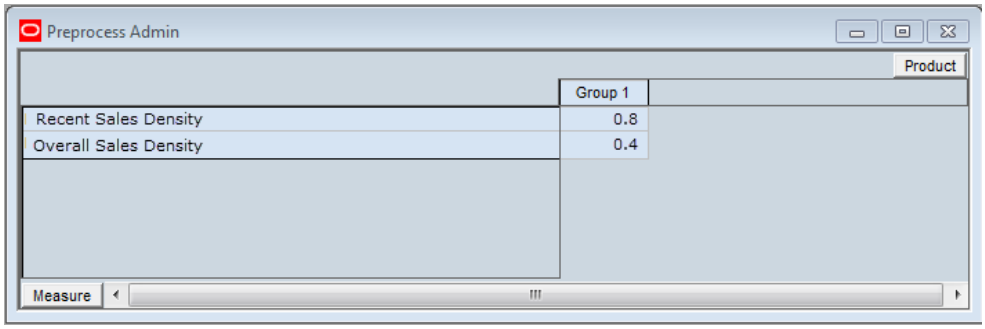

Measures: Preprocess Admin Worksheet

The Preprocess Admin worksheet contains the following editable measures:

> **Note:** The time series is preprocessed when both the Recent Sales Density and the Overall Sales Density results are larger than the thresholds for both measures.

Recent Sales Density

This measure shows a threshold value of sales density during a recent time period. The time period is defined by the Like TS Duration measure and the density is calculated as the number of populated cells during the Like TS Duration divided by the Like TS Duration. For example, if the Like TS Duration is five periods, and there are two periods with non-zero sales, then the density is 2/5, which is 40%.

Overall Sales Density

This measure shows a threshold value of sales density over the entire sales history (the time period). The time period is defined by the first non-zero data point until today. The density is calculated as the number of populated cells during the time period divided by the time period. For example, if the time period is 100 periods, and there are 20 periods with non-zero sales, then the density is 20/100, which is 20%.

Preprocess Parameters Tab

This tab has four worksheets:

- Preprocess Panel for Baseline Data Source Worksheet
- Preprocess Panel for Causal Data Source Worksheet
- Preprocess Method Parameters Worksheet
- Preprocess Method Parameters Override Worksheet

These worksheets make available the preprocessing parameters for four rounds of preprocessing runs, necessary to calculate the data source for baseline forecasting and promotional forecasting.

Preprocess Panel for Baseline Data Source Worksheet

The worksheet displays the measures necessary to create the data source for baseline forecasting. This involves four rounds of preprocessing that run in batch or online.in this order:

1. Correcting for stockouts
2. Correcting for outliers
3. Depromoting sales
4. Smooth sales

Figure 5–2 Preprocess Panel for Baseline Data Source Worksheet

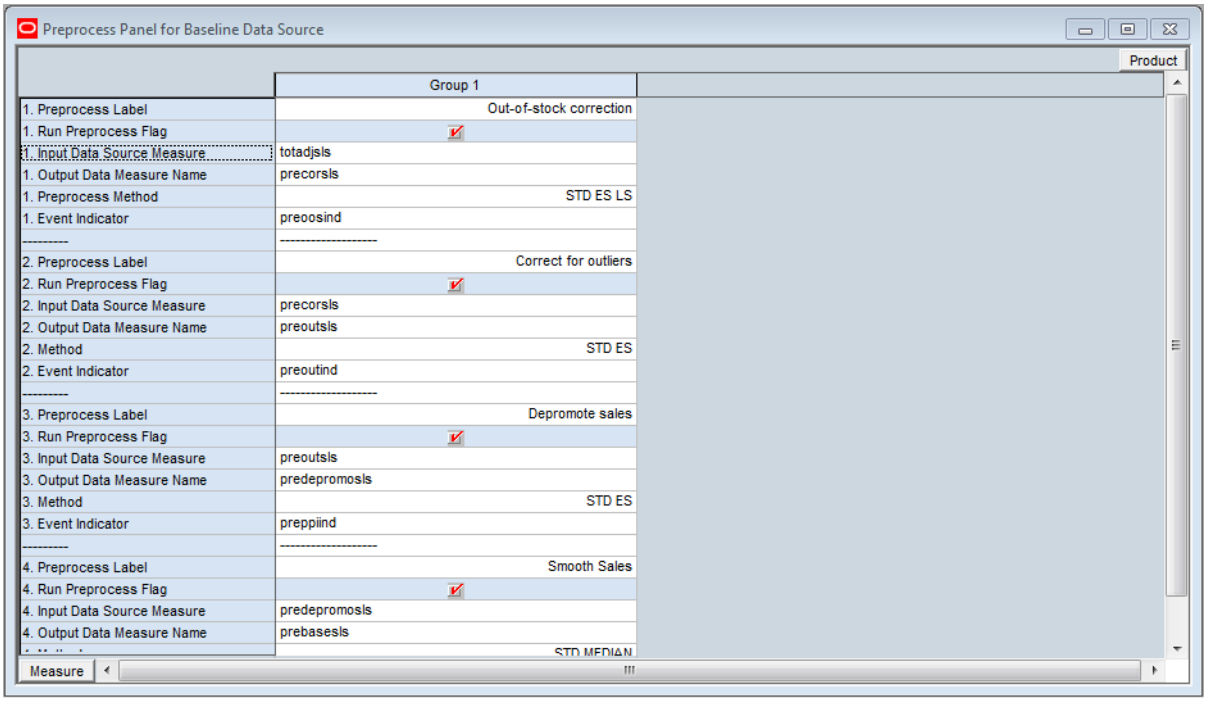

Measures: Preprocessing Panel for Baseline Data Source Worksheet

> **Note:** Measures are replicated for each round of preprocessing.

The Preprocessing Panel for Baseline Data Source worksheet contains the following measures:

Event Indicator

Indicates the measure name that represents the event flag.

Input Data Source Measure

Indicates the measures that will be corrected.

Output Data Measure

Indicates the measure that stores the result of the preprocessing run.

Preprocess Method

Name of the preprocessing method to be used. This method is fixed, but this can be expanded to a list.

Preprocessing Label

A label denoting the purpose of the preprocessing run, for example, Out-of-Stock Correction.

Run Preprocess Flag

Boolean measure indicating if this rule should be run or skipped.

Preprocess Panel for Causal Data Source Worksheet

This worksheet displays the measures necessary to create the data source for baseline forecasting. This involves three rounds of preprocessing that run in this order:

1. Correcting for stockouts
2. Correcting for outliers
3. Deseasonalizing the measure to create the causal data source

The first two runs are the same used to generate the baseline data source. Hence they use the same parameters as the ones in the Preprocess Panel for Baseline Data Source Worksheet. The third run deseasonalizes the demand.

Figure 5–3 Preprocess Panel for Causal Data Source Worksheet

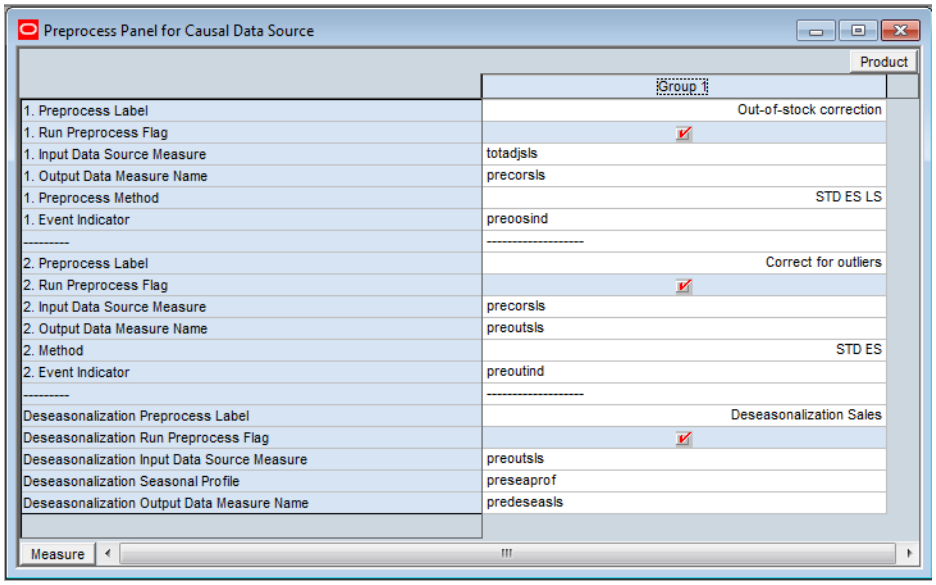

Measures: Preprocessing Panel for Causal Data Source Worksheet

> **Note:** Measures are replicated for each round of preprocessing.

The Preprocessing Panel for Causal Data Source worksheet contains the following measures:

Event Indicator

Indicates the measure name that represents the event flag.

Input Data Source Measure

Indicates the measures that will be corrected.

Output Data Measure

Indicates the measure that stores the result of the preprocessing run.

Preprocess Method

Name of the preprocessing method to be used. This method is fixed, but this can be expanded to a list.

Preprocessing Label

A label denoting the purpose of the preprocessing run, for example, Out-of-Stock Correction.

Run Preprocess Flag

Boolean measure indicating if this rule should be run or skipped.

Preprocess Method Parameters Worksheet

In this worksheet, you can enter values for parameters specific for some of the preprocessing methods available in the special expression. There are four sets of parameters, corresponding to the four runs of the preprocessing special expression. The parameters are entered at the class/store intersection

Figure 5–4 Preprocess Method Parameters Worksheet

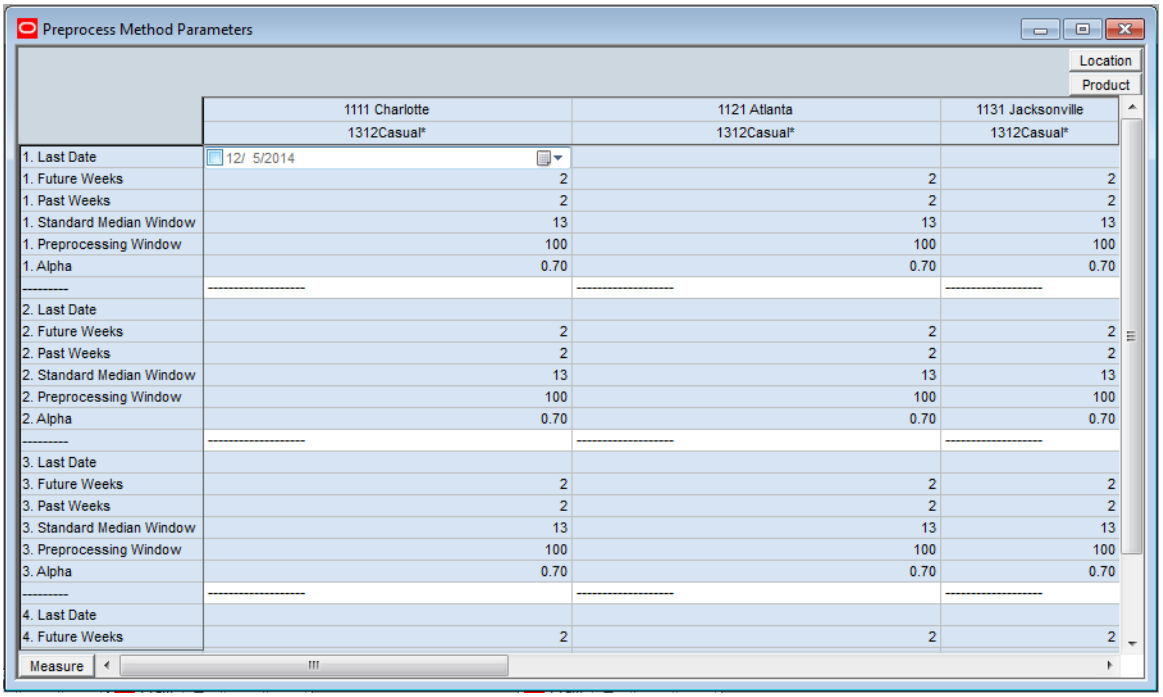

Measures: Preprocessing Methods Parameters Worksheet

The Preprocessing Methods Parameters worksheet contains the following measures:

Alpha

Exponential smoothing coefficient used to calculate past and future velocities.

Future Weeks

This represents the maximum number of data points to calculate the future velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Last Date

This represents the end date of the preprocessing window; it is typically today's date, but can be any date in the past.

Past Weeks

This represents the maximum number of data points to calculate the past velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Preprocessing Window

Number of historical data points that are preprocessed.

Standard Median Window

Filter window length for the Standard Median preprocessing method.

Preprocess Method Parameters Override Worksheet

In this worksheet, you can override values for parameters specific for some of the preprocessing methods available in the special expression. There are four sets of parameters, corresponding to the four runs of the preprocessing special expression.

After all parameters are set and committed back to the domain, usually a batch job will run the pre-processing tabs and prepare the source data for forecast generation.

Figure 5–5 Preprocess Method Parameters Override Worksheet

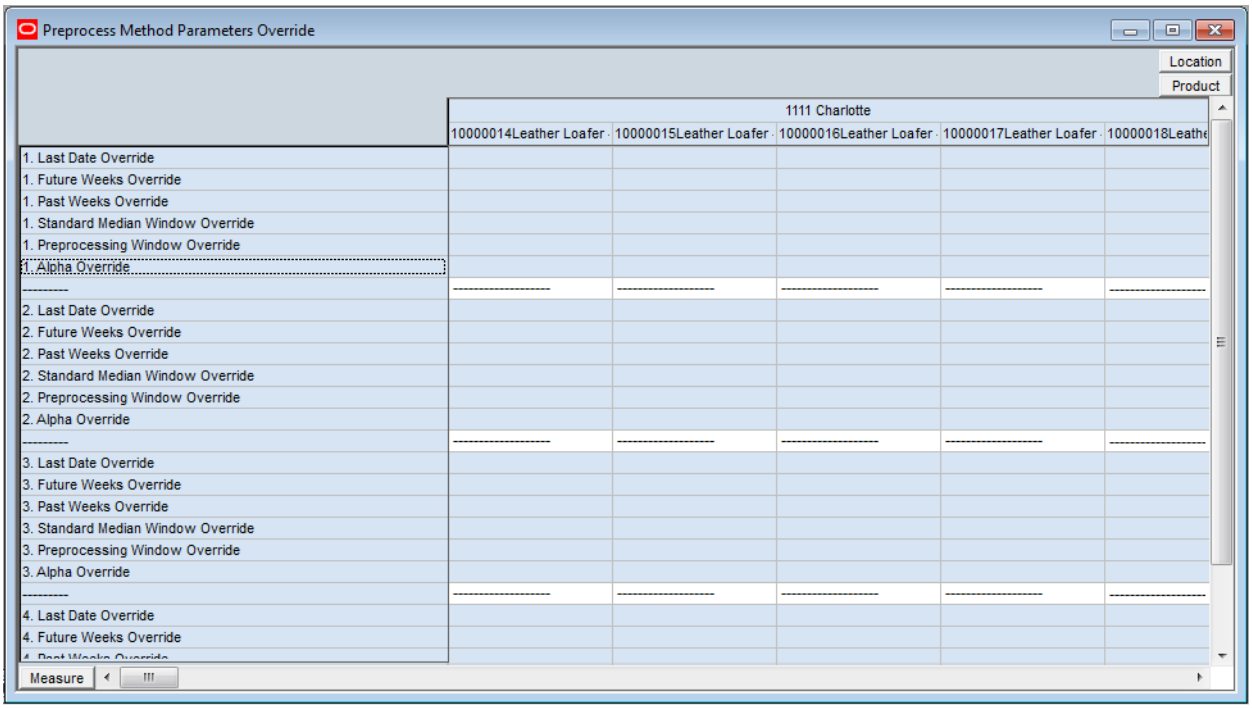

Measures: Preprocessing Method Parameters Override Worksheet

The Preprocessing Method Parameters Override worksheet contains the following measures:

Alpha Override

Exponential smoothing coefficient used to calculate past and future velocities.

Future Weeks Override

This represents the maximum number of data points to calculate the future velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Last Date Override

This represents the end date of the preprocessing window; it is typically today's date, but can be any date in the past.

Past Weeks Override

This represents the maximum number of data points to calculate the past velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Preprocessing Window Override

Number of historical data points that are preprocessed.

Standard Median Window Override

Filter window length for the Standard Median preprocessing method.

6

Source Measure Maintenance Workbook

This chapter describes how the preprocessing functionality is implemented in RDF using the Source Measure Maintenance Workbook.

The functionality in the Source Measure Maintenance Workbook is a superset of the functionality in the Preprocessing Administration Workbook. The purpose and functionality between the two is described in Source Measure Maintenance Functionality.

Note: For information about preprocessing and RDF, see Preprocessing Data in the RDF Workflow

---

Source Measure Maintenance Functionality

The Preprocess Admin Workbook and the Source Measure Maintenance Workbook have a large set of common content.

The main difference is that while the Source Measure Maintenance Workbook has the calendar hierarchy, on top of the product and location, and the Preprocess Admin Workbook has only the product and location.

The additional hierarchy allows the review of the time-phased preprocessing measures, as well as the calculated forecasting data sources.runs.

Due to their additional dimension of *week*, these measures add to the size of the workbook, and also make workbook operations slower. For instance, workbook build, refresh, commit, and so on, take longer than in the otherwise similar Preprocess Admin Workbook.

Preprocess Administration Workbook

The Preprocess Admin Workbook, described in Chapter 5 is at the product/location intersection, so it can be built with a lot of positions, without experiencing poor performance. The purpose is to set preprocessing parameters, which are inputs to the special expression that is run in batch.

Source Measure Maintenance Workbook

The Source Measure Maintenance Workbook, described in this chapter, is at the product/location/calendar intersection, and is a lot more data intensive. The purpose is to set preprocessing parameters and run the data filtering online, with the ability to review the results without having to wait for an overnight batch. If the results are not as expected, or you want to experiment with different settings, you can make changes to the parameters and rerun the custom menus. To achieve this it is expected that only a small subset of the available product/locations is included in the workbook.

Source Measure Maintenance Workbooks, Steps, and Views

The following table lists the workbooks, steps, and views for the Preprocess Review and What-if task.

Workbook	Tabs	Worksheets
Source Measure Maintenance Workbook	Preprocess Admin Tab	Preprocess Admin Worksheet
	Preprocess Parameters Tab	Preprocess Panel for Baseline Data Source Worksheet Preprocess Method Parameters (Baseline or Causal) Worksheet Preprocess Method Parameters Override Worksheet
	Review Tab	Source Maintenance Worksheet Source Maintenance Plot Worksheet
	Seasonal Profile for Causal Tab	Seasonal Profile for Causal Worksheet

Source Measure Maintenance Workbook

Use the following procedure to create a Source Measure Maintenance workbook:

1. Within the Master or Local Domain, select **New** from the **File** menu.
2. Select the **Preprocess** tab to display a list of workbooks. Select Source Measure Maintenance.
3. Click **OK**. The Source Measure Maintenance wizard opens.
4. Select the products to include in the workbook and click **Next**.

   Note: It is important to include all products that are members of the Merchandise dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at subclass/store/week, you must include all items that are members of the particular subclass to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the Merchandise hierarchy is employed if the task supports an AutoTask build.

5. Select the locations to include in the workbook and click **Next**.

   Note: It is important to include all locations that are members of the location dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at item/chain/week, you should include all locations that are members of the particular chain to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the location hierarchy is employed if the task supports an AutoTask build.

6. Select the weeks of the forecast to include in the workbook.
7. Click **Finish** to display the workbook.

Worksheets for the Source Measure Maintenance Workbook

The following sections describe the tabs and worksheets of the Source Measure Maintenance workbook.

Preprocess Admin Tab

This tab contains the Preprocess Admin Worksheet

Preprocess Admin Worksheet

The Preprocess Admin worksheet displays the scope of the preprocessing run.

Figure 6–1 Preprocess Admin Worksheet

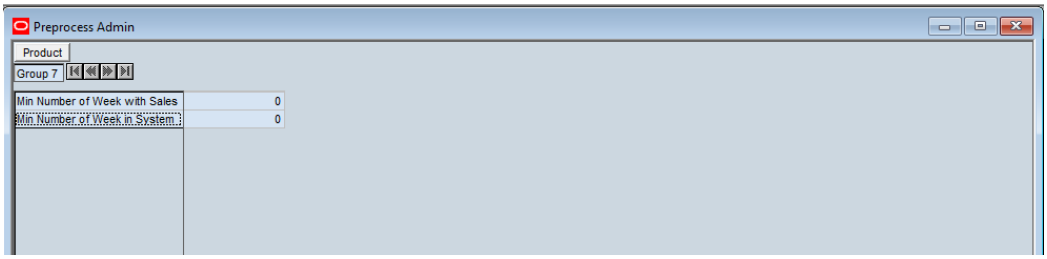

Measures: Preprocess Admin Worksheet

The Preprocess Admin worksheet contains the following editable measures:

Min Number of Weeks in System

This parameter defines the number of periods from when an item was introduced in the system. Usually the introduction time is considered to be the date when the item first sold. This check is also introduced to stop making data corrections for items that are very new, and where cleansing would be unreliable.

Min Number of Weeks with Sales

This parameter defines the number of weeks with sales that an item/store combination needs to have to qualify for data cleansing. The reasoning behind this check is that for items without enough data, corrections may not be reliable. Once there is enough data, and trends become clearer, corrections can be made.

Preprocess Parameters Tab

This tab contains these worksheets:

- Preprocess Panel for Baseline Data Source Worksheet
- Preprocess Method Parameters (Baseline or Causal) Worksheet
- Preprocess Method Parameters Override Worksheet

Preprocess Panel for Baseline Data Source Worksheet

Depending on your wizard selection for either baseline or causal, the worksheet shows the preprocessing parameters for each relevant run. For example if baseline is selected, the worksheet displays preprocessing information for the four runs that are configured. If causal is selected, the worksheet shows preprocessing parameters for three runs.

Baseline Worksheet

The baseline worksheet displays the measures necessary to create the data source for baseline forecasting. This involves four rounds of preprocessing that run in batch or online in this order:

1. Correcting for stockouts
2. Correcting for outliers
3. Depromoting sales
4. Smooth sales

Figure 6–2 Preprocess Panel for Baseline Worksheet

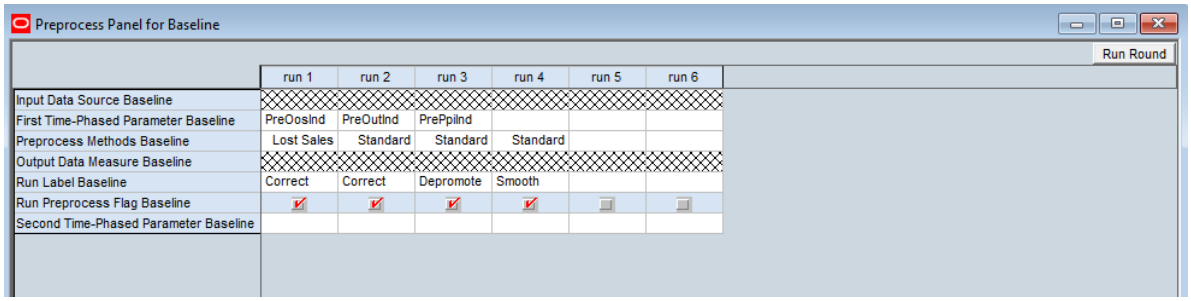

Causal Worksheet

The causal worksheet displays the measures necessary to create the data source for causal forecasting. This involves three rounds of preprocessing that are run in batch:

1. Correcting for stockouts
2. Correcting for outliers
3. Deseasonalizing the measure to create the causal data source

Figure 6–3 Preprocess Panel for Causal Data Source Worksheet

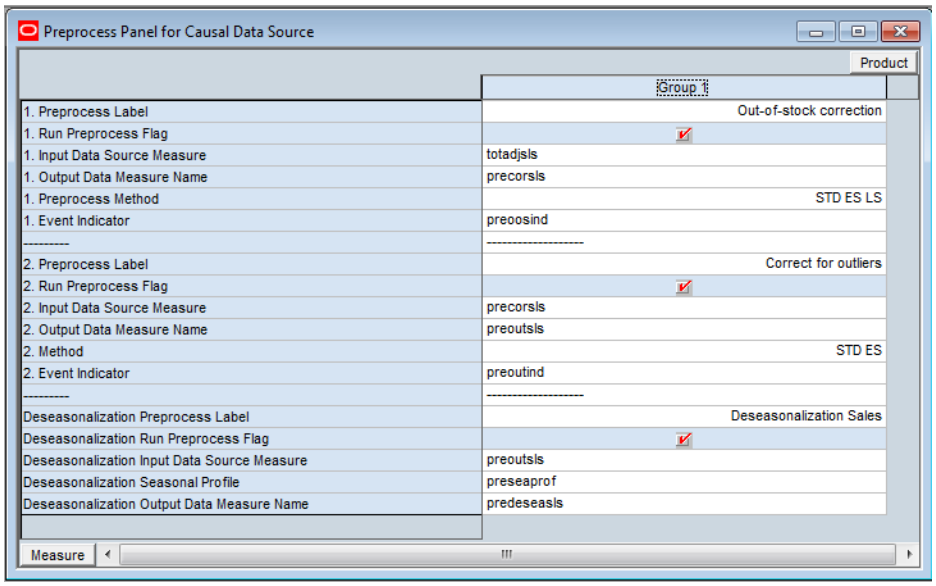

Measures: Preprocessing Panel (for Baseline or Causal) Worksheet

Note: Measures are replicated for each round of preprocessing.

The Preprocessing Panel (for Baseline or Causal) worksheet contains the following measures:

Input Data Source

Indicates the measure that will be corrected. This is the input to the first preprocessing run. There are no inputs available for other runs other than the first run.

First Time-Phased Parameter Baseline or Causal

This measure stores the first time-phased measure that is required for some preprocessing methods. For instance, for the STD ES LS method, this measure would store the measure name of the outage flag. Or for the STD ES method, it could store the name of the outlier flag.

Preprocess Method

Name of the preprocessing method to be used for each run. This method is selected in the Configuration Tools.

Output Data Measure

Indicates the measure that stores the result of the last configured preprocessing run. For instance, for the Preprocess Panel for Causal, the output comes from run 3.

Run Label

A label denoting the purpose of the preprocessing run, for example, Correct Outliers, or Smooth Sales.

Run Preprocess Flag

Boolean measure indicating if this run should be enabled or skipped.

Second Time-Phased Parameter Causal

This measure stores the second time-phased measure that is required for some preprocessing methods. For instance, for the Forecast Sigma method, this measure would store measure name of the confidence intervals. Or for the Override method it could store the measure name of the outage flag.

Preprocess Simulation Baseline or Causal

Depending on the level selected, the custom menu is called either Preprocess Baseline or Preprocess Causal and performs the following actions:

- Correcting for stockouts
- Correcting for outliers
- Depromoting sales (only applies to Preprocess Baseline)
- Deseasonalize Sales (only applies to Preprocess Causal)
- Smooth sales (only applies to Preprocess Baseline)

The ultimate goal for Preprocess Simulation is to create the source data for generating forecasts:

- Baseline Sales: source data for baseline forecast generation
- Deseasonalized Sales: source data for causal forecast generation

Preprocess Method Parameters (Baseline or Causal) Worksheet

In this worksheet, you can enter values for parameters specific for some of the preprocessing methods available in the special expression. There are four sets of parameters, corresponding to the four runs of the preprocessing special expression. The parameters are entered at the class/store intersection

Figure 6–4 Preprocess Method Parameters Worksheet

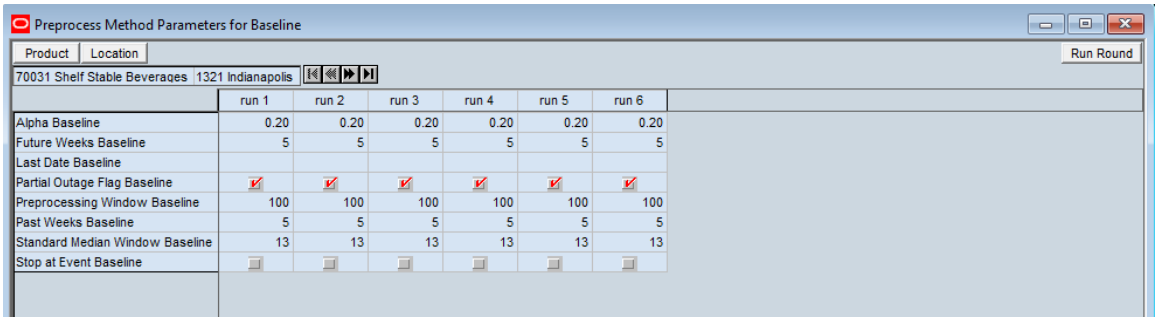

Measures: Preprocessing Methods Parameters (Baseline or Causal) Worksheet

The Preprocessing Methods Parameters worksheet contains the following measures:

Alpha

Exponential smoothing coefficient used to calculate past and future velocities.

Future Weeks

This represents the maximum number of data points to calculate the future velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Last Date

This represents the end date of the preprocessing window; it is typically today's date, but can be any date in the past.

Past Weeks

This represents the maximum number of data points to calculate the past velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Preprocessing Window

Number of historical data points that are preprocessed.

Standard Median Window

Filter window length for the Standard Median preprocessing method.

Preprocess Method Parameters Override Worksheet

In this worksheet, you can override values for parameters specific for some of the preprocessing methods available in the special expression. There are four sets of parameters, corresponding to the four runs of the preprocessing special expression.

After all parameters are set and committed back to the domain, usually a batch job will run the pre-processing tabs and prepare the source data for forecast generation.

Figure 6–5 Preprocess Method Parameters Override Worksheet

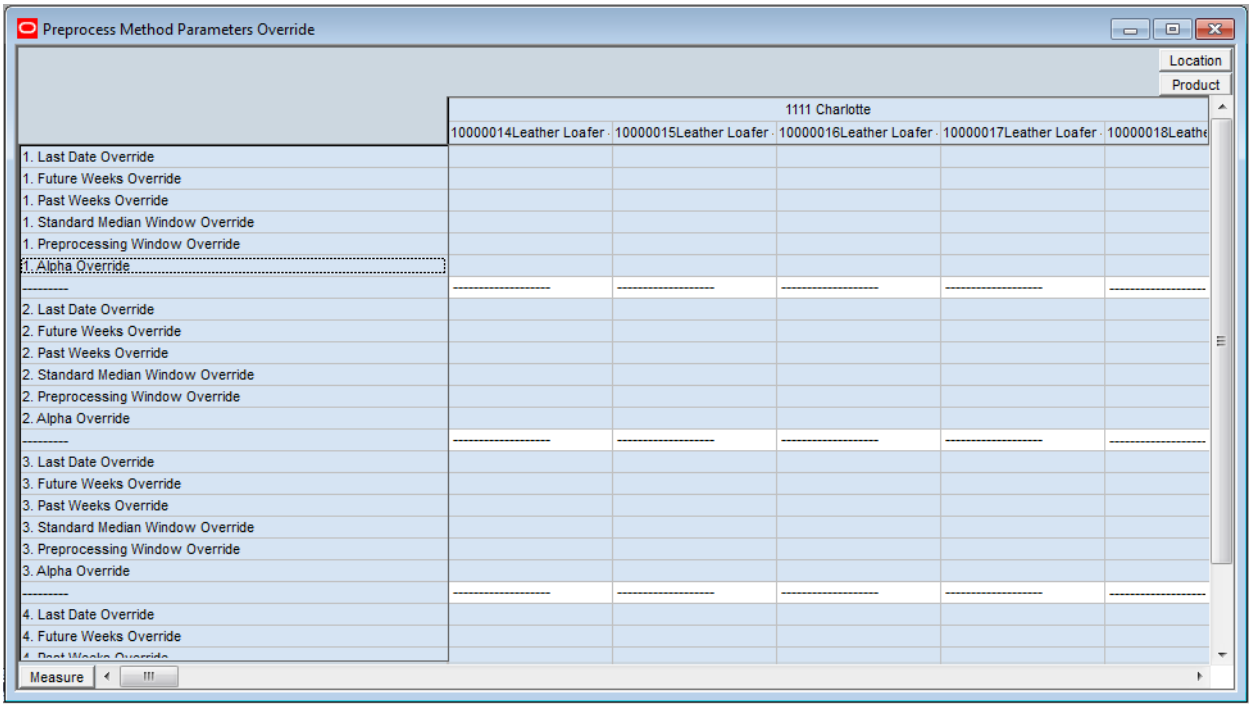

Measures: Preprocessing Method Parameters Override Worksheet

The Preprocessing Method Parameters Override worksheet contains the following measures:

Alpha Override

Exponential smoothing coefficient used to calculate past and future velocities.

Future Weeks Override

This represents the maximum number of data points to calculate the future velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Last Date Override

This represents the end date of the preprocessing window; it is typically today's date, but can be any date in the past.

Past Weeks Override

This represents the maximum number of data points to calculate the past velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Preprocessing Window Override

Number of historical data points that are preprocessed.

Standard Median Window Override

Filter window length for the Standard Median preprocessing method.

Review Tab

This tab contains these worksheets:

- Source Maintenance Worksheet
- Source Maintenance Plot Worksheet

The main purpose of this step is to display time-phased measures that represent input and output to the preprocessing stages, run in batch based on the settings selected in the Preprocessing Admin tasks.

Source Maintenance Worksheet

This worksheet displays measures that represent input and output of the preprocessing runs, in table format.

Figure 6–6 Source Maintenance Worksheet

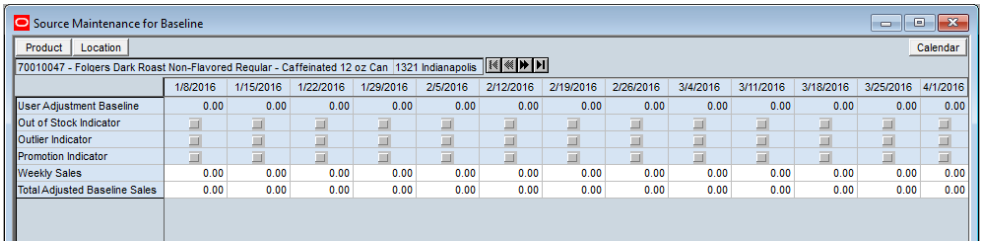

Measures: Source Maintenance Worksheet

The Source Maintenance worksheet contains the following measures:

User Adjustment

In this measure, you can enter values that are going to be added to the preprocessing adjustments to create the data sources.

The logic is: data source = weekly sales + preprocessing adjustments + user adjustment

This measure is read/write.

Weekly Sales

This measure stores the raw sales loaded in RDF. This is the input to the first run of preprocessing. This measure is read only.

Data Source

This measure represents the output of the preprocessed raw sales, as well as incorporates the user adjustments according to the formula:

data source = weekly sales + preprocessing adjustments + user adjustments.

Out of Stock Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects sales for lost sales.

Outliers Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects the sales for outliers.

Promotion Indicator

This measure is usually calculated as the or of all available Boolean promotional variables. It is used during the preprocessing run that removes promotional sales.

Source Maintenance Plot Worksheet

This worksheet displays measures that represent input and output of the preprocessing runs, in table format. The measures are the same as in the Source Maintenance Worksheet, but are displayed in chart form.

Figure 6–7 Source Maintenance Plot Worksheet

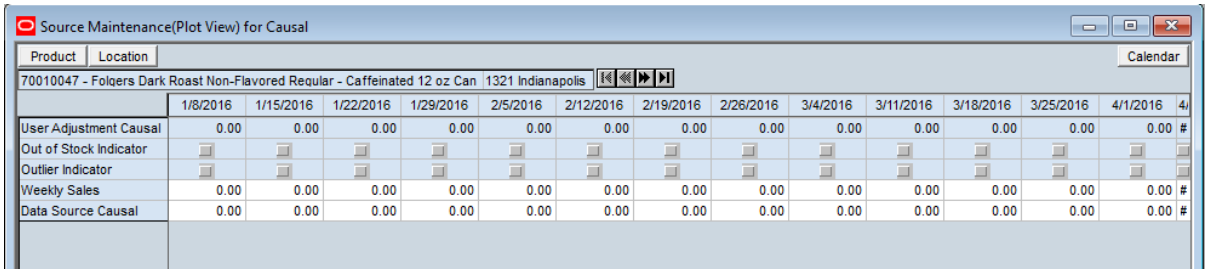

Measures: Source Maintenance Plot Worksheet

The Source Maintenance Plot worksheet contains the following measures:

User Adjustment

In this measure, you can enter values that are going to be added to the preprocessing adjustments to create the data sources.

The logic is: data source = weekly sales + preprocessing adjustments + user adjustment

This measure is read/write.

Weekly Sales

This measure stores the raw sales loaded in RDF. This is the input to the first run of preprocessing. This measure is read only.

Data Source

This measure represents the output of the preprocessed raw sales, as well as incorporates the user adjustments according to the formula:

data source = weekly sales + preprocessing adjustments + user adjustments.

Out of Stock Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects sales for lost sales.

Outliers Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects the sales for outliers.

Promotion Indicator

This measure is usually calculated as the or of all available Boolean promotional variables. It is used during the preprocessing run that removes promotional sales.

Seasonal Profile for Causal Tab

This tab contains this worksheet Seasonal Profile for Causal Worksheet

Seasonal Profile for Causal Worksheet

The purpose of this step is to display and allow for editing a week-of-year profile, that can be used during the preprocessing run to deseasonalize the demand that is used as data source for the causal run.

Note: This worksheet is only available if the workbook is built for the causal level.

The worksheet displays the seasonal profile that can be used to deseasonalize the data source for the causal run. One of the options is to load the profile, and the user can review and adjust it in this worksheet. If the profile is not available the user can manually key in the values. Since a profile is usually higher than the forecast intersection, that is, item/store, this worksheet is at subclass/region. However, the user can roll up the worksheet to an even higher intersection, for instance, department/chain, and there enter the values, which are evenly replicated to subclass/region by clicking **Calculate**.

Figure 6–8 Seasonal Profile for Causal Worksheet

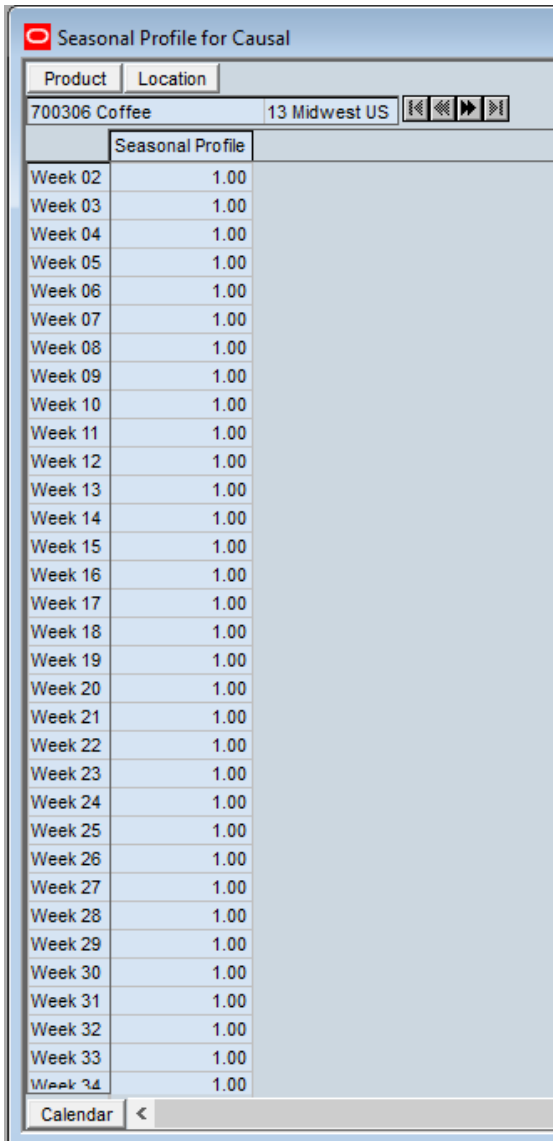

7

Preprocessing Review and What-if Workbook

This chapter describes how the preprocessing functionality is implemented in RDF using the Preprocess Review and What-if Workbook.

The functionality in the Preprocessing Review and What-if Workbook is a superset of the functionality in the Preprocess Admin Workbook. The purpose and functionality between the two is described inPreprocess Workbook Functionality.

Use the Preprocess Review and What-if Workbook to perform the tabs for this workbook:

- Preprocess Admin Tab
- Out of Stock and Outlier Threshold Tab
- Preprocess Parameters Tab
- What-if Tab

> **Note:** For information about preprocessing and RDF, see Preprocessing Data in the RDF Workflow

Preprocess Workbook Functionality

The main difference in the Preprocess Review and What-if Workbook versus the Preprocess Admin Workbook lies in the ability to run custom menus to create out-of-stock and outlier flags, as well as run the preprocessing special expression to create the two data sources: for the causal forecasting run, as well as for forecasting baseline demand. To do this, time-phased measures are necessary, which are inputs and outputs of the preprocessing runs.

Due to their additional dimension of *week*, these measures add to the size of the workbook, and also make workbook operations slower. For instance, workbook build, refresh, commit, and so on, take longer than in the otherwise similar Preprocess Admin Workbook.

Preprocessing Administration Workbook

The Preprocess Admin Workbook is at the product/location intersection, so it can be built with a lot of positions, without experiencing poor performance. The purpose is to set preprocessing parameters, which are inputs to the special expression that is run in batch.

Preprocess Review and What-if Workbook

The Preprocess Review and What-if Workbook is at the product/location/calendar intersection, and is a lot more data intensive. The purpose is to set preprocessing parameters and run the data filtering online, with the ability to Review the results without having to wait for an overnight batch. If the results are not as expected, or you want to experiment with different settings, you can make changes to the parameters and rerun the custom menus. To achieve this it is expected that only a small subset of the available product/locations is included in the workbook.

Preprocess Review and What-if Workbooks, Tabs, and Worksheets

The following table lists the workbooks, tabs, and worksheets for the Preprocess Review and What-if workbook.

Workbook	Tabs	Worksheets
Preprocess Review and What-if Workbook	Preprocess Admin Tab	Preprocessing Admin Worksheet
	Out of Stock and Outlier Threshold Tab	Out of Stock and Outlier Threshold Worksheet
	Preprocess Parameters Tab	Preprocess Panel for Baseline Data Source Worksheet Preprocess Panel for Causal Data Source Worksheet Preprocess Method Parameters Worksheet Preprocess Method Parameters Override Worksheet
	What-if Tab	What-if Data Worksheet What-if Plot Worksheet

Preprocess Review and What-if Workbook

Use the following procedure to create a Preprocess Review And What-if workbook:

1. Within the Master or Local Domain, select **New** from the **File** menu.
2. Select the **Preprocess** tab to display a list of workbooks. Select Preprocess Review And What-if.
3. Click **OK**. The Preprocess Review And What-if wizard opens.
4. Select the locations to include in the workbook.
5. Click **Next**.
6. Select the products to include in the workbook.
7. Click **Next**.
8. Select the weeks to include in the workbook.
9. Click **Finish** to display the workbook.

Preprocess Admin Tab

This tab contains the Preprocess Admin worksheet that allows you to define the scope of the preprocessing run, as well as filter out item/locations where preprocessing does not make sense because of lack of enough historical sales.

Preprocessing Admin Worksheet

The Preprocessing Admin worksheet displays the scope of the preprocessing run.

Figure 7–1 Preprocessing Admin Worksheet

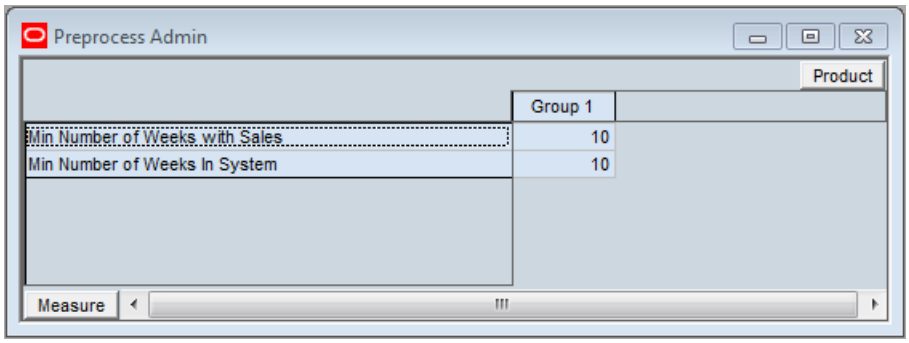

Measures: Preprocessing Admin Worksheet

The Preprocessing Admin worksheet contains the following measures:

Deseasonalized Flag

This parameter defines the scope of the preprocessing runs. The options are:

- Preprocess demand to create the data source necessary to forecast baseline demand
- Preprocess demand to create the data source for promotional forecasting
- Generate both data source measures

Min Number of Weeks in System

This parameter defines the number of periods from when an item was introduced in the system. Usually the introduction time is considered to be the date when the item first sold. This check is also introduced to stop making data corrections for items that are very new, and where cleansing would be unreliable.

Min Number of Weeks with Sales

This parameter defines the number of weeks with sales that an item/store combination needs to have to qualify for data cleansing. The reasoning behind this check is that for items without enough data, corrections may not be reliable. Once there is enough data, and trends become clearer, corrections can be made.

Out of Stock and Outlier Threshold Tab

This tab contains a worksheet that allows you to manage parameters used to create out of stock and outlier flags. Ideally, a retailer knows when stockouts occurred, or when special events created outliers in demand. These flags are usually interfaced in RDF and used in the preprocessing special expression. Many times, though, outages and outliers are not tracked, or they are not very exact. However, it is always beneficial to correct for those unwanted cases, to get closer to unconstrained demand. To make sure corrections are made, RDF has rules configured to create the flags, based on demand, and the parameters defined in this worksheet.

Out of Stock and Outlier Threshold Worksheet

The Out of Stock and Outlier Threshold worksheet displays stock and outlier flags.

Figure 7–2 Out of Stock and Outlier Threshold Worksheet

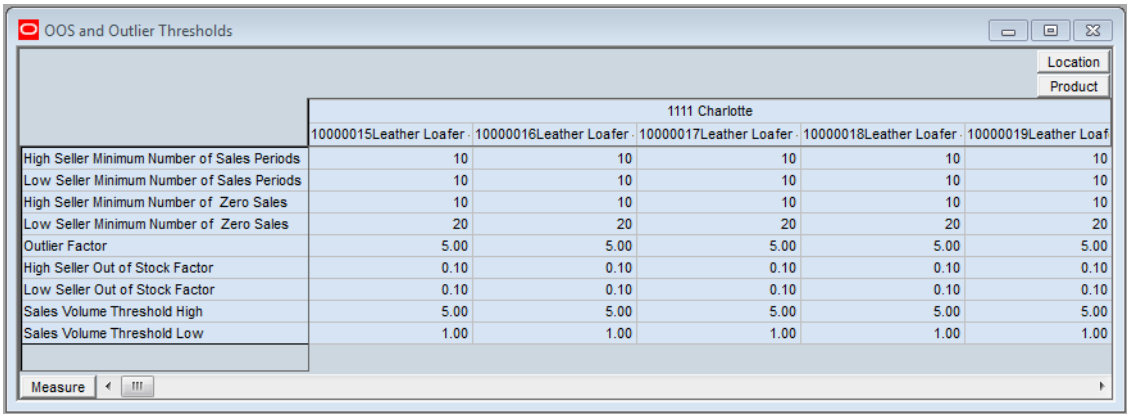

Measures: Out of Stock and Outlier Threshold Worksheet

The Out of Stock and Outlier Threshold worksheet contains the following measures:

High Seller Out of Stock Factor

This measure defines the threshold for when an high selling item is considered out of stock. For in-stance if the sales of an item are less than the rate of sales of the item multiplied with this value, this can be an indication that there was a stockout.

High Sellers Minimum Number of Sales Periods

This a condition on the number of periods a high seller has been selling. For instance if the number of periods the item has been selling is less than the value, the history is deemed to be too short, and no out of stock calculation will occur for the item.

High Sellers Minimum Number of Zero Sales

This a condition on the number of periods when a high selling item had zero sales. For instance if the number of periods with zero sales is larger than this value, this can be an indication that there was a stockout

Low Seller Out of Stock Factor

This measure defines the threshold for when a low selling item is considered out of stock. For instance if the sales of an item are less than the rate of sales of the item multiplied with this value, this can be an indication that there was a stockout.

Low Sellers Minimum Number of Sales Periods

This a condition on the number of periods a low seller has been selling. For instance if the number of periods the item has been selling is less than the value, the history is deemed to be too short, and no out of stock calculation will occur for the item.

Low Sellers Minimum Number of Zero Sales

This a condition on the number of periods when a low selling item had zero sales. For instance if the number of periods with zero sales is larger than this value, this can be an indication that there was a stockout

Outlier Factor

This measure defines the threshold for when the sales of an item are considered outliers. For instance if the sales of an item are more than the rate of sales of the item multiplied with this value, this can be an indication that there was an outlier.

Sales Volume Threshold High

If the rate of sale of an item is larger than value stored in this measure, it is considered a high selling item; this measure is part of the calculation when the out of stock indicator needs to be created using rules.

Sales Volume Threshold Low

If the rate of sale of an item is less than value stored in this measure, it is considered a low selling item; this measure is part of the calculation when the out of stock indicator needs to be created using rules. The rules that populate the Out of Stock and Outlier flags can be found in Appendix A, "Appendix: Preprocessing."

Preprocess Parameters Tab

This tab has four worksheets:

- Preprocess Panel for Baseline Data Source Worksheet
- Preprocess Panel for Causal Data Source Worksheet
- Preprocess Method Parameters Worksheet
- Preprocess Method Parameters Override Worksheet

These worksheets make available the preprocessing parameters for four rounds of preprocessing runs, necessary to calculate the data source for baseline forecasting and promotional forecasting.

Preprocess Panel for Baseline Data Source Worksheet

The worksheet displays the measures necessary to create the data source for baseline forecasting. This involves four rounds of preprocessing that run in batch or online.in this order:

1. Correcting for stockouts
2. Correcting for outliers
3. Depromoting sales
4. Smooth sales

Figure 7–3 Preprocess Panel for Baseline Data Source Worksheet

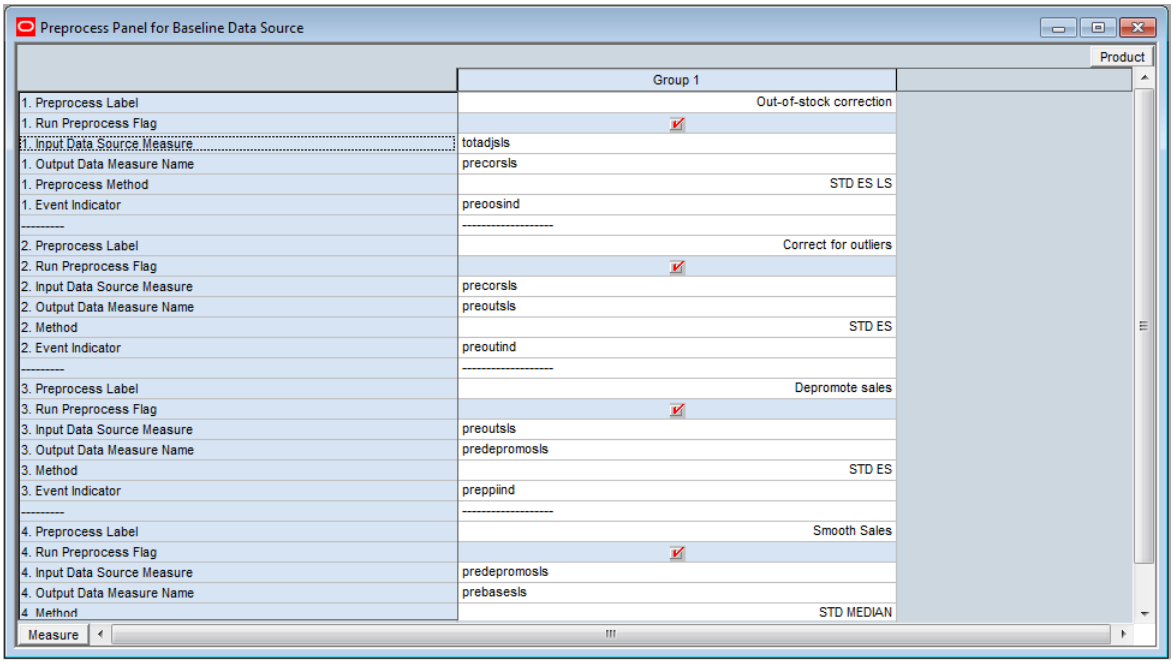

Measures: Preprocessing Panel for Baseline Data Source Worksheet

Note: Measures are replicated for each round of preprocessing.

The Preprocessing Panel for Baseline Data Source worksheet contains the following measures:

Event Indicator

Indicates the measure name that represents the event flag.

Input Data Source Measure

Indicates the measures that will be corrected.

Output Data Measure

Indicates the measure that stores the result of the preprocessing run.

Preprocess Method

Name of the preprocessing method to be used. This method is fixed, but this can be expanded to a list.

Preprocessing Label

A label denoting the purpose of the preprocessing run, for example, Out-of-Stock Correction.

Run Preprocess Flag

Boolean measure indicating if this rule should be run or skipped.

Preprocess Panel for Causal Data Source Worksheet

This worksheet displays the measures necessary to create the data source for baseline forecasting. This involves three rounds of preprocessing that run in this order:

1. Correcting for stockouts
2. Correcting for outliers
3. Deseasonalizing the measure to create the causal data source

The first two runs are the same used to generate the baseline data source. Hence they use the same parameters as the ones in the Preprocess Panel for Baseline Data Source Worksheet. The third run deseasonalizes the demand.

Figure 7–4 Preprocess Panel for Causal Data Source Worksheet

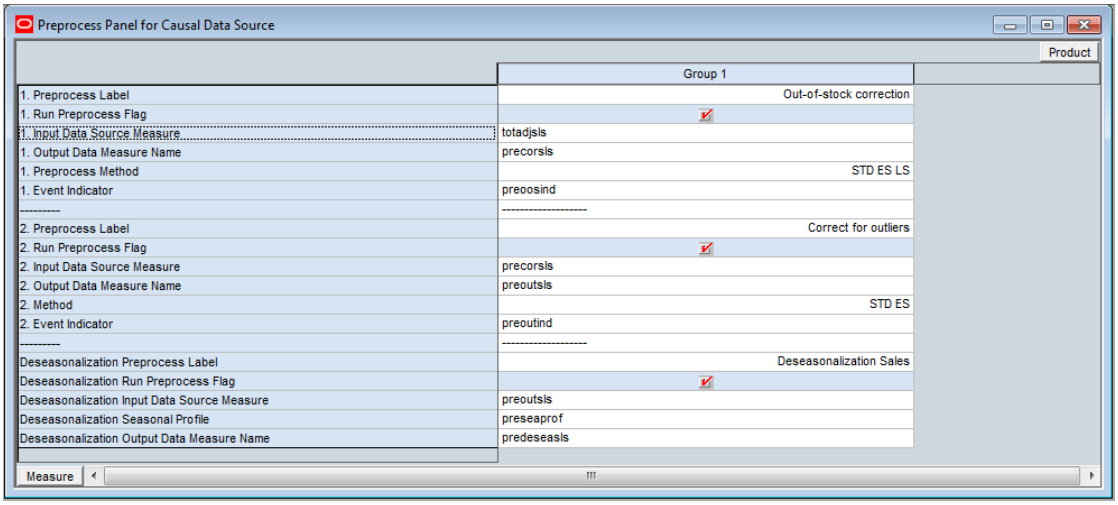

Measures: Preprocessing Panel for Causal Data Source Worksheet

> **Note:** Measures are replicated for each round of preprocessing.

The Preprocessing Panel for Causal Data Source worksheet contains the following measures:

Event Indicator

Indicates the measure name that represents the event flag.

Input Data Source Measure

Indicates the measures that will be corrected.

Output Data Measure

Indicates the measure that stores the result of the preprocessing run.

Preprocess Method

Name of the preprocessing method to be used. This method is fixed, but this can be expanded to a list.

Preprocessing Label

A label denoting the purpose of the preprocessing run, for example, Out-of-Stock Correction.

Run Preprocess Flag

Boolean measure indicating if this rule should be run or skipped.

Preprocess Method Parameters Worksheet

In this worksheet, you can enter values for parameters specific for some of the preprocessing methods available in the special expression. There are four sets of parameters, corresponding to the four runs of the preprocessing special expression. The parameters are entered at the class/store intersection

Figure 7–5 Preprocess Method Parameters Worksheet

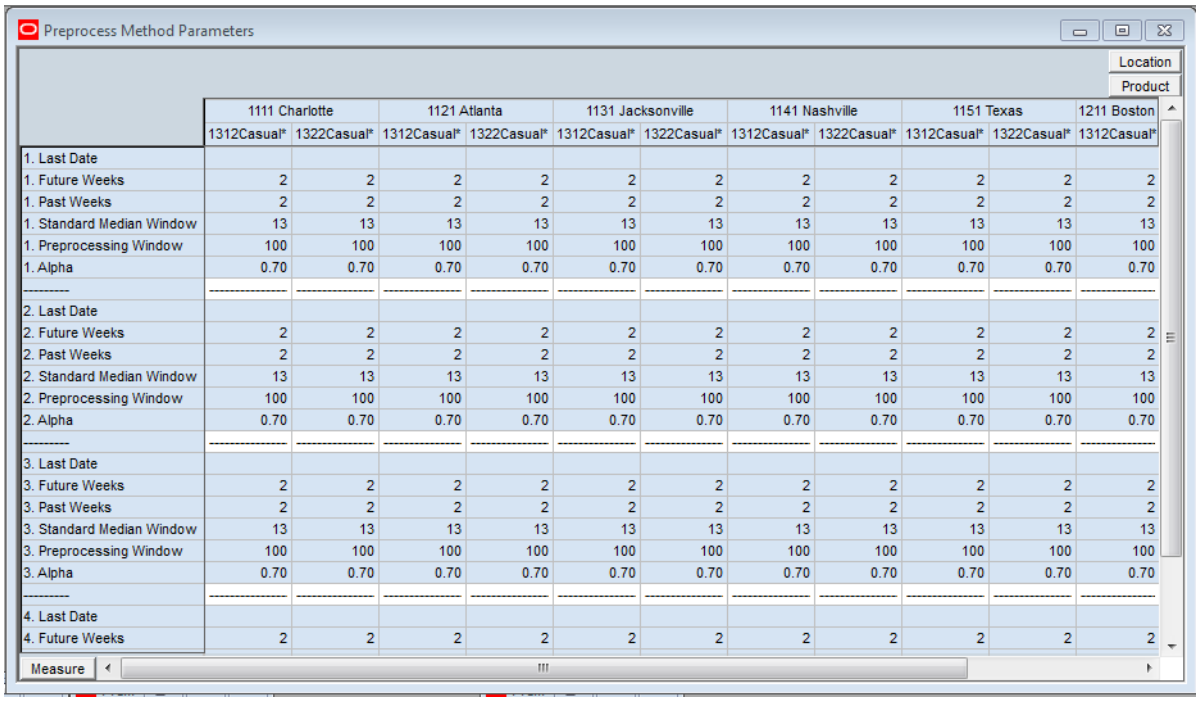

Measures: Preprocessing Methods Parameters Worksheet

The Preprocessing Methods Parameters worksheet contains the following measures:

Alpha

Exponential smoothing coefficient used to calculate past and future velocities.

Future Weeks

This represents the maximum number of data points to calculate the future velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Last Date

This represents the end date of the preprocessing window; it is typically today's date, but can be any date in the past.

Past Weeks

This represents the maximum number of data points to calculate the past velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Preprocessing Window

Number of historical data points that are preprocessed.

Standard Median Window

Filter window length for the Standard Median preprocessing method.

Preprocess Method Parameters Override Worksheet

In this worksheet, you can override values for parameters specific for some of the preprocessing methods available in the special expression. There are four sets of parameters, corresponding to the four runs of the preprocessing special expression.

After all parameters are set and committed back to the domain, usually a batch job will run the pre-processing steps and prepare the source data for forecast generation.

Figure 7–6 Preprocess Method Parameters Override Worksheet

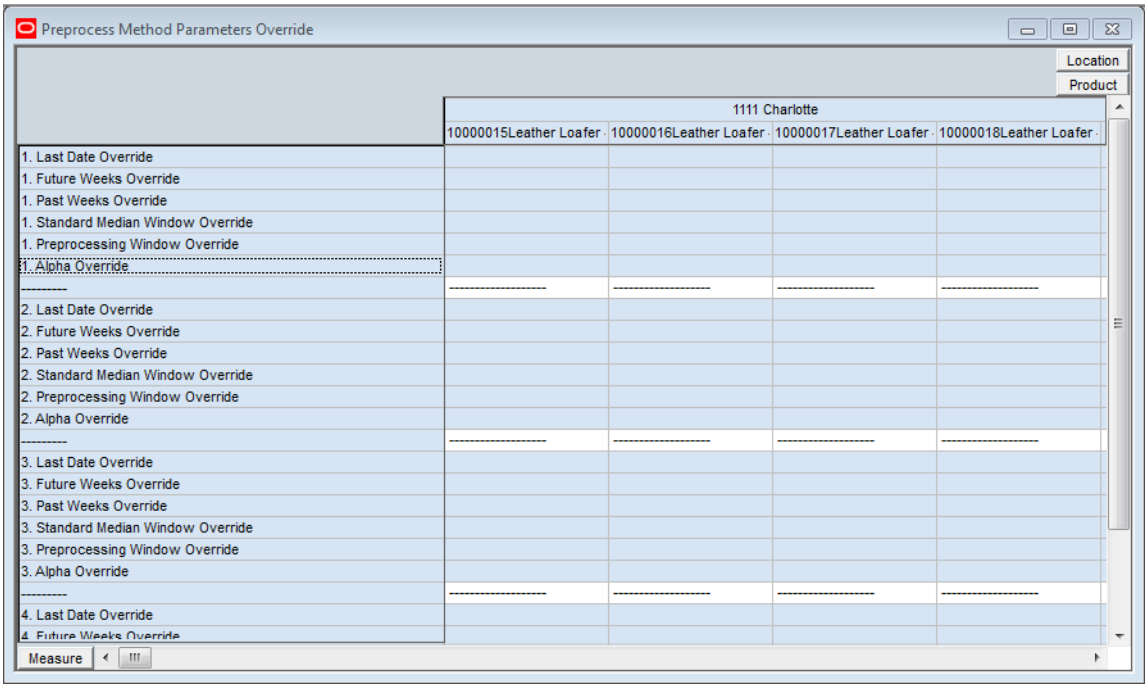

Measures: Preprocessing Method Parameters Override Worksheet

The Preprocessing Method Parameters Override worksheet contains the following measures:

Alpha Override

Exponential smoothing coefficient used to calculate past and future velocities.

Future Weeks Override

This represents the maximum number of data points to calculate the future velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Last Date Override

This represents the end date of the preprocessing window; it is typically today's date, but can be any date in the past.

Past Weeks Override

This represents the maximum number of data points to calculate the past velocity, when using the Standard Exponential Smoothing or Lost Sales Standard Exponential Smoothing preprocessing methods.

Preprocessing Window Override

Number of historical data points that are preprocessed.

Standard Median Window Override

Filter window length for the Standard Median preprocessing method.

What-if Tab

The main purpose of this step is to display time-phased measures that represent input and output to the preprocessing stages, run by the two available custom menus:

- Calculate Out of Stock (OOS) and Outliers
- Preprocess Simulation

Calculate OOS and Outliers

This menu generates the two measures necessary to correct for out of stock and outliers. The calculations triggered by the custom menu populate the:

- **Out of Stock Indicator:** the measure necessary for the out of stock correction
- **Outliers Indicator**: the measure necessary for outlier correction

> **Note:** For additional information, refer to Preprocessing for Outliers.

Preprocess Simulation

This custom menu performs the following actions:

- Correcting for stockouts
- Correcting for outliers
- Depromoting sales
- Smooth sales

The ultimate goal for Preprocess Simulation is to create the source data for generating forecasts:

- Baseline Sales: source data for baseline forecast generation
- Deseasonalized Sales: source data for causal forecast generation

What-if Data Worksheet

This worksheet displays measures that represent input and output of the preprocessing runs, in table format.

Figure 7–7 What-if Data Worksheet

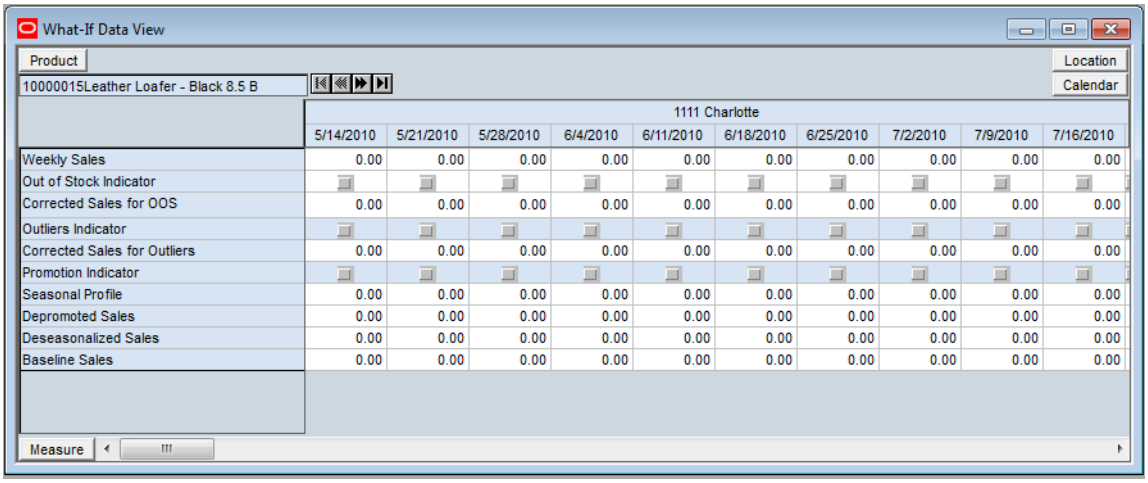

Measures: What-if Data Worksheet

The What-if Data worksheet contains the following measures:

Baseline Sales

This measure is the result of the preprocessing run that smoothes the Depromoted Sales. Note how seasonality is preserved. This measure is used as the data source for the baseline forecasting generation.

Corrected Sales for OOS

This measure is the result of the preprocessing run that corrects for lost sales. It becomes the input to the preprocessing run that corrects for outliers.

Corrected Sales for Outliers

This measure is the result of the preprocessing run that corrects for outliers. Its purpose is twofold. It can be used as input to the preprocessing run that removes promotional sales, to eventually become the data source for baseline forecasting. It can also be used in conjunction with a seasonal profile, to remove seasonality, to become the data source for promotional forecasting.

Depromoted Sales

This measure is the result of the preprocessing run that removes promotional sales. It becomes the input to the preprocessing run that will smooth the sales.

Deseasonalized Sales

This measure is the result of the step where the Corrected Sales for Outliers and the Seasonal Profile have been combined to create sales with seasonality removed, but with promo information. This measure is used as the data source for the causal forecasting generation.

Out of Stock Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects sales for lost sales.

Outliers Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects the sales for outliers.

Promotion Indicator

This measure is usually calculated as the or of all available Boolean promotional variables. It is used during the preprocessing run that removes promotional sales.

Seasonal Profile

This measure represents the seasonal profile that is used to deseasonalize the Corrected Sales for Outliers. This happens using a simple RPAS rule — a call to the preprocessing special expression is not necessary. Note how the promotional sales are not removed.

What-if Plot Worksheet

This worksheet displays measures that represent input and output of the preprocessing runs, in table format. The measures are the same as in the What-if Data Worksheet, but displayed in chart form.

Note: This worksheet needs to be formatted to display the measures in chart form. The user interface cannot be formatted using Configuration Tools.

Figure 7–8 What-if Plot Worksheet

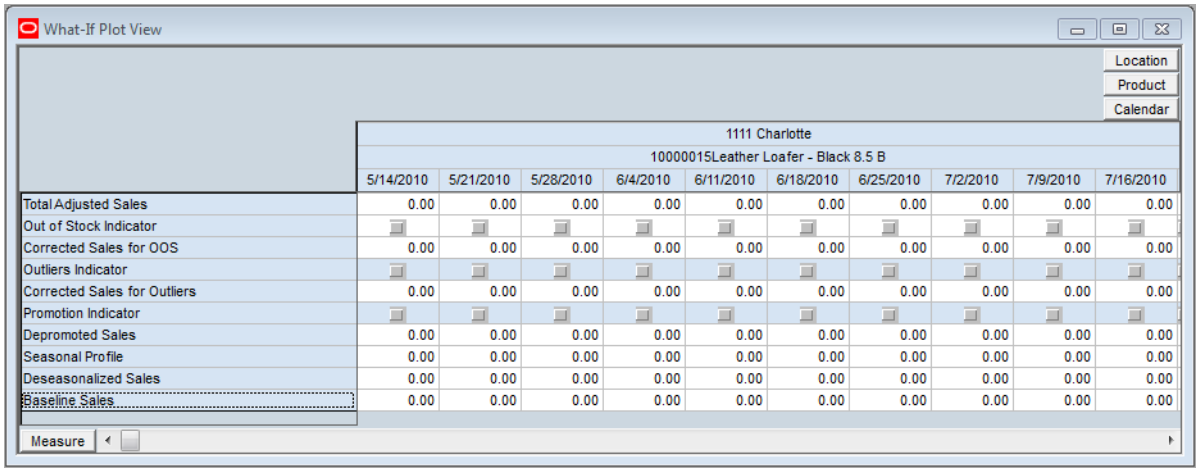

Measures: What-if Plot Worksheet

The What-if Plot worksheet contains the following measures:

Baseline Sales

This measure is the result of the preprocessing run that smoothes the Depromoted Sales. Note how seasonality is preserved. This measure is used as the data source for the baseline forecasting generation.

Corrected Sales for OOS

This measure is the result of the preprocessing run that corrects for lost sales. It becomes the input to the preprocessing run that corrects for outliers.

Corrected Sales for Outliers

This measure is the result of the preprocessing run that corrects for outliers. Its purpose is twofold. It can be used as input to the preprocessing run that removes promotional sales, to eventually become the data source for baseline forecasting. It can also be used in conjunction with a seasonal profile, to remove seasonality, to become the data source for promotional forecasting.

Depromoted Sales

This measure is the result of the preprocessing run that removes promotional sales. It becomes the input to the preprocessing run that will smooth the sales.

Deseasonalized Sales

This measure is the result of the step where the Corrected Sales for Outliers and the Seasonal Profile have been combined to create sales with seasonality removed, but with promo information. This measure is used as the data source for the causal forecasting generation.

Out of Stock Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects sales for lost sales.

Outliers Indicator

This measure is either loaded or calculated by the rules in the custom menu. It is used during the pre-processing run that corrects the sales for outliers.

Promotion Indicator

This measure is usually calculated as the or of all available Boolean promotional variables. It is used during the preprocessing run that removes promotional sales.

Seasonal Profile

This measure represents the seasonal profile that is used to deseasonalize the Corrected Sales for Outliers. This happens using a simple RPAS rule – a call to the preprocessing special expression is not necessary. Note how the promotional sales are not removed.

8

Floating Event Administration Workbook

This chapter describes managing events and holidays for forecast generation in RDF using the Floating Event Administration Workbook.

The Floating Event Administration Workbook includes this tab:

- Floating Event Admin Tab

About Floating Events

When managing periods when events are active, RDF has the Floating Event Administration Workbook. Its one worksheet lists the events and you have the ability to flag periods when certain product/location combinations are active during an event.

A possible workbook scheduled to run before forecast generation is managing events and holidays, which may heavily impact a retailer's business. Some annual events such as Easter occur at a slightly offset week of the year from year to year. RDF needs to have some control over how these floating events and holidays will be managed along the time dimension so that their impact is not forecasted in the wrong time period. A lift associated with the event needs to be calculated, so it can be applied to the correct periods in the forecast horizon.

Management of this issue causes customers the pain, time and cost of configuring their data as this happens every year. This is a basic retail system requirement that the system should be able to address.

In a RDF causal implementation, an annual event that does not occur in the same period every year, but has a spike in demand associated with it, is handled by associating a causal factor to it. The system then determines the associated lift, and applies it to the relevant point in time in the forecast horizon.

In a RDF implementation, where there is no causal, these events still exist. For instance, an increase for sales leading to Easter will happen even if there is no promotion specifically associated with Easter (chocolate bunnies sell more during Easter even though they are not specifically promoted). And because it does not happen in the same time period every year, the Easter spike will appear randomly any time in March or April. The spike is baked in the seasonality of each item/store, making its prediction, timing, and magnitude; inaccurate at best.

RDF is able to handle such situations, not only through its causal capabilities, but also for baseline forecasting.

To do that, a set of events has to be defined. This usually happens at configuration time. Then, a user has to flag the periods where an event was active in the past (history) and in the future (forecast horizon). Once the dates are set up, the effect of the

event is estimated based on past instances, and applied in the forecast horizon if the flag is active.

Floating Event Administration Workbooks, Tabs, and Views

The following table lists the workbooks, tabs, and views for the Floating Event Administration workbook.

Workbook	Tabs	Views
Floating Event Administration Workbook	Floating Event Admin Tab	Floating Event Calendar Maintenance Worksheet

Floating Event Administration Workbook

Use the following procedure to create a Floating Event Administration workbook:

1. Within the Master or Local Domain, select **New** from the **File** menu.
2. Select the **Preprocess** tab to display a list of workbooks. Select Floating Event Administration.
3. Click **OK**. The Floating Event Administration wizard opens.
4. Select the products to include in the workbook.
5. Click **Next**.
6. Select the locations to include in the workbook.
7. Click **Next**.
8. Select the weeks to include in the workbook.
9. Click **Finish** to display the workbook.

Floating Event Admin Tab

This tab contains the Floating Event Calendar Maintenance worksheet.

Floating Event Calendar Maintenance Worksheet

The Floating Event Calendar Maintenance worksheet allows you to flag periods when certain product/location combinations are active during an event.

Figure 8–1 Floating Event Calendar Maintenance Worksheet

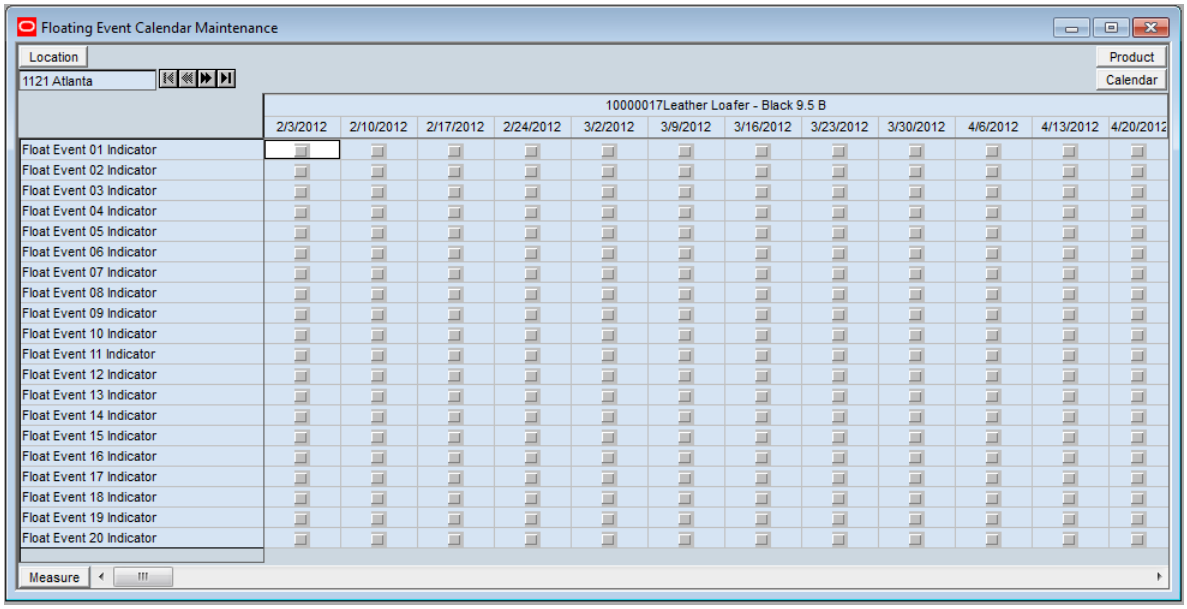

Measures: Floating Event Calendar Maintenance Worksheet

The Floating Event Admin worksheet contains the following measures:

Float Event Indicator (01-20)

These measures, numbered 1-20, allow you to flag up to 20 float events as active in history or to be active in the future.

9

New Item & Locations Workbook

This chapter describes these two features:

- Item attributes — which are used to make automatic like item recommendations for new items.
- New items — including:
  - workflow around handling
  - reviewing and approving their automatic like item recommendation
  - the manual process for like stores selection with new stores

In general, an item is eligible to be considered new if it satisfies the following condition:

- Forecast start date override is in the future *or* Sales history length is less than the time series duration parameter

In general, an item is eligible to be considered a like-item if its recent sales density is acceptable.

These workbooks are in the New Item & Locations Workbook:

- New Item Maintenance Workbook
- New Store Maintenance Workbook
- Attribute Maintenance Workbook

New Item & Locations Workbooks, Tabs, and Worksheets

The following table lists the workbooks, tabs, and worksheets for the New Item & Locations workbook.

Workbook	Tabs	Worksheets
New Item Maintenance Workbook	Like Store Assignment Tab	Select and Approve Worksheet Eligible Like Item Worksheet
New Store Maintenance Workbook	Like Store Assignment Tab	Product Like Store Assignment Worksheet

Workbook	Tabs	Worksheets
Attribute Maintenance Workbook	Review - Attributes Tab	Attribute Match Worksheet New Item Attributes Worksheet Existing Item Threshold Worksheet
	Review - Recommendations Tab	Aggregate Level Worksheet Item & Location Recommendation Worksheet Similarity Score Worksheet

New Item Maintenance Workbook

Note: The full functionality of the New Item Maintenance workbook is available only when item attributes are loaded and thus the like item recommendation is automated.

When item attributes are not available the like item has to be selected manually. However, the workbook is the same for both attribute and non-attribute cases. If at a later point in time, attributes become available, the automated like item recommendation can be used, without the need to patch the environment.

To build the New Item Maintenance workbook, perform these steps:

1. Within the Master or Local Domain, select New from the File menu.
2. Select the New Item Maintenance tab to display a list of workbooks.
3. Select New Item Maintenance.
4. Click **OK**. The New Item Maintenance wizard opens.
5. Select the products to include in the workbook.
6. Click **Finish** to display the workbook.

The New Item Maintenance workbook is built.

Like Item - Select And Approve Tab

This tab contains worksheets that allow you to review and approve like item recommendations, when item attributes are available. If attributes are not available, the like item selection is done manually. You also have visibility to which items are eligible to serve as like items for the new items by location.

The available worksheets are:

- Select and Approve Worksheet
- Eligible Like Item Worksheet

Select and Approve Worksheet

The Select and Approve worksheet is used to reviews the system generated like item recommendations if attributes are available. You can overwrite the recommendations, and/or trigger the approval, at the granular item / store intersection. This worksheet has a custom menu, Approve New Items, and it approves the like items

recommendations for new items. The recommendations are the system recommended or overrides, depending on the approve settings.

Measures: Select and Approve Worksheet

The Select and Approve worksheet contains the following measures:

Approve Date

This measure displays the date when the like item recommendation was approved by running the custom menu.

Approve

This measure determines which recommendations are approved the next time the custom menu is run. The options are:

- Null— no like item is recommended.
- Approve System — the system recommended like item is approved.
- Approve Override — your selected like item is approved. Note that if no user selections are available, then no like item is approved.

Substitute Method

This measure displays a list where you can select the substitute method. When a Substitute Method is used to forecast, the method set for an intersection is cleared once the Default Forecast Start Date is greater than the Forecast Start Date Override plus the Like TS Duration for the intersection. Valid options are:

Substitute Method	Description
None	No Forecast is created for the time series (product location combination)
Seasonal	You provide a like item/location that has a similar seasonality pattern. The new time series' forecast is the like item/locations demand forecast with the applied adjustment. The forecast is set to zero (0) for all dates before the new product/location's start date.
Lifecycle	You provide a like item/location that had a similar lifecycle pattern as the new item/location. The new item/location's forecast is the like/location's actual historical demand with the applied adjustment shifted such that the like item/locations first sales matches the new item/location's forecast start date.
Cloning	You provide a like item/location that has a similar selling pattern as the new item/location. The historical demand of the like item/location is copied into the historical sales of the new item/location. During forecasting, the forecast of the new item/location is generated based on the new item/location's own, copied historical demand.

Substitute Method	Description
Base Rate of Sales	You provide a base rate of sales for a new item/store combination. The new product's forecast is a combination of the seasonality of the forecast at the corresponding source level and the base rate of sales. Specifically, the formula to calculate the forecast is: *Forecast at time t = source level forecast at time t divided by the average of the source level forecast times the base rate of sales*. The base rate of sale is a convenient way to generate (seasonal) forecasts for a new item. No Like item or Clone item is needed. What is necessary is a value for how much the item is selling on average per period. That value works as the interim forecast used to spread the source level forecast down to final level. Because this item is new, the system knows that as you have specified a history start date. The base rate of sale forecast is generated until the TS duration is reached. The TS duration is the difference between forecast start date and history start date. Usually the default history start date is empty, which means it is the beginning of the calendar. If the history start date override is also empty, the forecast start date less the beginning of the calendar (>=2 years) is probably much more than the TS duration, so the item is not considered new. In this case base rate of sales or any new item method, will not work.

Percent Contribution of Like Item 1

This measure determines the percentage of Like item 1's forecast that is going to be applied towards the forecast of the new item.

Percent Contribution of Like Item 2

This measure determines the percentage of Like item 2's forecast that is going to be applied towards the forecast of the new item.

Percent Contribution of Like Item 3

This measure determines the percentage of Like item 3's forecast that is going to be applied towards the forecast of the new item.

Adjustment Factor

This measure determines how much the combined forecasts are scaled up or down to create the forecast of the new item.

System Recommended Like Item 1

This measure displays the top like item for a given store.

System Recommended Like Item 2

This measure displays the second like item for a given store.

System Recommended Like Item 3

This measure displays the third like item for a given store.

User Selected Like-Item 1

This measure allows you to override the top like item for a given store.

User Selected Like-Item 2

This measure allows you to override the second like item for a given store.

User Selected Like-Item 3

This measure allows you to override the third like item for a given store.

Base Rate of Sales

This measure represents the average sales of a new item/store combination. It is used when specifying *Base Rate of Sales New SKU* as the Substitution Method to create a forecast for the new item/store combination. The measure can be generated in another application and loaded into RDF, or it can be manually entered by a user.

Forecast Start Date Override

This measure represents the date to start forecasting for an item/location combination. This measure can be set in the future if using like-item or Sister-Store functionality, and, upon reaching that time, the forecast is 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 the Forecast Start Date for this intersection needs to be edited once it is no longer in the future. For like-item or sister store, the Forecast Start Date and the History Start Date should be set to the same date. It is important to understand how Forecast Start Date should be used in conjunction with Forecast End Date. No value is in this measure if the system default set in the Forecast Administration Workbook is to be used.

Note: This measure can also be set in the Forecast Maintenance Workbook. Changes to this measure can be seen in the Forecast Maintenance workbook. The most recent commit (between either of the tasks) is the value used by the system.

Eligible Like Item Worksheet

The Eligible Like Item worksheet is used to display what items are eligible to be like items for each new item included in the worksheet

Measure: Eligible Like Item Worksheet

The Eligible Like Item worksheet contains the following measure:

Product Location Eligibility

This measure displays what items are eligible to be like items for each new item included in the worksheet. Note how the eligibility of an item can be different by location.

New Store Maintenance Workbook

Use this workbook to assign like stores to handle forecasting for new stores. The like store assignment is manual and there is a good reason for it. New stores have a large financial impact, so it makes sense having a business person making the like store selection.

It is probably more appropriate than going with an automatic selection based on something like store attributes. Also, new store introductions are infrequent compared to new item introductions, so manually handling new stores is not a significant overhead activity.

To build the New Store Maintenance workbook, perform these steps:

1. Within the Master or Local Domain, select New from the File menu.
2. Select the New Store Maintenance tab to display a list of workbooks.
3. Select New Store Maintenance.
4. Click **OK**. The New Item Store wizard opens.
5. Select the locations to include in the workbook.
6. Click **Finish** to display the workbook.

The New Store Maintenance workbook is built.

Like Store Assignment Tab

This tab contains the Product Like Store Assignment Worksheet.

Product Like Store Assignment Worksheet

The Product Like Store Assignment worksheet is at the intersection of prod/location, so all parameters can vary by product. For example, a new store opening in the Midwest can have a Like Store from Alaska for items in the Shovels department. However, for summer items, the Like Store is picked from the Northeast region.

Measures: Product Like Store Assignment Worksheet

The Product Like Store Assignment worksheet contains the following measures:

Like Store 1

In this measure, you can specify the first like store. Note how the like store can be different by product. In RDF, a different first like store selection can be made for every subclass.

Like Store 2

In this measure, you can specify the second like store. Note how the like store can be different by product. In RDF, a different second like store selection can be made for every subclass.

Like Store 3

In this measure, you can specify the third like store. Note how the like store can be different by product. In RDF, a different third like store selection can be made for every subclass.

Adjustment Factor

This measure determines how much the combined forecasts are scaled up or down to create the forecast of the new item.

Substitute Method

This measure displays a list where you can select the substitute method. When a Substitute Method is used to forecast, the method set for an intersection is cleared once

the Default Forecast Start Date is greater than the Forecast Start Date Override plus the Like TS Duration for the intersection. Valid options are:

Substitute Method	Description
None	No Forecast is created for the time series (product location combination)
Seasonal	You provide a like item/location that has a similar seasonality pattern. The new time series' forecast is the like item/locations demand forecast with the applied adjustment. The forecast is set to zero (0) for all dates before the new product/location's start date.
Lifecycle	You provide a like item/location that had a similar lifecycle pattern as the new item/location. The new item/location's forecast is the like/location's actual historical demand with the applied adjustment shifted such that the like item/locations first sales matches the new item/location's forecast start date.
Cloning	You provide a like item/location that has a similar selling pattern as the new item/location. The historical demand of the like item/location is copied into the historical sales of the new item/location. During forecasting, the forecast of the new item/location is generated based on the new item/location's own, copied historical demand.
Base Rate of Sales	You provide a base rate of sales for a new item/store combination. The new product's forecast is a combination of the seasonality of the forecast at the corresponding source level and the base rate of sales. Specifically, the formula to calculate the forecast is: *Forecast at time t = source level forecast at time t divided by the average of the source level forecast times the base rate of sales*. The base rate of sale is a convenient way to generate (seasonal) forecasts for a new item. No Like item or Clone item is needed. What is necessary is a value for how much the item is selling on average per period. That value works as the interim forecast used to spread the source level forecast down to final level. Because this item is new, the system knows that as you have specified a history start date. The base rate of sale forecast is generated until the TS duration is reached. The TS duration is the difference between forecast start date and history start date. Usually the default history start date is empty, which means it is the beginning of the calendar. If the history start date override is also empty, the forecast start date less the beginning of the calendar (>=2 years) is probably much more than the TS duration, so the item is not considered new. In this case base rate of sales or any new item method, will not work.

Percent Contribution of Like Store 1

This measure determines the percentage of Like Store 1's forecast that is going to be applied towards the forecast of the new item.

Percent Contribution of Like Store 2

This measure determines the percentage of Like Store 2's forecast that is going to be applied towards the forecast of the new item.

Percent Contribution of Like Store 3

This measure determines the percentage of Like Store 3's forecast that is going to be applied towards the forecast of the new item.

Forecast Start

This measure specifies the first date for which forecast is generated for the new item.

Attribute Maintenance Workbook

This workbook is intended to review like item recommendations, as well as metrics that support the recommendations. The recommendations are driven by similarity among items, which in turns is based on how close a new item's attributes are compared to all existing items' attributes.

The workbook can be built by manually selecting the new items during the wizard process, or the new item selection can be based on the available new item alert.

To build the Attribute Maintenance workbook, perform these steps:

1. Within the Master or Local Domain, select New from the File menu.
2. Select the Attribute Maintenance tab to display a list of workbooks.
3. Select Attribute Maintenance.
4. Click **OK**. The Attribute Maintenance wizard opens.
5. Select the products to include in the workbook and click **Next**.
6. Select the locations to include in the workbook.
7. Click **Finish** to display the workbook.

The Attribute Maintenance workbook is built.

Review - Attributes Tab

This tab contains worksheets that allow you to review attributes for new and existing items. Also, they show how much alike items are and the best choices for like items.

The available worksheets are:

- Attribute Match Worksheet
- New Item Attributes Worksheet
- Existing Item Threshold Worksheet

Attribute Match Worksheet

The Attribute Match worksheet is used to review the attributes of new and existing items and how well they match.

Measures: Attribute Match Worksheet

The Attribute Match worksheet contains the following measures:

Attribute Value

This measure displays the content of the attribute. For instance the unit of measure (UOM) attribute can have different values. It can be inch or XL for fashion items. Or it can be ounce or grams for grocery items.

Attribute Scores

This measure displays the quantitative fit of the attribute values between new and existing items. For instance we can compare the UOM attribute between a coffee pack and a pair of jeans. The relevance of matching ounces and inches may not be very high, and the attribute score is likely zero. However, if we compare the color attribute of a shirt and a t-shirt, the match be more relevant and the score is larger than zero.

This measure incorporates the goodness of the fit in attributes between new and existing items, as well as how important an attribute is for the new item. If the new item is a carbonated drink, the brand attribute may be much more relevant than the price tier, because the consumer is prepared to pay a higher price for a brand item.

New Item Attributes Worksheet

The New Item Attributes worksheet displays attribute information about new items.

Measures: New Item Attributes Worksheet

The New Item Attributes worksheet contains the following measures:

Attribute Weight

This measure displays the relative importance of the attributes for a given new item. While the flavor family may not be important for a t-shirt, and the attribute has a weight of zero, the brand and color attributes definitely are. Their relative importance may be 0.2 for the color and 0.3 for the brand.

Note: That the sum of all attribute weights does not need to be 1 for every item. The automatic like item recommendation algorithm is taking care of it. The most likely scenario is that this measure is loaded, but it can be review and adjusted in this worksheet.

Attribute Value

This measure displays the content of the attribute. For instance the unit of measure (UOM) attribute can have different values. It can be inch or XL for fashion items. Or it can be ounce or grams for grocery items.

Existing Item Threshold Worksheet

The Existing Item Threshold worksheet is used to adjust the parameter that decides if an item is eligible to be selected as like item for a new item.

Measure: Existing Item Threshold Worksheet

The Existing Item Threshold worksheet contains the following measure:

Threshold Recent Sales Density

The value of this measure decides if an item can be recommended as like item for new items. The idea behind determining the eligibility is that an item needs to be actively selling to be eligible. We don't want a stale item, or item with no sales to be selected as like item even if item attributes match very well. To determine the eligibility we first calculate the count of non-zero sales in the most recent periods given by the TS duration. Then we divide the count by the TS duration to get the sales density. Finally, the density is compared with the value of the threshold. If it is larger, then the item can be assigned as like item.

Review - Recommendations Tab

This tab contains worksheets that allow you to review the top choices for like items for the new items, based on similarity among items.

The available worksheets are:

- Aggregate Level Worksheet
- Item & Location Recommendation Worksheet
- Similarity Score Worksheet

Aggregate Level Worksheet

The Aggregate Level worksheet is at the item/item RHS intersection. The item represents the new items, while the item RHS represents the like items. The item RHS dimension has only positions that were identified as like items for new items.

Measures: Aggregate Level Worksheet

The Aggregate Level worksheet contains the following measure:

Aggregated Store Count

This measure displays the number of stores at which an item RHS position was identified as the most suitable like item for a new item. For instance, for new item A, item BA was identified to be the like item for 20 stores. For the rest of 42 stores, it was item CA. This can happen simply because the best fit - in this case item BA - is not sold in all stores. Item CA is the second best fit, but it's sold in more stores, and thus the subjective better fit.

Item & Location Recommendation Worksheet

The Item & Location Recommendation worksheet is at the item / store intersection. For new items it displays the top three matching items for a certain store.

Measures: Item & Location Recommendation Worksheet

The Item & Location Recommendation worksheet contains the following measures:

System Recommended Like Item 1, 2, 3

This measure displays the top three like items (based on similarity and store ranging) for every store. For a given new item, the like items may be different by store. The reason is that for some stores, the existing item with the highest similarity is not available for sale. Then, the algorithm is picking the existing item with the highest score for that store.

Similarity Score Worksheet

The Similarity Score worksheet is at the item/item RHS/attribute intersection. The item dimension contains new item which are manually selected during the wizard process. The new item alert can also be used to range down the items to a relevant selection. The item RHS dimension is ranged to only relevant existing items that are eligible to be like items. The measure shows the calculated similarity between new and existing items.

Measures: Similarity Score Worksheet

The Similarity Score worksheet contains the following measure:

Similarity Scores

The Similarity Score is a measure of how well a new item's demand behavior can be modeled after a certain existing item's. The higher the score, the better the fit, and the better the chance that the existing item is going to be the like item

10

Extra Week Administration Workbook

The Extra Week Administration workbook allows you to manage forecasts for years that include an extra 53rd week. For additional information, see the section, 53 Week Calendar.

For forecasting methods which only require demand as input, the Extra Week functionality works by setting up the Extra Week Indicator. However, there are methods that require additional inputs, like Causal, Bayesian or profile-based.

For these, the Extra Week is handled by the additional inputs. For the Causal method, it is handled by the baseline, which is an input to the method.

For Bayesian and profile-based methods, it is handled by the plan and the profile, respectively.

Extra Week Administration Workbook

Use the following procedure to create a Extra Week Administration workbook:

1. Within the Master or Local Domain, select **New** from the **File** menu.
2. Select the **Preprocess** tab to display a list of workbooks. Select Extra Week Administration.
3. Click **OK**. The Extra Week Administration wizard opens.
4. Select the weeks to include in the workbook.
5. Click **Next**.
6. Select the groups to include in the workbook.
7. Click **Finish** to display the workbook.

Extra Week Indicator Setup Worksheet

The Extra Week Indicator Setup worksheet allows you to set which week is the 53rd, or extra week. The measure can be selected in the Extra Week Indicator Data Source found in the Final or Source System Parameters Worksheet.

Figure 10–1 Extra Week Indicator Setup Worksheet

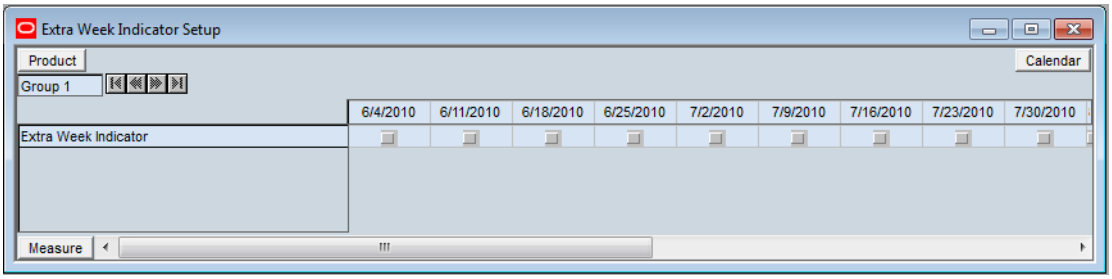

Measures: Extra Week Indicator Setup Worksheet

The Extra Week Indicator Setup worksheet contains the following measure:

Extra Week Indicator

This indicator sets which week is the 53rd, or extra week. The Extra Week Indicator Data Source references the Extra Week Indicator.

11

Alert Thresholds Task

The Alert Thresholds task allows you to manage the thresholds used in the batch and workbook alerts included in RDF. The thresholds can be set at three levels:

- Domain level
- Intermediate level (lower than all Product in the merchandise hierarchy)
- Item/store level (allows threshold values to be overwritten)

Defined Alerts

The following sections provide the complete list of alerts defined for RDF. Most of the alerts have thresholds that can be adjusted in the Alert Threshold Task. The alerts that do not have any adjustable thresholds defined are still listed, and their calculation expressions are mentioned.

The alert thresholds are available for the baseline as well as causal levels. However, based on the measures in the calculation expressions we specify for which level each alert is most relevant.All alerts are batch alerts, and can be specified in the Forecast Administration workbook as forecast approval alerts. Some of the batch alerts go along with real time alerts, which have similar calculation expressions. The calculations can't be identical, because the batch alerts are at the production/location intersection, while the real time alerts also have the time dimension.

The way the batch and the real time alerts can be used together, is detailed in the following steps. For example purposes, we will use the Current Forecast versus Last Approved Forecast alert, for both batch and real time. However, once we describe each alert, we specify if it has a real time counterpart and the levels for which it is relevant.

1. First the user selects a batch alert, for instance Forecast versus Last Approved Forecast, to be the choice for the Approval method.
2. Then the batch is run. After the forecast is generated, it is compared to the last approved forecast. For item/stores for which the two forecasts are different with respect to attached thresholds, the alert is triggered, and the time series is not approved.
3. Then the batch is run. After the forecast is generated, it is compared to the last approved forecast. For item/stores for which the two forecasts are different with respect to attached thresholds, the alert is triggered, and the time series is not approved.
4. Once in the workbook, the user can use the batch alert to navigate from time series to time series, or she can use the real time alert to navigate to alerted time periods.

   For additional information, refer to the Forecast Administration Workbook.

Forecast versus Recent Sales

Usually it is not expected that demand values differ very much from period to period. This also implies that the forecast magnitude generally is in line with the magnitude of the most recent sales. There are exceptions to this case. For instance, when an item enters a season, the forecast is probably higher than the sales in periods leading to the season. Or when an item is towards the end of the season, the forecast will be lower than sales in peak periods. For these exceptions, you can be alerted to review the forecast, rather than automatically approving it.

This alert a batch alert, and can be used for forecast approval.

Alert Expression

Where *length*, *threshold1*, and *threshold2* are adjustable parameters.

$$\frac{\sum_{today}^{today+length-1} system\ forecast(t)}{length} > threshold1$$

$$\&\&$$

$$\left| \frac{\sum_{today}^{today+\ length-1} adjusted\ forecast(t)}{\sum_{today-length}^{today-1} forecast\ source(t)} - 1 \right| > threshold2$$

The calculation expression for this alert involves thresholds which can be adjusted in the Alert Thresholds task.*Threshold1* allows filtering out time series by rate of sales. For instance, if it is set to 1, the alert expression is only run for time series that sell at least one item per week. The reason is that we don't want the user to be overwhelmed by reviewing item/store/week information for low sellers. *Threshold2* defines how different the forecast and the recent sales may be before they get alerted. For instance, if *threshold2=0.1*, a difference of 10% is considered acceptable, but everything above will trigger the alert, and the time series is not approved.

The alert is mostly relevant for the baseline level. The reason is that there can be large differences in recent sales versus forecast due additional demand coming from past or upcoming promotions.

Note: The calculations are not performed for the entire forecast horizon, but rather by the number of periods determined by the *length* parameter. The reason is that the forecast horizon can sometimes be very long (52 weeks) and average demand over such a long time period can not be used as in-season versus out of season rate of sales.

Current Forecast versus Last Approved Forecast

Typically RDF generates forecasts weekly. Every week, new sales data is loaded and the forecast is regenerated. While the latest data points are expected to make the forecast more accurate, it is not expected that the difference in forecasts generated in two consecutive weeks to vary too much. If the forecasts differ, you are alerted to review the forecasts.

This alert is both a batch and a workbook alert.

Alert Expression

Where *length*, *threshold1*, and *threshold2* are adjustable parameters.

For the batch alert

$$\frac{\sum_{today}^{today+length-1} system\ forecast(t)}{length} > threshold1$$

$$\&\&$$

$$\left| \frac{\sum_{today}^{today+\ length-1} adjusted\ forecast(t)}{\sum_{today}^{today+length-1} last\ approved\ forecast(t)} - 1 \right| > threshold2$$

For the real time alert:

$$\frac{\sum_{today}^{today+length-1} system\ forecast(t)}{length} > threshold1$$

$$\&\&$$

$$\left| \frac{adjusted\ forecast(t)}{lappf(t)} - 1 \right| > threshold2$$

Note: The batch alert does not have a time component, while the real time alert is time-phased.

The calculation expression for this alert involves thresholds which can be adjusted in the Alert Thresholds task.*Threshold1* allows filtering out time series by rate of sales. For instance, if it is set to 1, the alert expression is only run for time series that sell at least one item per week. The reason is that we do not want the user to be overwhelmed by reviewing item/store/week information for low sellers. *Threshold2* defines how different the two forecasts may be before they get alerted. For instance, if *threshold1=0.1*, a difference of 10% is considered acceptable, but everything above will trigger the alert, and the time series is not approved.

The values of *Threshold1*, and *Threshold2* are the same for both batch and real time alerts.

This alert is relevant for both baseline and causal forecasts.

Note: Summation of forecasts are not performed over the entire forecast length. If the forecast horizon is very long, the user may only be interested in periods closer to the forecast generation date, and not want to check periods far out. This can be controlled by the *length* parameter

Forecast Accuracy Alert

The following steps implement the accuracy and alert calculations:

1. Calculate rate of sales (ROS) per item/store. There are several places where average sales are calculated. If available, reuse an available measure that stores the desired rate of sales.
2. Define thresholds for buckets of ROS. For example:

Low ROS value = 1
High ROS value =10

3. Define error values by low/medium/high ROS. For example:

Low ROS error value = 10000%
Medium ROS error value = 100%
High ROS error value = 40%

4. Define the alert thresholds per rate of sales at the item/store intersection. For example, an item/store that has average sales of half a unit per week, is not worth the effort to be reviewed even if the forecast error is more than 100%. The logic for calculating the threshold is:

   If ROS < Low ROS error value

   Alert threshold = Low error value

   Else if ROS < High ROS value

   Alert threshold = Medium ROS error value

   Else

   Alert threshold = High ROS error value

5. Calculate the forecast error:

   If ROS is less or equal to Low ROS error value for error calculation

   Forecast error = (abs(sum of forecast)-abs(sum of sales))/sum of sales * 100

   Else

   Forecast error = (sum(abs(forecast-sales))/sum of sales * 100

   Note: The second formula yields a higher error. This formula is applied to medium and high ROS items.

   The calculation of the forecast error should happen on a window starting the period prior to today and going back an adjustable number of periods.

6. Trigger the alert if the error is larger than the threshold (calculated in Step 4).

   If Forecast error > Alert threshold

   Trigger alert

Forecast versus Last Year Sales

The most reliable forecasts are generated from data that has a repeatable pattern year over year. However, this is not always the case. A change in business strategy,

merchandise reclassifications, new items can all lead to changing selling patterns over time.

To detect possible changes in selling patterns, the following alert will compare the last year's sales volume with the forecasted sales volume. If they are different by an adjustable percent, the alert is triggered.

Alert Expression

Where *length*,*threshold1*, and *threshold2* are adjustable parameters.

For the batch alert:

$$\left\{\frac{\sum_{today}^{today+length-1} system\ forecast(t)}{length} > threshold1\right.$$

$$||$$

$$\left.\frac{\sum_{today}^{today+length-1} forecast\ source\ LY(t)}{length} > threshold1\right\}$$

$$\&\&$$

$$\left|\frac{\sum_{today}^{today+forecast\ length} adjusted\ forecast(t)}{\sum_{today}^{today+forecast\ length} forecast\ source\ LY\ (t)} - 1\right| > threshold2$$

For the real time alert:

$$\left\{\frac{\sum_{today}^{today+length-1} system\ forecast(t)}{length} > threshold1\right.$$

$$||$$

$$\left.\frac{\sum_{today}^{today+length-1} forecast\ source\ LY(t)}{length} > threshold1\right\}$$

$$\&\&$$

$$\left|\frac{adjusted\ forecast(t)}{forecast\ data\ source\ LY(t)} - 1\right| > threshold2$$

Note: The batch alert does not have a time component, while the real time alert is time-phased.

The calculation expression for this alert involves thresholds which can be adjusted in the Alert Thresholds task.First, we compare the average forecast and average sales last year to *Threshold1*. If *Threshold1= 1*, we make sure to run the alert only for item/stores that sell on average at least one per week. If the average is not too low, we proceed

comparing the current adjusted forecast to the sales last year. If they are very different (defined by *Threshold2*), the user is alerted to review the forecast.

The alert is mostly relevant for the baseline level. The reason is that there can be large differences in last year sales versus forecast due additional demand coming from past or upcoming promotions.

Note: Summation of forecasts are not performed over the entire forecast length. If the forecast horizon is very long, the user may only be interested in periods closer to the forecast generation date, and not want to check periods far out. This can be controlled by the *length* parameter.

New Item Alert

When a new item for any store is introduced into the assortment, forecasts need to be generated such that it is correctly replenished. Since no historical demand is available, RDF has a few mechanisms to create the forecasts. To do that, it first needs to be aware that new items have entered the system.

Alert Expression

If forecast startdate override is in the future

and POS length < TS duration

and substitution method is blank

and no like item is assigned

then item is considered new

The calculation expression for this alert does not involve thresholds which can be adjusted in the Alert Thresholds task.

The alert is not designed specifically for a forecast level, but rather to trigger the new item functionality available in RDF.

Causal Peaks Alert

The purpose of this alert is to check how large the forecast peaks are compared to historical demand. The peaks can come from various effects like promotions, price discount, demand transference due to assortment changes, and so on. The most common, though, are due to price changes and promotions.

There is a batch and a workbook alert, with similar calculations, but different intersections;

Batch Alert Expression

The batch alert has the following calculation:

Calculate the maximum historical causal data source:

-Max_hist_dmnd = maximum of source data for causal forecasting (item/store)

Calculate the adjusted peaks in forecast region:

-adj frcst_pks = causal forecast –baseline forecast (item/store/week)

Calculate alert:

-If max(fadjrcst_pks) > Max_hist_dmnd then trigger alert (item/store)

The business case this addresses is to alert you when the peaks in the forecast region are larger than any observed sales in the past. There may be valid justification for this, for instance, several events are active in the same time period, thus creating a huge spike in demand. You can review the alert and take action.

Real Time Alert Expression

The workbook alert condition, based on the information already calculated in the batch version, is:

Adj frcst_pks(t) > Max_hist_dmnd

The target measure for the alert should be the system forecast.

The advantage of the workbook alert is that you can immediately spot the values where the business rule is violated, instead of checking the values for the entire forecast region.

The calculation expression for this alert does not involve thresholds which can be adjusted in the Alert Thresholds task.

If the peaks only come from promotions and/or price discounts, then the alert should be enabled only for the causal level.

Note: The batch alert does not have a time component, while the real time alert is time-phased.

Alert Thresholds Workbook

To build the Alert Thresholds workbook, perform these steps:

Use the following procedure to create a Alert Thresholds workbook:

1. Within the Master or Local Domain, select **New** from the **File** menu.
2. Select the **Forecast** tab to display a list of workbooks. Select Forecast Alert.
3. The Forecast Alert 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.
4. Select the products to include in the workbook.
5. Click **Next**.
6. Select the locations to include in the workbook.
7. Click **Finish** to display the workbook.

Forecast Alert Default Parameters Worksheet

The Forecast Alert Default Parameters worksheet allows you to adjust the default values for alert thresholds.

Figure 11–1 Forecast Alert Default Parameters Worksheet

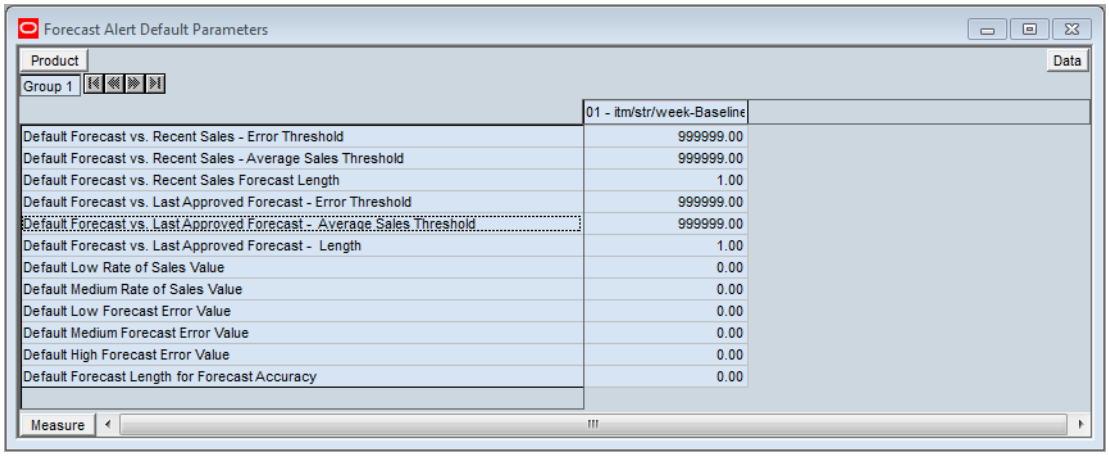

Measures: Forecast Alert Default Parameters Worksheet

Note: All information in this worksheet can be edited and represent domain level values.

The Forecast Alert Default Parameters worksheet contains the following measures:

Default Forecast Length for Forecast Accuracy

This parameter defines the time frame over which the calculations are performed for the forecast accuracy alert.

Default Forecast vs Last Approved Forecast - Average Sales Threshold

This parameter defines the minimum rate of sales a product/location must have in order to be considered in the alert calculation. If the rate of sales condition is met and the forecast error is higher than allowed, the alert is triggered for the product/location.

Default Forecast vs Last Approved Forecast Length

This parameter defines the time frame over which the alert calculations are performed. The value should not exceed the forecast length. The calculation is still performed but will likely not be accurate.

Default Forecast vs Last Approved Forecast Threshold

This parameter defines the allowable relative error between system forecast and last approved fore-cast. If the error is higher than the threshold and the condition on the average rate of sales is met, the alert is triggered for the product/location.

Default Forecast vs Recent Sales - Average Sales Threshold

This parameter defines the minimum rate of sales a product/location must have in order to be considered in the alert calculation. If the rate of sales condition is met and the forecast error is higher than allowed, the alert is triggered for the product/location.

Default Forecast vs Recent Sales - Forecast Length

This parameter defines the time frame over which the alert calculations are performed. The value should not exceed the forecast length. The calculation is still performed but will likely not be accurate.

Default Forecast vs Recent Sales Threshold

This parameter defines the allowable relative error between system forecast and forecast data source. If the error is higher than the threshold and the condition on the average rate of sales is met, the alert is triggered for the product/location.

Default High Forecast Error Value

Forecast accuracy for items with high rate of sales, as defined by the Default Medium Rate of Sales Value, is tested against this parameter to determine if an alert is triggered.

Default Low Forecast Error Value

Forecast accuracy for items with low rate of sales, as defined by the Default Medium Rate of Sales Value, is tested against this parameter to determine if an alert is triggered.

Default Low Rate of Sales Value

If the rate of sales is less than this value, the item is considered to have a low rate of sale. Its forecast accuracy is tested against the Default Low Forecast Error Value, to determine if an alert is triggered.

Default Medium Forecast Error Value

Forecast accuracy for items with medium rate of sales, as defined by the Default Medium Rate of Sales Value, is tested against this parameter to determine if an alert is triggered.

Default Medium Rate of Sales Value

If the rate of sales is less than this value, the item is considered to have a medium rate of sale. Its forecast accuracy is tested against the Default Medium Forecast Error Value, to determine if an alert is triggered. If the rate of sales is more than this value, the item is considered to have a high rate of sale. Its forecast accuracy is tested against the Default High Forecast Error Value, to determine if an alert is triggered.

Forecast Alert Final Level Intersection Parameters Worksheet

The Forecast Alert Final Level Intersection Parameters worksheet allows you to adjust the values for alert thresholds at the final level intersection.

Figure 11–2 Forecast Alert Final Level Intersection Parameters Worksheet

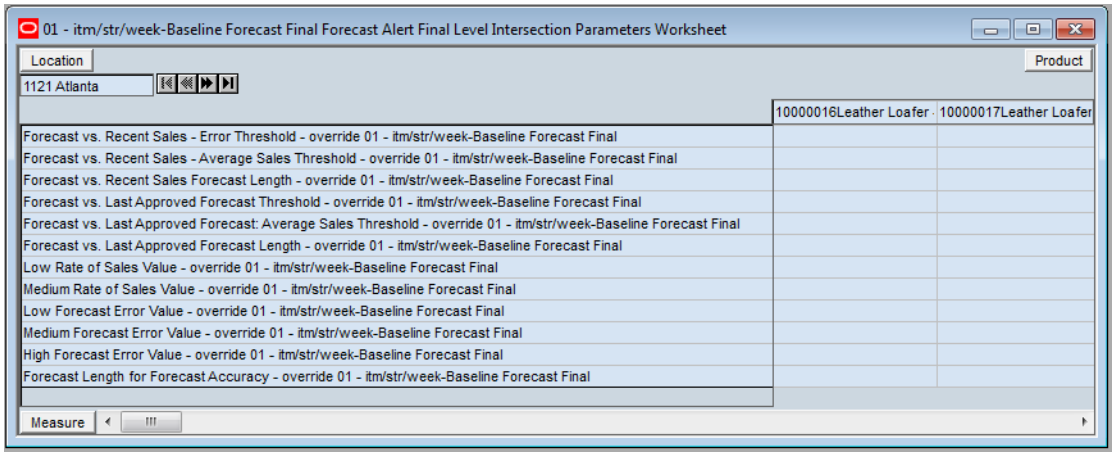

Measures: Forecast Alert Final Level Intersection Parameters Worksheet

Note: All information in this worksheet can be edited. All are defined for measures at the final level intersection. and are meant to serve as an override of the default and intermediate intersection values.

The Forecast Alert Final Level Intersection Parameters worksheet contains the following measures:

Forecast Length for Forecast Accuracy - override

This parameter defines the time frame over which the calculations are performed for the forecast accuracy alert.

Forecast vs Last Approved Forecast Average Sales Threshold - override

This parameter defines the minimum rate of sales a product/location must have in order to be considered in the alert calculation. If the rate of sales condition is met and the forecast error is higher than allowed, the alert is triggered for the product/location.

Forecast vs Last Approved Forecast Length - override

This parameter defines the time frame over which the alert calculations are performed. The value should not exceed the forecast length. The calculation is still performed but will likely not be accurate.

Forecast vs Last Approved Forecast Threshold - override

This parameter defines the allowable relative error between system forecast and last approved fore-cast. If the error is higher than the threshold and the condition on the average rate of sales is met, the alert is triggered for the product/location.

Forecast vs Recent Sales Average Sales Threshold - override

This parameter defines the minimum rate of sales a product/location must have in order to be considered in the alert calculation. If the rate of sales condition is met and the forecast error is higher than allowed, the alert is triggered for the product/location.

Forecast vs Recent Sales Forecast Length - override

This parameter defines the time frame over which the alert calculations are performed. The value should not exceed the forecast length. The calculation is still performed but will likely not be accurate.

Forecast vs Recent Sales Threshold - override

This parameter defines the allowable relative error between system forecast and forecast data source. If the error is higher than the threshold and the condition on the average rate of sales is met, the alert is triggered for the product/location.

High Forecast Error Value - override

Forecast accuracy for items with high rate of sales, as defined by the Medium Rate of Sales Value – override, is tested against this parameter to determine if an alert is triggered.

Low Forecast Error Value - override

Forecast accuracy for items with low rate of sales, as defined by the Low Rate of Sales Value – override, is tested against this parameter to determine if an alert is triggered.

Low Rate of Sales Value - override

If the rate of sales is less than this value, the item is considered to have a low rate of sale. Its forecast accuracy is tested against the Low Forecast Error Value – override, to determine if an alert is triggered.

Medium Forecast Error Value - override

Forecast accuracy for items with medium rate of sales, as defined by the Medium Rate of Sales Value – override, is tested against this parameter to determine if an alert is triggered.

Medium Rate of Sales Value - override

If the rate of sales is less than this value, the item is considered to have a medium rate of sale. Its forecast accuracy is tested against the Medium Forecast Error Value – override, to determine if an alert is triggered. If the rate of sales is more than this value, the item is considered to have a high rate of sale. Its forecast accuracy is tested against the High Forecast Error Value – override, to determine if an alert is triggered.

12

Forecast Analysis Tools

RDF provides you with a number of tools that may be used for additional forecast analysis. The Forecast workbook group includes these workbooks:

- Interactive Forecasting Workbook
- Forecast Scorecard Workbook

Interactive Forecasting Workbook

The Interactive Forecasting workbook is a forecast simulation tool that allows you to make changes to forecast parameters and see the results without having to wait for the batch run. In this workbook, you can edit various forecast parameters including sales history and forecast method. A new forecast is produced based on the changed parameters. In addition to forecasts, the Interactive Forecasting workbook can also generate fit in historical region and the system picked model if an ES forecast method is used.

Opening the Interactive Forecasting Workbook

Use the following procedure to open the Interactive Forecasting workbook.

1. Within the Local domain, select **New** from the **File** menu.
2. Select the Forecast tab to display a list of workbooks.
3. Select Interactive Forecasting. Click **OK**.
4. Select the forecast level. Click **Next**.
5. Select the forecast starting date. Click **Next**.
6. Select the first date in the historical data to use for generating the forecast. Click **Next**.
7. Select the end date for the forecast horizon. Click **Next**.
8. Perform one of the following options:
   - Specify future time periods beyond the forecast horizon to include in your workbook.
   - Select **Do not include dates** after the horizon if no future dates are to be included in the workbook.
9. Click **Next.**
10. Select the products to be included in the workbook. Click **Next**.
11. Select the locations to be included in the workbook. Click **Next**.

12. Select extra measures (if needed) to be included in the forecasting workbook.Click **Finish**.

Once the wizard is completed, the forecast is generated based on your selections in the wizard and the Default Forecast Method set for the specified forecast level in the Forecast Administration Workbook.

After the forecast generation is complete, the Interactive Forecasting workbook is displayed.

Worksheets for the Interactive Forecasting Workbook

The following sections describe the components of the Interactive Forecasting workbook.

The Interactive Forecasting workbook includes these 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 for the select forecast level.

The Forecasting Parameter worksheet contains the following measures. Of these measures, all but System picked model are editable. Changes can be made to the editable parameters, that is, History Start Date, Forecast Method, Forecast Start Date, and Forecast End Date, to regenerate a forecast of the time series.

Figure 12–1 Forecast Parameter Worksheet

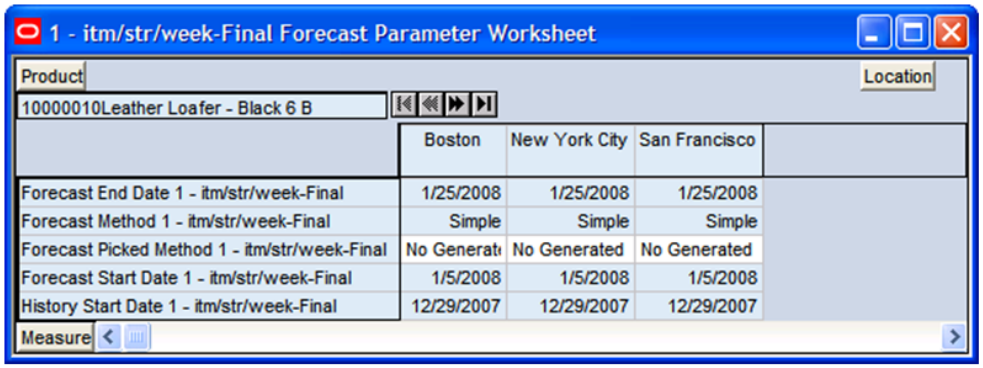

Measures: Forecasting Parameter Worksheet

The Forecasting Parameter worksheet contains the following measures:

Forecast Method

A list from which you can select the method used to generate the forecast. The workbook defaults to the method selected in the Forecast Administration Workbook for the specified level if the method is supported by Interactive Forecasting. Otherwise, it defaults to *No Forecast*.

Note: The Forecast Method list does not include Causal, Bayesian, Components, Copy, LoadPlan, or Profile-based forecast methods.

Forecast Picked Method

This is the method that was used to generate the forecast for the given product/location combination. This field is useful when combined methods are requested (for example, ES methods). Then this field displays the actual method the system picked from the combined methods. In case stand-alone methods are chosen, generally 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, the system defaults to a simpler method, and that method is displayed here.

History Start Date

This is 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 2011, you need to set your History Start Date to January 1, 2011. 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 , RDF 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, 2011, and an item/location does not have sales history until February 1, 2011, the system considers 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

This is first date of the forecast horizon. The default is the start date selected in the wizard.

Forecast End Date

This is the last date of the forecast horizon. The default is the end date selected in the wizard.

Additional Information about the Interactive Forecasting Workbook Parameters Interactive Forecasting assumes that the historical data is contained in the measure *posxlxb*. All forecasts generated by interactive forecasting are based on this data.

Performance may be improved by utilizing a mask measure. Interactive forecasting uses the measures `mask<forecast level>xb`; for example *mask01xb*. Here, `<forecast level>` corresponds to the forecast level chosen in the wizard when opening the Interactive Forecasting workbook. The mask is a boolean measure. Only cells with the value True are included in the forecast. This is useful when desiring to focus on only a few product/locations.

The Interactive Forecast method is intended to be a quick way to view results produced by different forecast methods. It allows the user to gauge, in a qualitative way, the suitability of a particular forecast method for a given set of sales history data. Even though it uses the exact same forecast methods as the full forecast, only the parameters listed previously will affect the results. No other user-input values are used. Users wishing to vary a larger number of parameters are encouraged to use the regular forecast functionality.

Also, the Winters Method has a pre-established minimum history of 104 periods.

Interactive Forecasting Worksheet

The Interactive Forecasting worksheet is based on the intersection of the Product, Location, and Calendar dimensions for the forecast level selected in the wizard process. The Interactive Forecasting worksheet contains the History Data and the System Forecast. Of these measures, only History Data is editable.

Note: The Interactive Forecast workbook calls the same Forecast Special expression that the Batch Forecast calls.

Interactive Forecast for group levels (example level 13) should be run, only after the corresponding Batch Forecast is run since a required measure is populated during the batch run.

Figure 12–2 Interactive Forecasting Worksheet

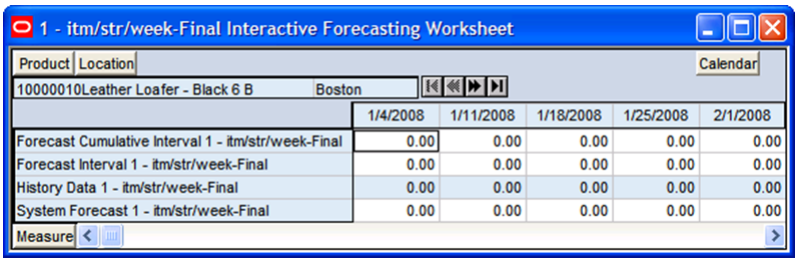

Measures: Interactive Forecasting Worksheet

The Interactive Forecasting worksheet contains the following measures:

History Data

This is the historical sales data set in the Forecast Administration Workbook. 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 adversely affects your forecast. However these changes are for simulation purposes only and cannot be committed.

System Forecast

The quantity that the system predicts is required for the product, location, and calendar combination displayed. Changes to History Data, History Start Date, Forecast Start Date, Forecast End Date, and Forecast Method causes the System Forecast to re-calculate when **Calculate** is selected from the toolbar.

Forecast Scorecard Workbook

This section describes the purpose and content of the Forecast Scorecard and the steps required in order to create and access this workbook. This section also discusses the 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 forecast period, statistical

information can be reviewed to help you analyze the accuracy of forecasting models and methods.

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.

Opening or Creating a Forecast Scorecard Workbook

Select **Open** from the File menu to bypass the Forecast Scorecard wizard. Or, perform the following:

1. Within the Master or Local domain, select **New** from the **File** menu.
2. Select the **Forecast** tab to display a list of workbooks.
3. Select **Forecast Scorecard**. Click **OK**. The Forecast Scorecard wizard opens.
4. You are prompted to select the forecast level that you want to evaluate. Make a selection. Click **Next**.

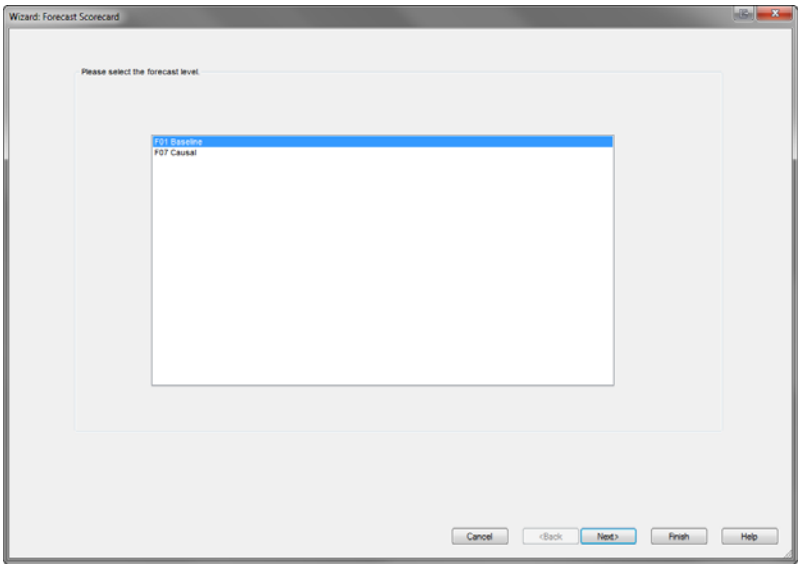

5. Select the time periods for the calculation of errors.

Note: You are responsible to select time periods that have available actuals and forecasts.

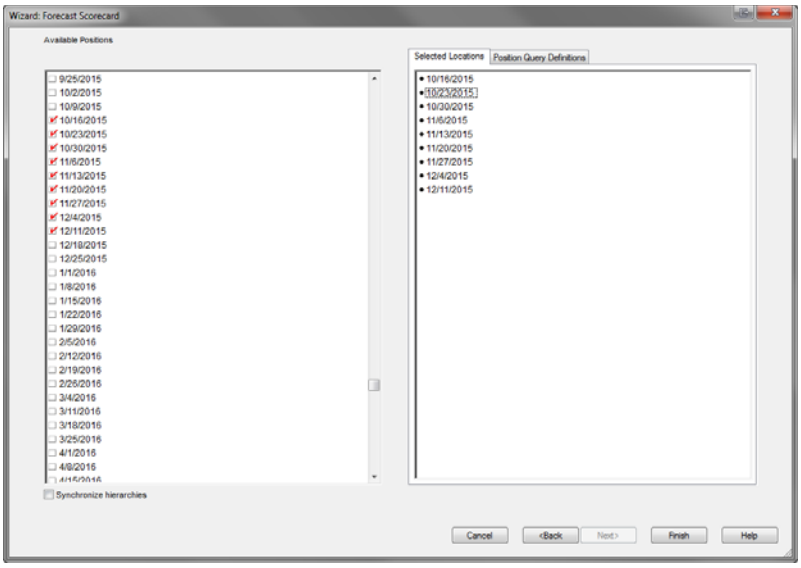

6. Click **Next**.

7. Select the specific locations that you want to view. It is important to include all locations that are members of the location dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at item/chain/week, you should include all locations that are members of the particular chain to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the location hierarchy is employed if the workbook supports an AutoWorkbook build. Make your selection. Click **Next**.

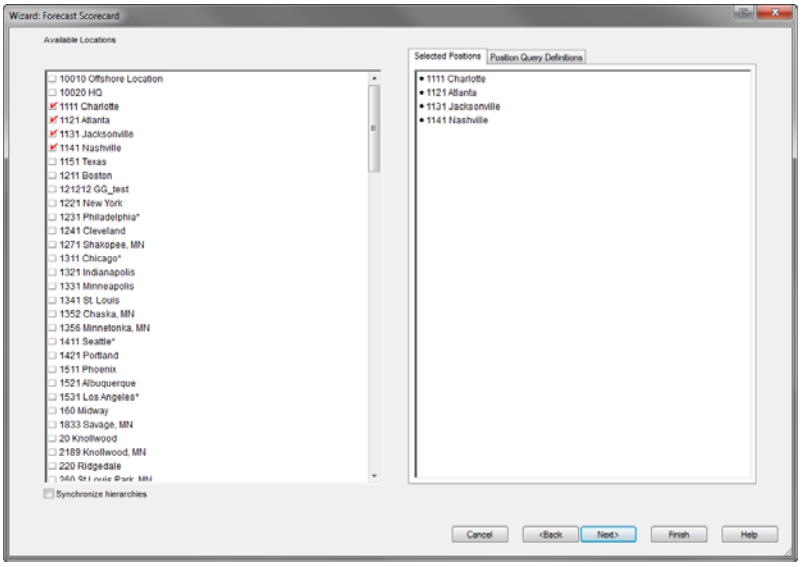

8. Select the merchandise you want to include.view. It is important to include all products that are members of the Merchandise dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at subclass/store/week, you must include all items that are members of the particular subclass to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the Merchandise hierarchy is employed if the workbook supports an AutoWorkbook build. Make your selection. Click **Next**.

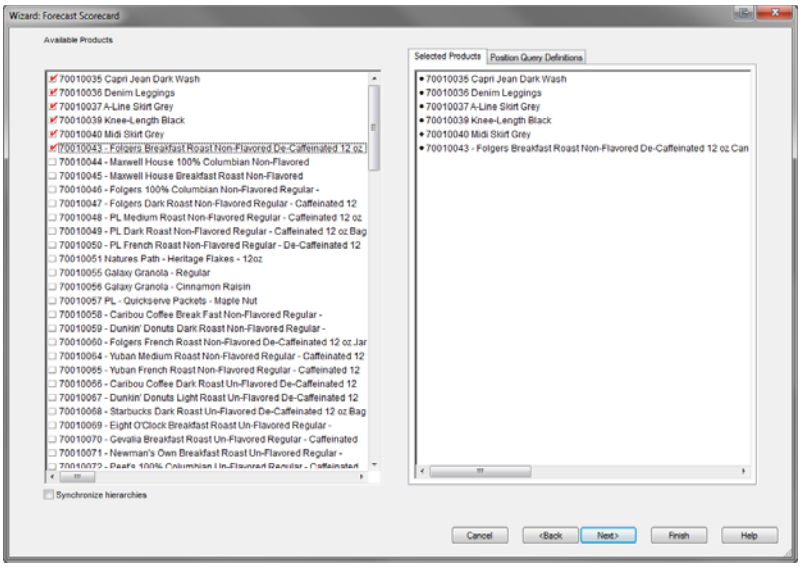

9. Click **Finish**.

The Forecast Scorecard workbook displays after the wizard process is completed.

Worksheets for the Forecast Scorecard Workbook

The following sections describe the components of the Forecast Scorecard workbook.

The Forecast Scorecard consists of a wizard and these worksheets:

- Forecast Errors - Low Level Worksheets
- Forecast Errors - Aggregate Level Worksheets
- Actuals versus Forecasts Worksheet

Forecast Errors - Low Level Worksheets

There are two Forecast Error worksheets at the Low Level:

- Approved Forecast Errors - Low Level
- Approved System Forecast Errors - Low Level

These worksheets are intended to report on the errors at the final forecast intersection. Ultimately, this is the level where the user wants the accuracy criteria to be met.

> **Note:** All error statistics are calculated for the periods, selected in the wizard.

Measures: Approved Forecast Errors - Low Level and Approved System Forecast Errors - Low Level Worksheet

The Forecast Errors - Low Level worksheets contain the following measures:

History Data for Forecast Window

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.

Forecast Measures

The forecast errors are calculated between the historical demand, and two forecast metrics: the approved forecast and the approved system forecast. The errors between historical demand and each forecast measure are displayed in two worksheets which can be arranged side by side. The error of the approved forecast takes into account adjustments of the forecast user, while the error of the approved system forecast truly reflects the quality of the forecasts generated by the system, without human intervention.

Mean Error

The error of a forecast observation is the difference between the forecast value and the actual POS value. The Mean Error statistic is a measure of the average error over time. This is calculated by summing the errors for all observations and then 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 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.

Mean Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecast value and the actual POS value. The Mean Absolute Error statistic is a measure of the average absolute error. This is calculated by summing the absolute errors for all observations and then 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.

Root Mean Squared Error

This is 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 it is best used to compare alternative forecasting models for a given series.

Mean Absolute Percentage Error

The percentage error of a forecast observation is the difference between the actual POS value and the forecast 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).

Percentage Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecast 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.

Forecast Errors - Aggregate Level Worksheets

There are two Forecast Error worksheets at the Aggregate Level:

- Approved Forecast Errors - Aggregate Level
- Approved System Forecast Errors - Aggregate Level

In these worksheets, although the data is still at the low level, it can be rolled up to any intersection, and the error metrics are recalculated. This information can be useful in several scenarios, For example, if you decide that the forecast accuracy at the low level is inadequate, then you can roll up to different intersections, and check at what level, the forecast accuracy starts to deteriorate. Or the rolled-up level can be preset, and you first check the aggregate level accuracy. You can expect that the low level accuracy will be at most on par with the accuracy at the aggregate level.

> **Note:** All error statistics are calculated for the periods, selected in the wizard.

Measures: Approved Forecast Errors - Aggregate Level and Approved Forecast Errors - Aggregate Level Worksheet

The Forecast Errors - Aggregate Level worksheets contains the following measures:

History Data for Forecast Window

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.

Forecast Measures

The forecast errors are calculated between the historical demand, and two forecast metrics: the approved forecast and the approved system forecast. The errors between historical demand and each forecast measure are displayed in two views which can be arranged side by side. The error of the approved forecast takes into account adjustments of the forecast user, while the error of the approved system forecast truly reflects the quality of the forecasts generated by the system, without human intervention.

Mean Error

The error of a forecast observation is the difference between the forecast value and the actual POS value. The Mean Error statistic is a measure of the average error over time. This is calculated by summing the errors for all observations and then 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 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.

Mean Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecast value and the actual POS value. The Mean Absolute Error

statistic is a measure of the average absolute error. This is calculated by summing the absolute errors for all observations and then 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.

Root Mean Squared Error

This is 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 it is best used to compare alternative forecasting models for a given series.

Mean Absolute Percentage Error

The percentage error of a forecast observation is the difference between the actual POS value and the forecast 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).

Percentage Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecast 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.

Actuals versus 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 historical demand quantities.

Figure 12–3 Actuals versus Forecast Worksheet

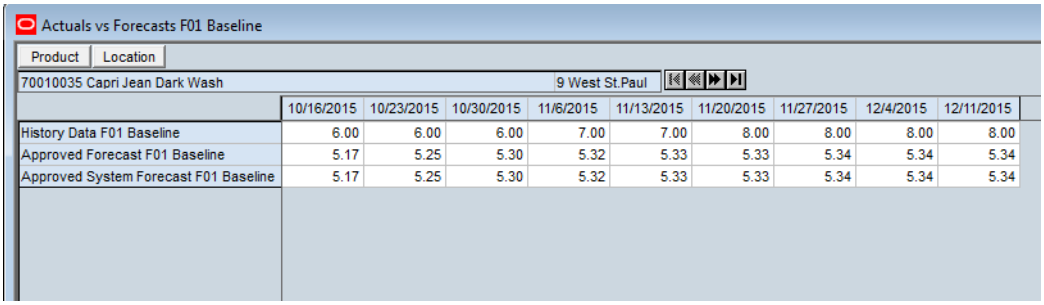

	10/16/2015	10/23/2015	10/30/2015	11/6/2015	11/13/2015	11/20/2015	11/27/2015	12/4/2015	12/11/2015
History Data F01 Baseline	6.00	6.00	6.00	7.00	7.00	8.00	8.00	8.00	8.00
Approved Forecast F01 Baseline	5.17	5.25	5.30	5.32	5.33	5.33	5.34	5.34	5.34
Approved System Forecast F01 Baseline	5.17	5.25	5.30	5.32	5.33	5.33	5.34	5.34	5.34

The Actuals versus Forecasts worksheet contains the following measures:

History Data

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

Approved Forecast

This is the reviewed/adjusted forecast approved by the user.

Approved System Forecast

This forecast is wholly system generated.

13

Promote (Promotional Forecasting)

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 it outlines the use of Oracle Retail's Causal method of forecasting demand. It describes terminology used in the context of promotional forecasting, and it concludes with detailed descriptions of the workbooks and worksheets contained in the Promote workbook group.

About 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.

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, and so on 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. There often 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 influence spikes and dips of past demand. 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 Oracle Retail'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 (1 indicates 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 Oracle Retail 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.

Promotional Forecasting Approach

Oracle Retail combines time series forecast methods with causal forecast methods, resulting in a new forecast method supported by the promotional forecasting 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 Promotional Forecasting 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, Promotional Forecasting provides a cost-effective means for producing forecasts using promotional information.

Promotional Forecasting Terminology and Workflow

Promotional Forecasting 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. 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.

Note: More detailed information on the Causal forecasting algorithm is provided in the Forecasting Methods Available in RDF section.

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 13–1 Promotion Event 1

Christmas Day applies to all products/locations. The Christmas promotional variable therefore has only one dimension, Day." Because Christmas Day falls on the 359th day of each non-leap year, the Day359 variable is set to True for every such year (all other days it is set to False).

Example 13–2 Promotion Event 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 has two associated dimensions, State and Week. Week13, Week14, Week15, and Week16 is set to True only for states NY and NJ (all other states/weeks is set to False). Since no product dimension exists, the TV ad is assumed to have an effect on all products.

Example 13–3 Promotion Event 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 - item and Week. For this year, Week23 and Week24 is set to True only for individual items related to sporting goods items (all other weeks and all other items is 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 RDF, you can use the Promo 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. In addition, you can perform What-if, evaluate the results and commit back to the database the desired settings.

Note: External to Promote, the Forecast Administration Workbook includes several parameters that may be used as additional configuration options for promotional forecasting. For additional information, see Procedures in Promotional Forecasting.

The Promote workbooks are listed in Table 13–1.

Table 13–1 Promote Workbooks

Workbooks	Description	Worksheets
Promo Planner Workbook	Allows you to specify when certain promotional events were active in the past, and when they are active in the future.	Promotion Worksheet
Promo Effect Maintenance Workbook	Allows you to review the system-calculated promotional lift effects, edit these effects, and determine how changes are factored into the promotional model.	Final Promotion Effect Parameters Worksheet
Promotion Effectiveness Workbook	In addition to all tasks that can be performed in the other Promo workbooks, the user can perform what-if, evaluate the promotional forecast and commit the parameters.	Default Causal Parameters Worksheet Final Level Causal Parameters Worksheet Pooling Levels Causal Parameters Worksheet Causal Forecast Worksheet Promotion Model Type Worksheet Promotion Calendar Worksheet
Promotion Management Workbook	Allows you to enable and disable promotions for causal forecast levels and to specify whether promotions can have a negative effect.	Promo Variable Model Type Worksheet Promotion Enable Worksheet Accept Negative Lift Worksheet

Promo Planner Workbook

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 Promo Planner workbook allows you to indicate to the system when certain events were active in the past and when they is 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 Promo 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 may be loaded at the item/store/week intersection, while an event such as Christmas is loaded at the [Day] level.

In this setup, the Advertisement and Gift with Purchase promotions would appear on one worksheet, and Christmas would appear on another. Whenever a hierarchy is not included in the base intersection (as in the case of the Christmas promotional event) the event is assumed to apply to all positions in the undefined hierarchy. Thus, Christmas is assumed to apply to all products and all locations, but only to the Day-level calendar positions specified in the Promo Planner Workbook.

Opening or Creating a Promo Planner Workbook

The Promo Planner wizard steps you through the process of creating a new Promo Planner workbook. To access the Promo Planner, select Open from the File menu to

bypass the wizard and open an existing Promo Planner workbook, or perform the following steps:

1. Within the Local domain, select **New** from the **File** menu.
2. On the Promote tab, select **Promo Planner**. Click **OK**.
3. The Promo Planner wizard opens and prompts you to select the promotion events to edit or review. Click **Next**.

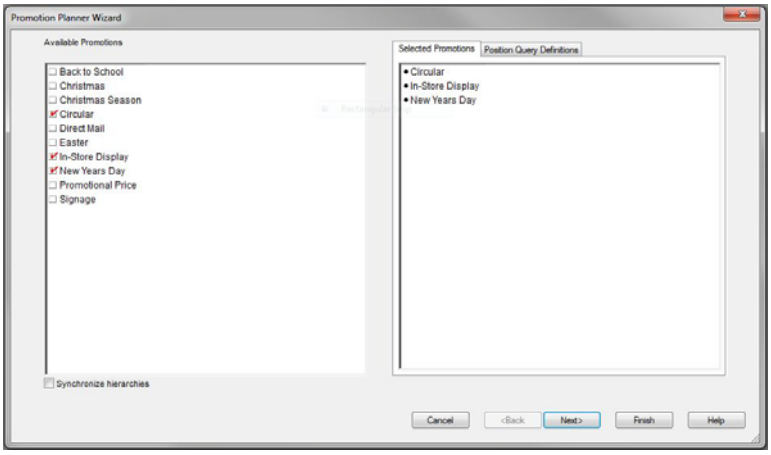

4. Select the products that need to have promotions planned. Click **Next**.

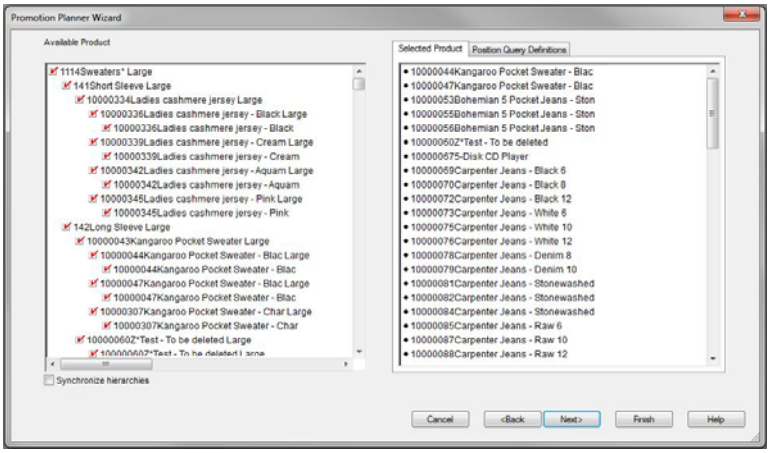

5. Select the locations that need to have promotions planned. Click **Next**.

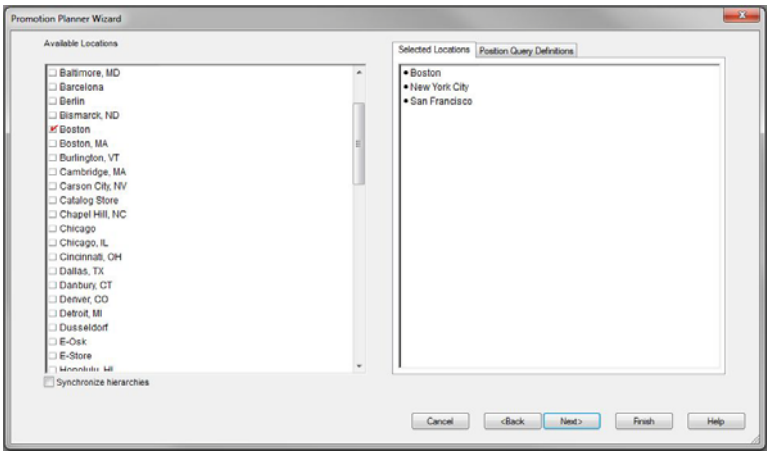

6. Select the dates that need to have promotions planned. Click **Next**.

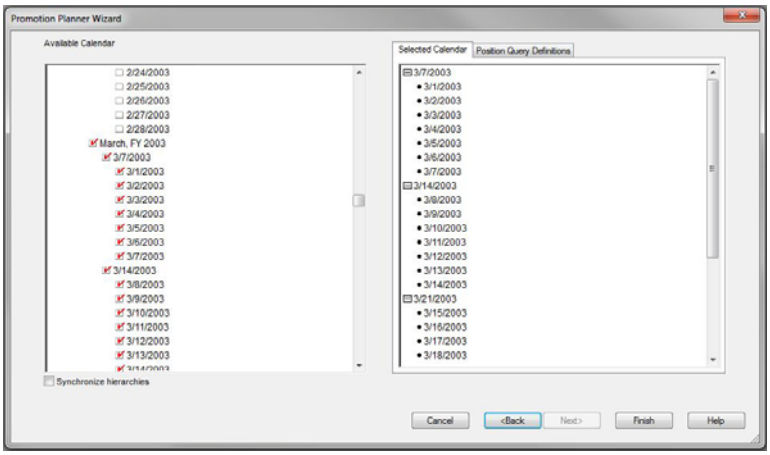

7. Click **Finish** to build the workbook.

Worksheets for the Promo Planner Workbook

The following sections describe the components of the Promo Planner workbook.

The Promo Planner workbook includes this worksheet:

- Promotion Worksheet

Promotion Worksheet

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

Figure 13–1 Promotion Worksheet

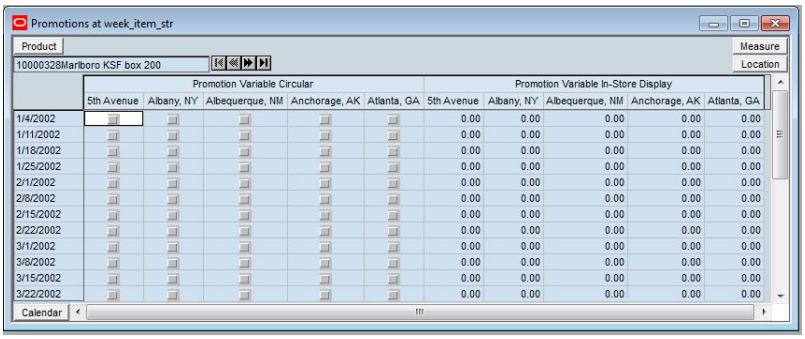

Measures: Promotion Worksheet

The Promotion worksheet contains the following measures:

Promotion Variables

Promotion Variable Type is defined through setting both Data Type and Model Type during the configuration process in the Promote Plug-in. The Promotion Variable Type includes Boolean, Exponential, and Real. The following table explains how Promotion Variable Type is defined through Data Type and Model Type settings.

Table 13–2 Promotion Variable Types

When the Type is	And the Model is	Then the Promotion Variable Type is
Boolean	Linear	Boolean
Real	Linear	Real
Real	Exponential	Exponential

The Data Type includes Boolean and Real types. The Model Type includes Linear and Exponential types. When Data Type is set to Boolean, the Promotion Variable Type is Boolean. When Data Type is set to Real and the Model Type is set to Linear, the Promotion Variable Type is Real. If Data Type is set to Real and the Model Type is set to Exponential, the Promotion Variable Type is Exponential.

Refer to the *Oracle Retail Demand Forecasting Configuration Guide* for additional details on configuring the Promote solution.

Promotion Variable Type	Description
Boolean	A check in a given cell indicates that the associated promotion event's status is on (or 100% of the lift effect applies) for that intersection. If no check is indicated, the event's status is off. Boolean promotions can be enabled together with other types of promotions at the same time.

Promotion Variable Type	Description
Exponential	Real numbers are expected in the cells. The value specified here is used in an exponential model which is an input to the stepwise regression sub-routine. A typical example of a metric modeled as exponential is the percent price change. Note how the percent price change can be positive (price increase) or negative (price decrease). Exponential promotions can be enabled together with Boolean promotions at the same time. However, they can not be enabled together with Real promotions.
Real	Non-negative values are expected in the cells. A value of 0 indicates that the event's status is off. A value greater than 0 will act as a weight when calculating the effects during the stepwise regression. Real promotions can be enabled together with Boolean promotions at the same time. However, they can not be enabled together with Exponential promotions.

About Implementation

Among the ways Causal variables can be implemented include price or discount percent. Your Oracle Retail Consultant can best determine the most accurate setup of promotion variables based upon your promotional forecasting requirements.

Promo Effect Maintenance Workbook

The Promo Effect Maintenance workbook 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 are factored into the promotional models. The Promo Effect Maintenance workbook contains one worksheet. There may be multiple versions of this worksheet, defined at various causal levels.

Opening the Promo Effect Maintenance Workbook

The Promo Effect Maintenance wizard steps you through the process of creating a new Promo Effect Maintenance workbook.

1. Within the Local domain, select **New** from the **File** menu.
2. On the Promote tab, select **Promo Effect Maintenance**. Click **OK**.
3. Select the promotion events to analyze. Click **Next**.
4. Select the causal forecast level for analysis. Click **Next**.

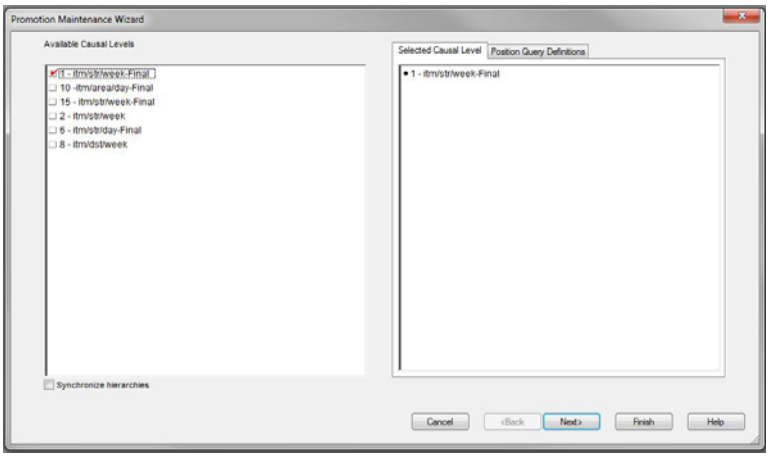

5. Select the products to analyze. Click **Next**.

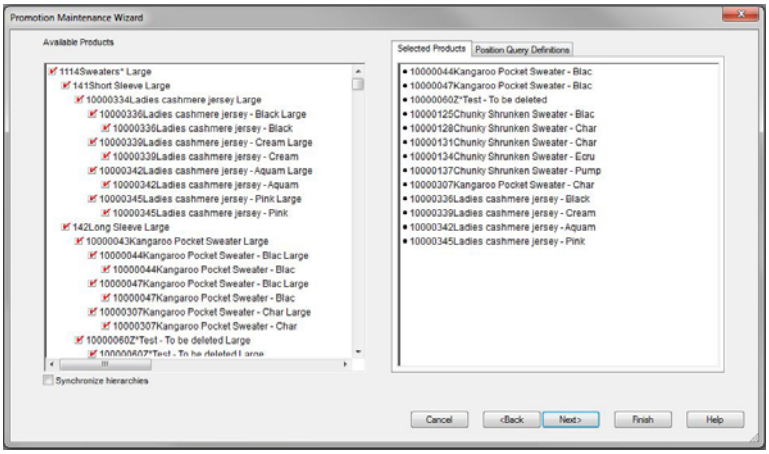

6. Select the locations to analyze. Click **Next**.

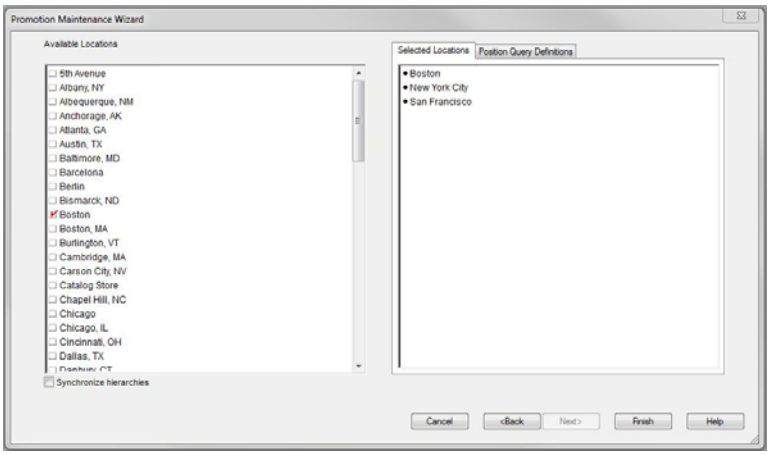

7. Select additional measures to view in the workbook (if necessary). Click **Finish**.

Worksheets for the Promo Effect Maintenance Workbook

The following sections describe the components of the Promo Effect Maintenance workbook.

The Promo Effect Maintenance workbook includes this worksheet

- Promotion Effect Parameters Worksheet
- Final Promotion Effect Parameters Worksheet

Promotion Effect Parameters Worksheet

The Promotion Effect Parameters worksheet allows you to worksheet and modify the system-calculated effects of a given promotion at the causal source level.There is one worksheet for every causal source level.

> **Note:** Changes that are committed in the Promo Effect Maintenance workbook are not incorporated into the forecast until another forecast is generated in RDF batch forecast process.

Measures: Promotion Effect Parameters Worksheet

The Promotion Effect Parameters worksheet contains the following measures:

Promo Effect 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. Table 13–4 lists the options.

System Calculated Effect

The System Calculated Effect is a read-only measure indicating the lift effect generated by the system.

System Effect Override

The user-specified lift effect. This user-entered effect is active if used in conjunction with the Override All and Override Future Only Promotion Effect Types. Otherwise, it equals 1.00.

Table 13–3 Promo Effect Types

Causal Variable Type	Description
Automatic	The inclusion of the Promo Effect is decided by regression. If the Promo Effect is found to be significant on the training set, it is included in the model. Otherwise, it is rejected. Automatic is the system default Promotion Effect Type.
Forced In	The Promo Effect 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 even if it has a negative impact to the demand forecast.
Disabled	The variable is excluded from the model, hence no effect is returned either.
Override All	This type allows you to specify a causal effect that is used during the forecasting process. For Promo Effects specified as Override All, you also specify the corresponding causal effect in the Promotion Effect Override. The causal engine de-causalizes the training data using the user-specified effect. During forecast generation, the user-specified effect is used to determine the causal forecast. Therefore, you must change the Promotion Effect Type when this user-specified effect is no longer to be used.

Table 13–3 (Cont.) Promo Effect Types

Causal Variable Type	Description
Override Future Only	This type allows you to specify a causal effect that is used during the forecasting process. For Promo Effects specified as Override Future, you also specify the corresponding causal effect in the Promotion Effect Override. The calculated effect is not written back to the effects measure, but it is used to de-causalize the data. During forecast generation the calculated effect is ignored, and, instead, the user-specified effect is used to produce a causal forecast. Therefore, you must change the Promotion Effect Type when this user-specified effect is no longer to be used.
Override from Higher Level	This promotion type is used in conjunction with the Causal Higher Intersection set in the Forecast Administration Workbook. If the Causal Higher Intersection is not specified, no promotional effect is calculated. It allows the system to use the geometric mean of causal effects computed from product/location combination in the same group (the intersection level specified in the Causal Higher Intersection during system setup) for product/location combinations for which from sales history alone a causal effect is unable to be computed for that Promo Effect. The way the system handles this, is by having the override effects measure filled in with higher-level effects for those variables specified as Override Higher Level. During forecast generation the effect calculated as average of low level effects is used to produce a causal forecast. Note that if no effects are found for any of the product/locations in the Causal Higher intersection group, then the system uses the geometric mean of causal effects calculated from all Product/Locations as the override effect being passed in.
Automatic Boolean	If the promotion variable is always be set to 0.00 or 1.00 (meaning inactive or active, respectively), this type produces the same results as the Automatic type; however, Automatic Boolean improves the performance (speed) of the forecasting engine during the batch run of the forecast.

Final Promotion Effect Parameters Worksheet

The Final Promotion Effect Parameters worksheet allows you to view and modify the system-calculated effects of a given promotion.

Figure 13–2 Final Promotion Effect Parameters Worksheet

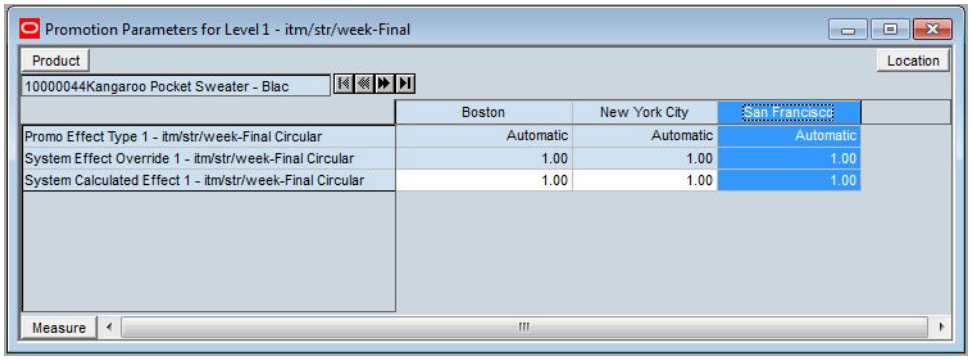

Measures: Final Promotion Effect Parameters Worksheet

The Final Promotion Effect Parameters worksheet contains the following measures:

System Calculated Effect

The System Calculated Effect is a read-only measure indicating the lift effect generated by the system.

System Effect Override

The user-specified lift effect. This user-entered effect is active if used in conjunction with the Override All and Override Future Only Promotion Effect Types. Otherwise, it equals 1.00.

Promo Effect 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. Table 13–4 lists the options.

Table 13–4 Promo Effect Types

Causal Variable Type	Description
Automatic	The inclusion of the Promo Effect is decided by regression. If the Promo Effect is found to be significant on the training set, it is included in the model. Otherwise, it is rejected. Automatic is the system default Promotion Effect Type.
Forced In	The Promo Effect 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 even if it has a negative impact to the demand forecast.
Disabled	The variable is excluded from the model, hence no effect is returned either.
Override All	This type allows you to specify a causal effect that is used during the forecasting process. For Promo Effects specified as Override All, you also specify the corresponding causal effect in the Promotion Effect Override. The causal engine de-causalizes the training data using the user-specified effect. During forecast generation, the user-specified effect is used to determine the causal forecast. Therefore, you must change the Promotion Effect Type when this user-specified effect is no longer to be used.
Override Future Only	This type allows you to specify a causal effect that is used during the forecasting process. For Promo Effects specified as Override Future, you also specify the corresponding causal effect in the Promotion Effect Override. The calculated effect is not written back to the effects measure, but it is used to de-causalize the data. During forecast generation the calculated effect is ignored, and, instead, the user-specified effect is used to produce a causal forecast. Therefore, you must change the Promotion Effect Type when this user-specified effect is no longer to be used.
Override from Higher Level	This promotion type is used in conjunction with the Causal Higher Intersection set in the Forecast Administration Workbook. If the Causal Higher Intersection is not specified, no promotional effect is calculated. It allows the system to use the geometric mean of causal effects computed from product/location combination in the same group (the intersection level specified in the Causal Higher Intersection during system setup) for product/location combinations for which from sales history alone a causal effect is unable to be computed for that Promo Effect. The way the system handles this is by having the override effects measure filled in with higher-level effects for those variables specified as Override Higher Level. During forecast generation the effect calculated as average of low level effects is used to produce a causal forecast. Note that if no effects are found for any of the product/locations in the Causal Higher intersection group, then the system uses the geometric mean of causal effects calculated from all Product/Locations as the override effect being passed in.
Automatic Boolean	If the promotion variable is always be set to 0.00 or 1.00 (meaning inactive or active, respectively), this type produces the same results as the Automatic type; however, Automatic Boolean improves the performance (speed) of the forecasting engine during the batch run of the forecast.

Note: Changes that are committed in the Promo Effect Maintenance workbook are not incorporated into the forecast until another forecast is generated in RDF batch forecast process.

Promotion Effectiveness Workbook

In this task you can perform every step necessary to plan promotions. You can determine which events should be active, the value of the price discount, review and possibly overwrite causal effects, and run simulations to evaluate the plan.

The simulations are a very useful tool to determine how effective a promotion is. For instance, an item with a price discount of 25% and advertised in a flyer may yield a lift of 150%. The same item, with the same price discount, advertised in a flyer, and on radio, yields a lift of 155%. The user has the information to decide if the cost of the radio ad is worth the extra 5% increase in sales.

Once you decide on the setup for the promotion, based on the what-if simulations, the information can be committed and the next forecast batch generates the baseline and lifts that are exported to downstream applications, such as replenishment solutions.

Opening the Promotion Effectiveness Workbook

The Promotion Effectiveness wizard steps you through the process of creating a new Promotion Effectiveness workbook.

1. Within the Local domain, select **New** from the **File** menu.
2. On the Promote tab, select Promotion Effectiveness. Click **OK**.
3. Select the causal level. Click **Next**.

Note: This wizard only displays the causal final levels for selection.

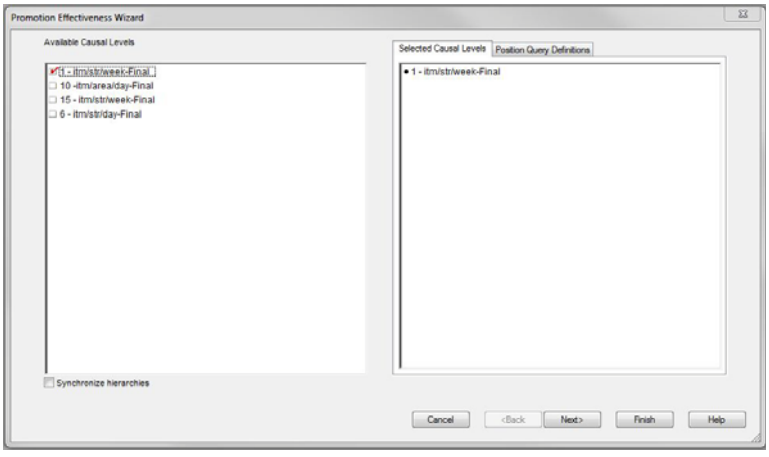

4. Select the products to analyze. Click **Next**.

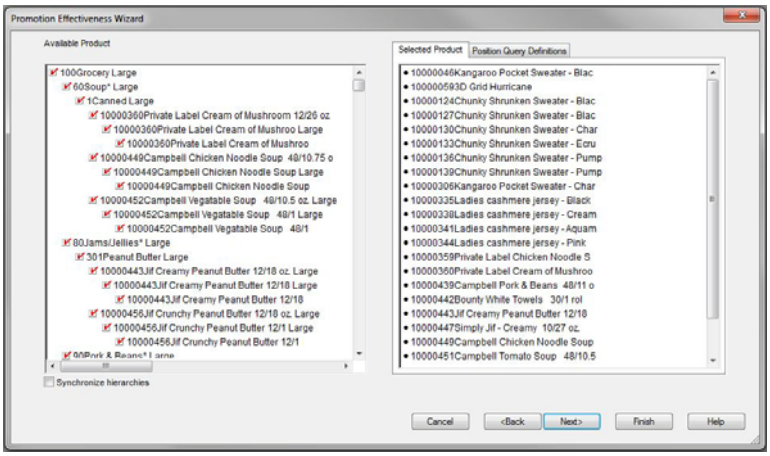

5. Select the locations to analyze. Click **Next**.

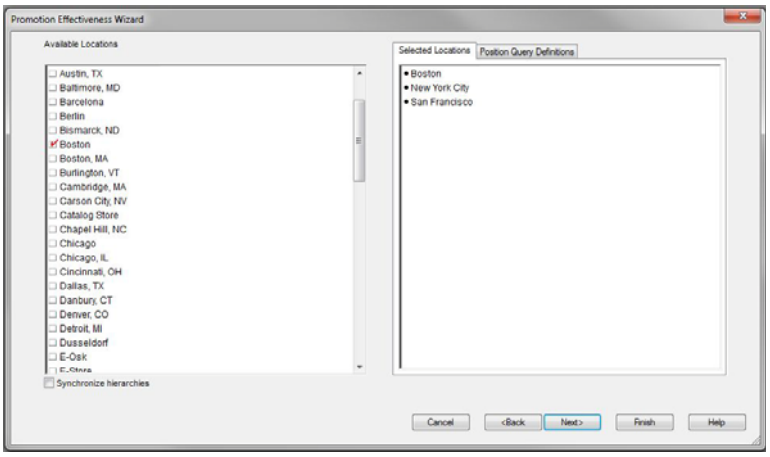

6. Select the dates to analyze. Click **Next**.

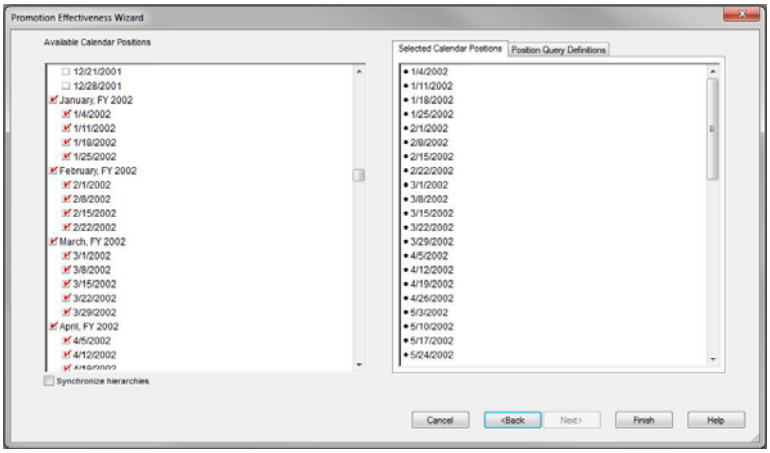

7. Select any additional measures that need to be included in the workbook (if necessary). Click **Finish**.

This workbook has two tabs:

- Promotional Parameters Tab
- Promotional Forecast Tab

Promotional Parameters Tab

The Promotional Parameters tab has several worksheets:

- Default Causal Parameters Worksheet
- Final Level Causal Parameters Worksheet
- Pooling Levels Causal Parameters Worksheet

In this tab, you can review all the information related to promotion effects. For example, you can check the calculated effects at the final as well as pooling levels. You can decide if the calculated effects need to be adjusted, and how to combine final level and pooling level effects to create a robust causal forecast.

The worksheets in this Promotional Parameters tab give you the tools to build the blocks necessary to build a robust and accurate causal forecast.

- Good item: blend param -> 0
- Bad item: blend param -> 1
- New item: get pooling
- Item not new, but never promoted: get pooling effect

Default Causal Parameters Worksheet

The Default Causal Parameters worksheet contains the following measures.

Default Blending Parameter

This parameter sets the weights for combining the Final and Source Level promotion effects, when calculating the blended effect. The range of the parameter is 0 to 1. A value closer to 1 will yield a blended effect closer to the Pooling Level effect. You pick a high value if you want t robust causal forecast, although the final level promotion information is not very reliable. A value closer to zero yields an effect closer to the Final Level effect. You pick a low value if the final level promotion information is accurate and you want the causal forecast to reflect item-specific effects.

The measure can also be edited in the Forecast Administration workbook, and can be overridden in the Forecast Maintenance workbook.

Default Source Level

The list of values displayed in this field allows you to change the pooling level that is used for calculating the blended effect. The pooling levels are set up in the RPAS Configuration Tools. A value from the pick list is required in this field at the time of forecast generation.

The measure can also be edited in the Forecast Administration workbook, and can be overridden in the Forecast Maintenance workbook.

Default Overlapping Promotions Adjustment Factors

This Default Overlapping Promotion Adjustment Factor specifies at a high level how the individually calculated promotions interact with each other when they are overlapping in the forecast horizon. This parameter serves as a global setting, but can be overridden at lower levels. The default value is 1.

- A value greater than 1 means the promotion effects will be compressed when applied in the model, instead of linearly summing up to get the total promotion effect. The larger the value is, the larger the compression effect will be, meaning the smaller the total effect will be.

- A value between 0 and 1 means the promotion effects will be amplified when applied in the model, instead of linearly summing up to get the total promotion effect. The smaller the value is, the larger the amplification effect will be, meaning the larger the total effect will be.
- If a value less than or equal to 0 is put in the cell, the calculation engine will find the best adjustment factor to fit the history data. This factor is then used to combine overlapping promotions in the forecast horizon. With few exceptions, the value should be anywhere from 1 to 5. This factor is then used to combine overlapping promotions in the forecast horizon.
- The measure can also be edited in the Forecast Administration workbook.

Caution: The process to find the best adjustment factor can be time consuming and introduce performance issues.

Final Level Causal Parameters Worksheet

This worksheet includes all the information available in the Final Level Causal Effect Parameters in the Promo Effect Maintenance workbook.

Additionally there are a few measures that give total flexibility to you to generate the causal forecast you want.

All of these measures are described and are available for editing in the Forecast Administration workbook.

Pooling Levels Causal Parameters Worksheet

This worksheet includes all the information available in the Pooling Level Causal Effect Parameters in the Promo Effect Maintenance workbook

Promotional Forecast Tab

The Promotional Forecast tab has several worksheets:

- Causal Forecast Worksheet
- Promotion Model Type Worksheet
- Promotion Calendar Worksheet

Causal Forecast Worksheet

The Causal Forecast worksheet allows you visibility to the baseline and promotion lifts components of the forecast. This is the place you review the impact of the promo-related parameters set in the Promotions and Promotion Parameters worksheets. The process of changing parameters, re-generating the promo forecast, and reviewing the impact in this worksheet, implements the promotional What-if functionality.

The Causal Forecast worksheet contains the following measures:

Forecasted Baseline

The Forecast Baseline is a forecast generated on past sales data that contains no promotions (that is, normal demand given no causal effects).

Promotion Lift Units

This measure displays the total effect that promotions can have on an item's sales units. It includes the effects of all the promotions active for that time period, as well as the price effect.

Total Promotional Forecast

This measure displays the sum of the baseline and the promotion lift.

Custom Menu: Calculate Forecast

The Promotional Forecast tab has a custom Menu which implements the forecast simulation (what-if). First, it calculates the promotional lifts based on the effects set in the Promotional Parameters tab. Then it applies the effects on top of the baseline to create the Total Promotional Forecast.

Promotion Model Type Worksheet

The Promotion Model Type worksheet contains the following measure:

Model Type for Promotional Variable

This measure stores the model type for the selected promotion variables. The choices are exponential and linear. An exponential model type is displayed if the promotion type is exponential. A linear model type is displayed when the promotion types are Real or Boolean.

Promotion Calendar Worksheet

The Promotions worksheet allows you to review and enable or disable promotions for item/locations for desired periods. Once the changes in the promo calendar have been made, click Calculate to update the forecast values in the Promotional Forecast worksheet.

The promotions must be enabled for the chosen causal level.Note that multiple Promotions worksheets are possible. Namely, there is one worksheet for every distinct promotion intersection. Promotions with identical intersections are displayed in the same worksheet.

The Promotion Calendar worksheet contains the following measure.

Promotion Variables

Promotion Variables are defined as either Boolean, Exponential, or Real types during the configuration process in the Promote Plug-In. Refer to the *Oracle Retail Demand Forecasting Configuration Guide* for additional details on configuring the Promote solution.

Promotion Management Workbook

The Promotion Management workbook allows you to enable and disable promotions for causal forecast levels and to specify whether promotions can have a negative effect.

Opening the Promotion Management Workbook

The Promotion Management wizard helps you create a new Promotion Management workbook.

1. Within the Local domain, select **New** from the **File** menu.

2. On the Promote tab, select **Promotion Management**.

   Click **OK**.

3. Select the causal level. Click **Next**.

4. Select the promotions. Click **Finish**.

Worksheets for the Promotion Management Workbook

The following sections describe the components of the Promotion Management workbook.

The Promotion Management workbook includes these worksheets:

- Promo Variable Model Type Worksheet
- Promotion Enable Worksheet
- Accept Negative Lift Worksheet

Promo Variable Model Type Worksheet

The Promo Variable Model Type worksheet allows you to view the model type for the selected promotion variables.

Figure 13–3 Promo Variable Model Type Worksheet

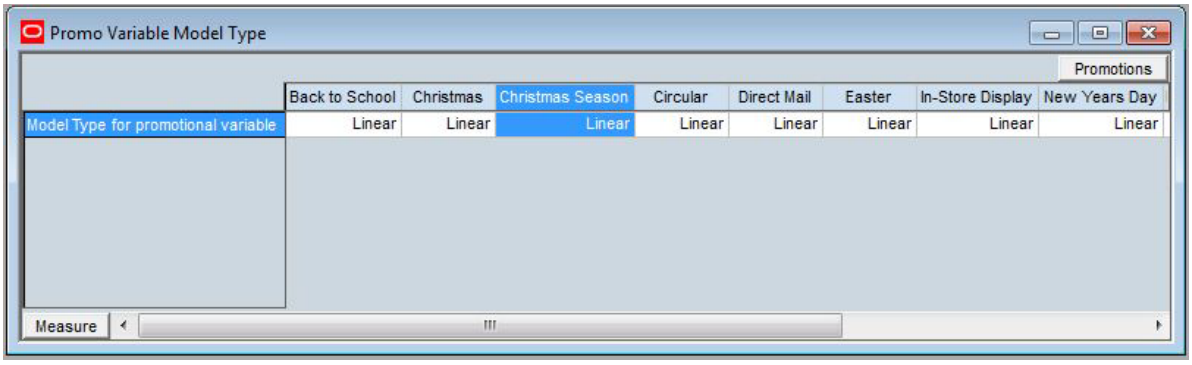

Measure: Promo Variable Model Type Worksheet The Promo Variable Model Type worksheet contains the following measure:

Model Type for Promotional Variable

This measure stores the model type for the selected promotion variables. The choices are exponential and linear. An exponential model type is displayed if the promotion type is exponential. A linear model type is displayed when the promotion types are Real or Boolean.

For additional information, see "Promotion Variables" on page 13-8.

Promotion Enable Worksheet

The Promotion Enable worksheet allows you to enable or disable a promotion for the causal forecast levels. The Promotion Enable worksheet is built at the hierarchy intersections of the promotion variables and the causal levels selected during the wizard process.

Figure 13–4 Promotion Enable Worksheet

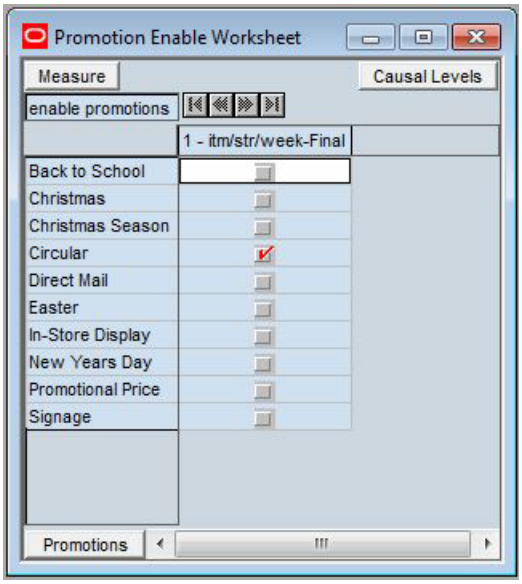

Note: Exponential promotions and Real promotions can not be enabled at the same time for the same level. In another word, Real promotions and Exponential promotions are not valid combinations. The valid promotion combinations are Boolean and Exponential promotions, as well as Boolean and Real promotions.

Each promotion's type is defined in the Configuration Tool. For additional information, see "Promotion Variables" on page 13-8.

Measures: Promotion Enable Worksheet

The Promotion Enable worksheet contains the following measure:

Enable Promotions

This is a read/write Boolean measure. The Enable Promotions measure allows you to enable a subset of promotions for a certain causal forecast level. It defaults to True for all causal forecast levels.

Accept Negative Lift Worksheet

The Accept Negative Lift worksheet allows you to specify whether a promotion is allowed to have negative effect. The Promotion Enable worksheet is built at the hierarchy intersections of the promotion variables selected during the wizard process.

Figure 13–5 Accept Negative Lift Worksheet

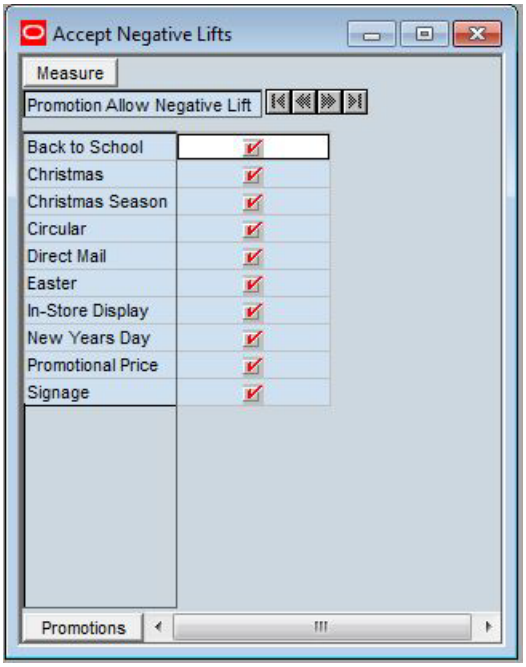

Measures: Accept Negative Lift Worksheet

The Accept Negative Lift worksheet contains the following measure:

Promotion Allow Negative Lift

This is a read/write Boolean measure. The Promotion Allow Negative Lift measure allows you to specify whether a promotion is allowed to have negative effect. When this measure is set to False for a promotion, the forecast engine does not allow that promotion to have any negative lift (Effect < 1). It defaults to True for all promotions.

Procedures in Promotional Forecasting

The following list outlines the standard procedures performed in order to set up the system to run a promotional forecast (more detailed steps are described in the next sections):

- Set up the system to run a promotional forecast
- Set forecast parameters in the Forecast Administration Workbook
- Set forecast parameters in the Forecast Maintenance Workbook
- Set promotions to be active in the Promo Planner Workbook
- Run the batch forecast
- View and Edit Causal Forecast results
- Analyze forecasts in the Forecast Approval Workbook
- Analyze and edit causal effects in the Promo Effect Maintenance Workbook
- Promotion Simulation (What-if) and Analysis

Setting Up the System to Run a Promotional Forecast

Perform the following procedure to set up the system to run a promotional forecast.

1. On the Forecast tab, select and build a Forecast Administration Workbook.
    a. In the Forecast Administration Workbook: Final Level Parameters Worksheet - Advanced Settings, select one of these options.
    b. Optional: **Set the Default Forecast Method** to **Causal** for the desired level if the level is to be use only for Promotional Forecasting.
    c. Optional: **Set the Causal Higher Intersection** for the desired level if the Override from Higher Level promotion type is used.
    d. Optional for use with Daily Causal Forecasting: Set the values for the following parameters:

        Causal Aggregation Profile

        Causal Calculation Intersection

        Causal Calculation Intersection Periodicity
    e. Commit your changes to the master database by selecting **Commit Now** from the File menu.
2. On the Forecast tab, select and build a Forecast Maintenance Workbook.
    - Set the **Forecast Method Override** to **Causal** for any items/locations at the desired levels that use Promotional Forecasting.
    - Commit your changes to the master database by selecting **Commit Now** from the File menu.
3. On the Forecast tab, select and build a Promo Planner Workbook.
    - Set causal variables for items and locations historically in that the selected promotions are active.
    - Commit your changes to the master database by selecting **Commit Now** from the File menu.
4. On the Forecast tab, select **Run Batch**.
    - Generate a Forecast.

Viewing a Forecast That Includes Promotion Effects

Perform this procedure to view a forecast that includes promotion effects.

1. On the Forecast tab, select and build a Forecast Approval Workbook and include System Baseline in your workbook.
2. In the Final Level Worksheet, review the System Baseline and the System Forecast. The System Baseline is predicted demand given no causal effects. The System Forecast is the sum of the System Baseline and the Promotional Peak calculated during the forecast generation process based on the causal data and settings.

Viewing and Editing Promotion System-Calculated Effects

Perform this procedure to view and edit promotion system-calculated effects.

1. On the Promote tab, select and build a Promo Effect Maintenance Workbook.

- In the Final Promotion Effect Parameters Worksheet, review the System Calculated Effect.
- If the System Calculated Effect is to be modeled using a method other than Automatic, edit the Promo Effect Type.
- If you choose to adjust the system-calculated effect, adjustments can be made to the System Effect Override. You must also set the Promo Effect Type to **Override All** or **Override Future Only**.

2. On the Forecast tab, select **Run Batch**.
   - Generate a Forecast.

Promotion Simulation (What-if?) and Analysis

Follow this procedure to perform analysis on past promotions and simulate the effects of historic or future promotions:

1. On the Promote tab, select and build a Promotion Effectiveness Workbook.
2. In the Final Promotion Effect Parameters Worksheet, review the Promotion System Calculated Effect and edit the System Effect Override then select **Calculate** to recalculate the Future Forecast.

Note: Changes to the Promotion Effectiveness Workbook cannot be committed; however, this workbook is useful for:

- Reporting on the performance of past promotions.
- Simulating the effects of future promotions to support more accurate promotional modeling.

14

Grouping Management

Grouping Management allows you to review the automatic group number assignments set for item/location.You can also change any group number assignments based on the information provided in the Grouping Management Workbook.

Grouping Management Workbook

This workbook includes the Review and Approve Group Assignments worksheet

Review and Approve Group Assignments Worksheet

This worksheet displays the group assignments and some historical demand information.

Measures: Review and Approve Group Assignments Worksheet

The Review and Approve Group Assignments worksheet contains the following measures:

Group Assignment

This measures determines to which group a certain item/location combination belongs.

There are various ways to make the assignment. First, the assignment can be determined outside of RDF, and then loaded. Second, rules can be configured to determine item location combinations that belong to certain groups. Third, you can manually assign item location combinations to a group.

Note: Any item/location combination can only belong to one group.

History Length

This measure contains the number of periods in history since the first sale. It is calculated as:

History Length = today - date of first sale

This measure is useful when you want to manually assign time series to groups. For example, you may want to assign every item/location with 55 weeks of sales to Group number 3. Or you may want to assign every item/location with sales history less than three weeks to Group 5.

History Length Tier

This measure displays the classification for a range of history length. For example, RDF can display short, medium, or long history length tiers, that can be set as 0-10 weeks, 11-56 weeks, and 57 and higher.

15

Cross Promotional Effects Module (CPEM)

This chapter provides an introduction to Cross Promotional Effects Module (CPEM). CPEM is a data mining solution that determines promotional Cannibalization, or Halo relationships, or both between items or groups of items.

About Promotional Effects

We refer to Cannibalization effects when a promoted item causes a drop in sales of a similar item. An example of this behavior is when a brand of mustard is promoted, it likely experiences an increase in sales and other mustard brands' demand decreases. After the promotional period the demand pattern of the entire mustard category returns to pre-promotion levels.

Conversely, we refer to Halo effects when a promoted item, or group of items causes an increase in sales of a complementary item or items. A typical case is when hot dogs are promoted, hot dog buns are likely to see a peak in sales as well.

Cannibalization is expected to happen at low levels in the merchandise hierarchy, like style or subclass. However, Halo is detected at much higher levels, like category or department.

CPEM is a module within RDF, but it is implemented separately of a RDF forecasting environment. The nature of the problem that mining for cross promotional relationships poses, that is, search for relationships among items in different areas of business, is not compatible with the master/local domain structure of a typical RDF environment. CPEM is implemented in a single domain, such that mining for Halo effects can span multiple departments.

The inputs to CPEM are sales and promotion information for which there is formal integration with the forecasting solution.

The output of CPEM is the promotional elasticity among items, which is location specific. The intersection of the output is item/item/location.

Currently, RDF Causal is predicting promotional sales based on item Self Promo Lifts. Basically the promotion of an item is considered to affect only sales of that item. With information about cross promotional elasticities, the promotion planners have the complete picture on how a promotion affects the entire business, and can act accordingly.

RDF was launched with Promotional Halo and Cannibalization capabilities in January 2013. The functionality has had limited adoption due to various factors including necessary data required to support reliable results, market maturity and implementation complexity. In an attempt to share lessons learnt with our experience so far, we would recommend the following to our customers considering implementing this functionality:

Starting Conditions

Both Halo and Cannibalization deal with the cross item effects of running promotions, that is, the impact of promoting an item on the sales of other related items. The most important inputs for estimating these cross effects are a robust baseline forecast and self-promotion effects from RDF Causal. Hence, we recommend that CPEM be considered for implementation only after RDF Causal has been up and live for a 12-18 month period, delivering reliable baseline and causal forecasts to drive business decisions.

- **Replacement SKUs** — All historical SKUs that are replacement or promotional variants of each other need to be grouped together into a Plan SKU. The Plan SKU level needs to be used within RDF for estimating both baseline forecasts and causal forecasts.
- **Appropriate pre-processing** — That corrects for missing sales periods and other factors that might bias the estimated baseline forecast and hence impact that quality of CPEM output.

Set Up for Cannibalization

Careful configuration and pre-implementation set up is required for Cannibalization to run. It is important to note that Cannibalization (though configurable) is designed to run at an item-group level (as defined in the following list) and not the individual SKU level.

Note: Cannibalization needs to stay inside a local domain.

- **Item-group or L1 Set Up** — This is the level at which Cannibalization is estimated. SKUs have to be grouped (outside the system and fed in). These are groups of items that are typically promoted together. For example, in the yogurt Category, SKUs of a particular brand, size, fat content but different flavor variants. An attribute analysis exercise, should inform this L1-grouping. What are the key combinations of attributes that need to be the same for all items in an L1-group and what attributes can be different? It is important to note that these attributes will vary significantly, from Category to Category.
- **Cannibalization-group Or L2 Set Up** — L2 is the level within which Cannibalization is estimated. Only items within a Cannibalization group are analyzed for possible Cannibalization, when one or more items are promoted. This could be a grouping of one or more classes or sub-classes. Note that L2 grouping needs to be a roll up of L1 groupings, that is, every L1 item-group needs to cleanly map to one and only one L2 grouping (many to one mapping from L1 to L2).
- **L1-pairing Set Up** — In addition to set up of L1 or Item-groups, RDF's CPEM also needs to be told of legal/possible L1-pairs. This requires careful analysis on a Category by Category basis. For example. SKUs within a certain regular price range could be considered Cannibalistic and hence valid L1-pairs. It is recommended that this analysis be conducted carefully, ideally in conjunction with Customer Decision Trees (CDTs) to ensure the correct pairings are set up for each Category.
- **Co-promotion Effects are Turned Off** — A system flag that effectively looks for conditions when both L1 pairings are promoted at the same time and accounts for it appropriately.

Halo effects

Current functionality is designed for a limited use case for businesses with minimal promotional activity. The solution aims to estimate Halo effects by observing the impact on sales of other subclasses when a subclass is promoted. As we vetted this with a broad set of retail use cases with significant promotional activities, we recognized that this functionality is not adequate. We are revisiting our approach to estimating Halo effects and welcome partnerships with retailers.

CPEM Workbooks and Wizards

The CPEM wizards allow you to setup the Effect Estimation Administration Workbook and Effect Estimation Review and Approval Workbook.

CPEM includes these workbooks and worksheets:

Workbooks	Description	Worksheets
Effect Estimation Administration Workbook	Allows you to setup parameters and thresholds, and their overrides, before the batch run.	Default Parameters Worksheet Cannibalization Parameter Override Worksheet Halo Parameter Override Worksheet Halo Products Mapping Worksheet
Effect Estimation Review and Approval Workbook	Allows you to review the cross promotional elasticities and statistical information related to the analysis. Based on the statistics, you can approve the elasticities or adjust them based on business knowledge.	Cannibalization Effects Worksheet Cannibalization Effects Percentage Worksheet Cannibalization Effects Drive View Worksheet Cannibalization History Data Worksheet Halo Effects Worksheet Halo Effects Percentage Worksheet Halo Effects Drive View Worksheet Halo History Data Worksheet Cannibalization Statistic Report Worksheet Halo Statistic Report Worksheet

Effect Estimation Administration Workbook

The Effect Estimation Administration workbook allows you to setup parameters and thresholds, and their overrides, before the batch run.

Opening or Creating a Effect Estimation Administration Workbook

The Effect Estimation Administration wizard steps you through the process of creating a new Effect Estimation Administration workbook. To access the Effect Estimation Administration, select **Open** from the File menu to bypass the wizard and open an existing Effect Estimation Administration workbook, or perform the following steps:

1. Within the Local domain, select **New** from the **File** menu.
2. On the CPEM tab, select **Effect Estimation Administration**. Click **OK**.
3. The Effect Estimation Administration wizard opens and prompts you to select the products for promotion effect estimation. Click **Next**.

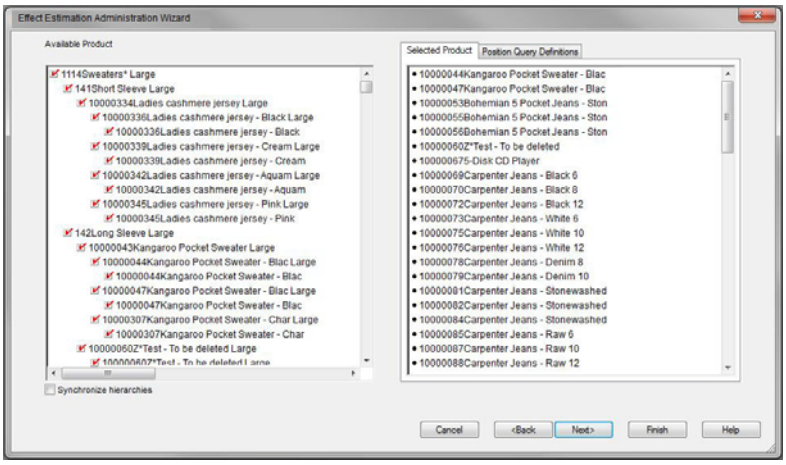

4. Select the locations for promotion effect estimation. Click **Next**.

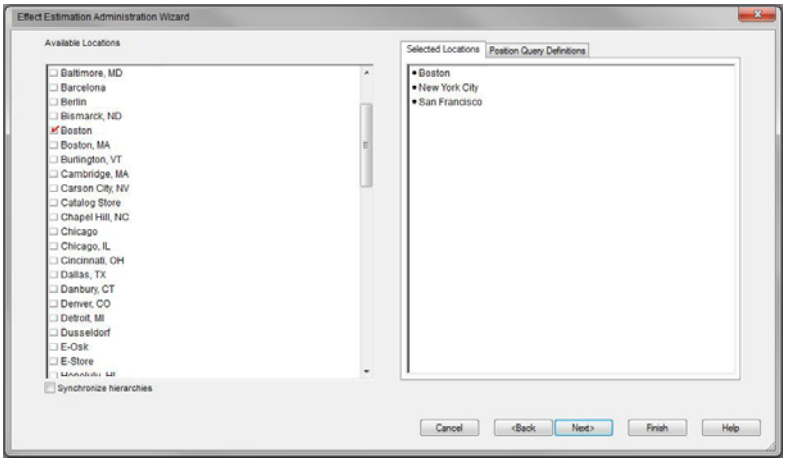

5. Select the right-hand-side (RHS) products for promotion effect estimation. Click **Next**.

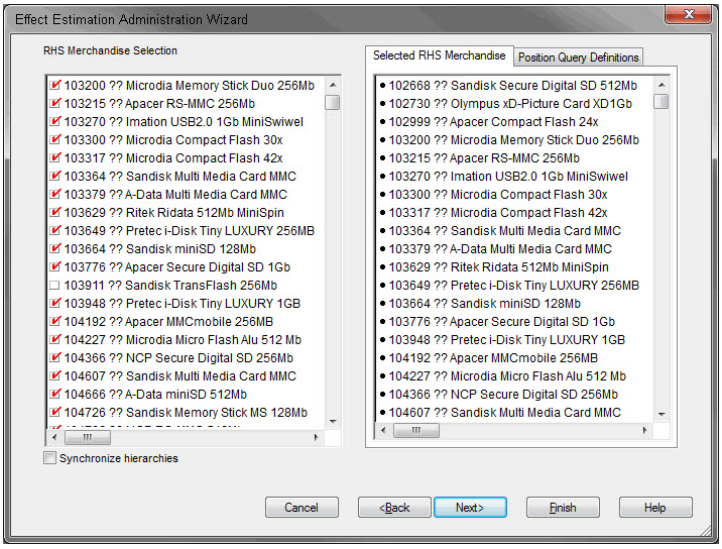

6. Click **Finish** to build the workbook.

Worksheets for the Effect Estimation Administration Workbook

The following sections describe the components of the Effect Estimation Administration workbook.

The Effect Estimation Administration workbook includes these worksheets:

- Default Parameters Worksheet
- Cannibalization Parameter Override Worksheet
- Halo Parameter Override Worksheet
- Halo Products Mapping Worksheet

Default Parameters Worksheet

The Default Parameters worksheet allows you to setup Default parameters and thresholds, and their overrides, before the batch run.

Figure 15–1 Default Parameters Worksheet

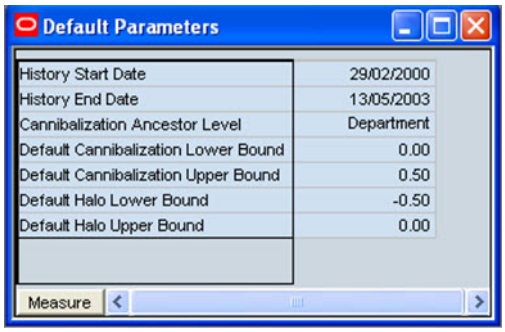

Measures: Default Parameters Worksheet

The Default Parameters worksheet contains the following measures:

Default Cannibalization P-Value

The P-value is used in statistics to measure the confidence one can have in an estimate. The Default Cannibalization P-value represents the minimum confidence threshold. During the cannibalization effects estimation process, any effects with P-values below the threshold are deemed insignificant and discarded.

Default Cannibalization Self Elasticity Check

Place a check in this parameter (set to True) if you want to ensure that the cannibalization elasticity is less than its shelf elasticity. If checked, the calculated elasticity which is higher than its own elasticity will be discarded. This parameter serves as a global override.

Default Halo P-Value

The P-value is used in statistics to measure the confidence one can have in an estimate. The Default Halo P-value represents the minimum confidence threshold. During the Halo effects estimation process, any effects with P-values below the threshold are deemed insignificant and discarded.

Default Halo Self Elasticity Check

Place a check in this parameter (set to True) if you want to ensure that the halo elasticity is less than its self elasticity. If checked, the calculated elasticity which is higher than its own elasticity will be discarded. This parameter serves as a global override.

History Start Date

The first calendar date considered in the analysis. All measure data prior to this date is discarded.

History End Date

The last calendar date considered in the analysis. All measure data after this date is discarded.

Cannibalization Ancestor Level

This is the dimension in the merchandise hierarchy where Cannibalization is mined for. Only item pairs within this dimension are evaluated.

Cannibalization Lower Bound

Minimum Cannibalization value threshold. Cannibalization values that are lower than the threshold are discarded.

Cannibalization Upper Bound

Maximum Cannibalization value threshold. Cannibalization values that are higher than the threshold are discarded (set to zero).

Halo Lower Bound

Minimum Halo value threshold. Halo values that are lower than the threshold are discarded.

Halo Upper Bound

Maximum Halo value threshold. Halo values that are higher than the threshold are discarded (set to zero).

Cannibalization Parameter Override Worksheet

The Cannibalization Parameter Override worksheet allows you to setup Cannibalization parameters and thresholds, and their overrides, before the batch run.

Figure 15–2 Cannibalization Parameter Override Worksheet

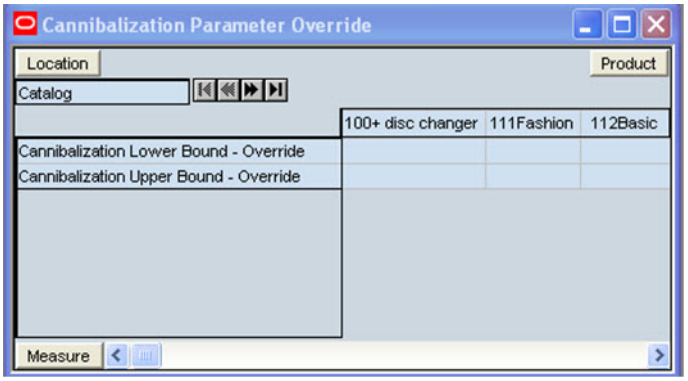

Measures: Cannibalization Parameter Override Worksheet

The Cannibalization Parameter Override worksheet contains the following measures:

Cannibalization P-Value - Override

Override of the Default Cannibalization P-Value threshold. The cannibalization effects whose p-values are below the threshold are discarded.

Cannibalization Self Elasticity Check - Override

This is an override of the global setting. If checked (set to True), the calculated cannibalization elasticity which is higher than its self elasticity will be discarded.

Cannibalization Lower Bound - Override

Override of the minimum Cannibalization value threshold. Cannibalization values that are lower than the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, subclass/region.

Cannibalization Upper Bound - Override

Override of the maximum Cannibalization value threshold. Cannibalization values that are higher than the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, subclass/region.

Halo Parameter Override Worksheet

The Halo Parameter Override worksheet allows you to setup Halo parameters and thresholds, and their overrides, before the batch run.

Figure 15–3 Halo Parameter Override Worksheet

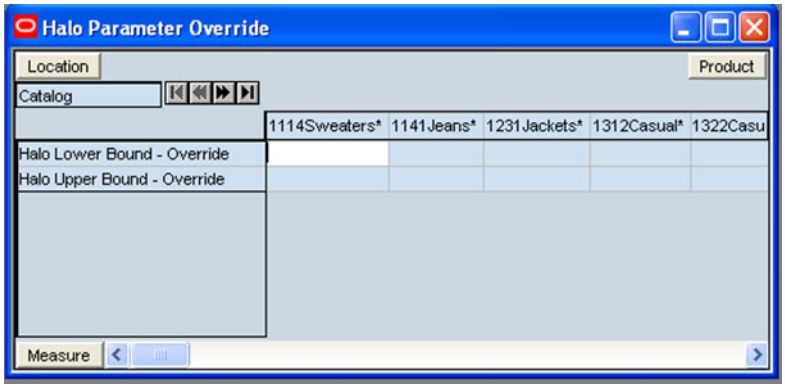

Measures: Halo Parameter Override Worksheet

The Halo Parameter Override worksheet contains the following measures:

Halo P-Value - Override

Override of the Default Halo P-Value threshold. The Halo effects whose P-values are below the threshold are discarded.

Halo Self Elasticity Check - Override

This is an override of the global setting. If checked (set to True), the calculated halo elasticity which is higher than its self elasticity will be discarded.

Halo Lower Bound - Override

Override of the minimum Halo value threshold. Halo values that are lower than the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, class/region.

Halo Upper Bound - Override

Override of the maximum Halo value threshold. Halo values that are higher than the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, class/region.

Halo Products Mapping Worksheet

The Halo Products Mapping worksheet allows you to specify item pairs for which a Halo relationship is mined for. For item pairs that do not have the check box selected, the system will not attempt to look for a Halo relationship.

Figure 15–4 Halo Products Mapping Worksheet

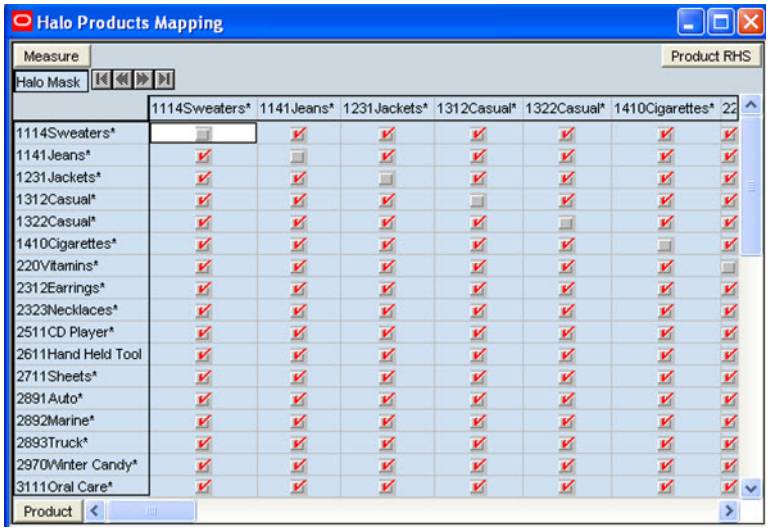

Measures: Halo Products Mapping Worksheet

The Halo Products Mapping worksheet contains the following measures:

Halo Mask

This measure defines item groupings that are analyzed for Halo relationship. The item groupings are typically bound to the merchandise hierarchy and can be as in Figure 15–4, higher than class to class. This is a writeable measure, and you are expected to use your business knowledge to mark pairs that could have Halo relationships.

Effect Estimation Review and Approval Workbook

The Effect Estimation Review and Approval workbook provides you the ability to review the cross promotional elasticities and statistical information related to the analysis. Based on the statistics, you can approve the elasticities or adjust them based on your business knowledge.

This workbook has two tabs:

Effect Estimation Review and Approval Tab	Worksheets
Cross Promo Effect	Cannibalization Effects Worksheet Cannibalization Effects Percentage Worksheet Cannibalization Effects Drive View Worksheet Cannibalization History Data Worksheet Halo Effects Worksheet Halo Effects Percentage Worksheet Halo Effects Drive View Worksheet Halo History Data Worksheet
Cross Effect Statistics Reports	Cannibalization Statistic Report Worksheet Halo Statistic Report Worksheet

Opening the Effect Estimation Review and Approval Workbook

The Effect Estimation Review and Approval wizard steps you through the process of creating a new Effect Estimation Review and Approval workbook. To access the Effect Estimation Review and Approval workbook, select **Open** from the File menu to bypass the wizard and open an existing Effect Estimation Review and Approval workbook, or perform the following steps:

1. Within the Local domain, select **New** from the **File** menu.
2. On the CPEM tab, select **Effect Estimation Review and Approval**. Click **OK**.
3. The Effect Estimation Review and Approval wizard opens and prompts you to select the products for promotion effect estimation review and approval. Click **Next**.

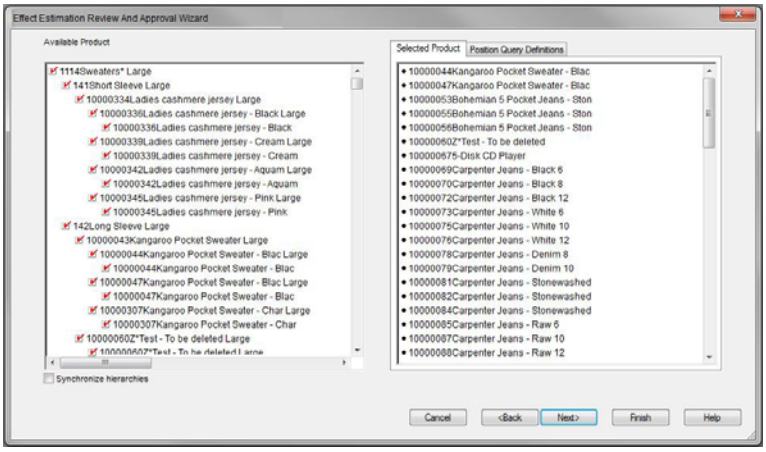

4. Select the locations for promotion effect estimation review and approval. Click **Next**.

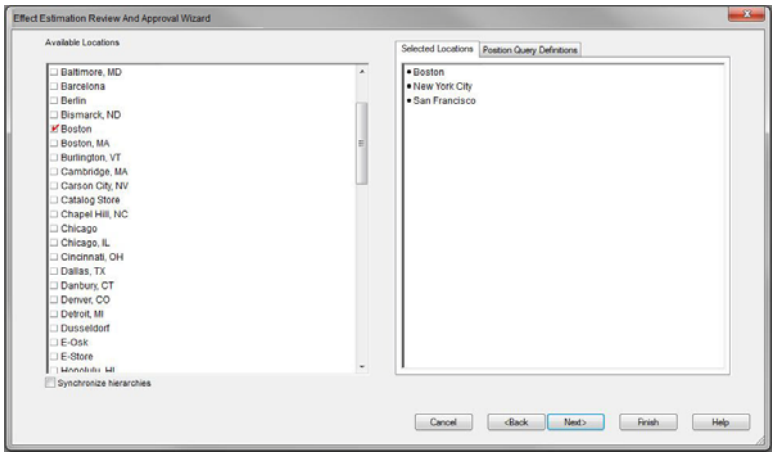

5. Select the right-hand-side (RHS) products for promotion effect estimation review and approval. Click **Next**.

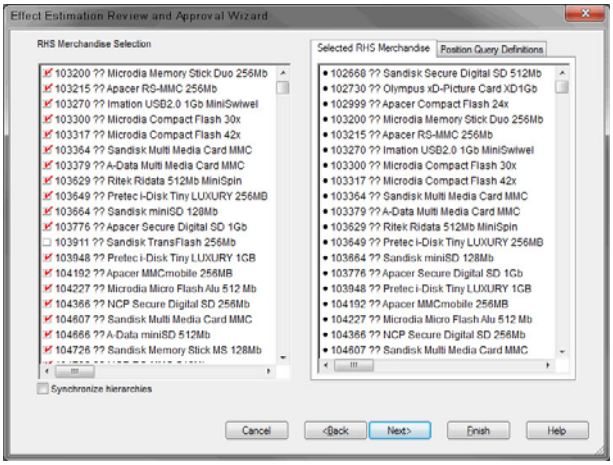

6. Select the dates for promotion effect estimation review and approval. Click **Next**.

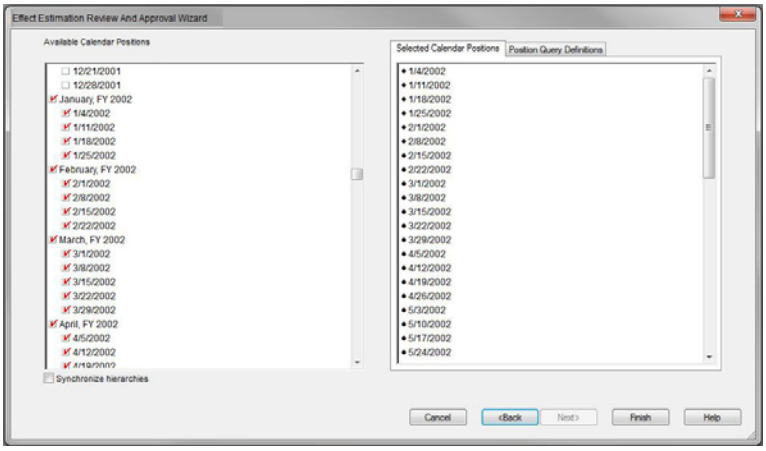

7. Click **Finish** to build the workbook.

Worksheets for the Effect Estimation Review and Approval Workbook

The following sections describe the components of the Effect Estimation Review and Approval workbook.

The Effect Estimation Review and Approval workbook includes these worksheets:

- Cannibalization Effects Worksheet
- Cannibalization Effects Percentage Worksheet
- Cannibalization Effects Drive View Worksheet
- Cannibalization History Data Worksheet
- Halo Effects Worksheet
- Halo Effects Percentage Worksheet
- Halo Effects Drive View Worksheet
- Halo History Data Worksheet
- Cannibalization Statistic Report Worksheet
- Halo Statistic Report Worksheet

Cannibalization Effects Worksheet

The Cannibalization Effects worksheet allows you to review and adjust the Cannibalization cross elasticities.

Figure 15–5 Cannibalization Effects Worksheet

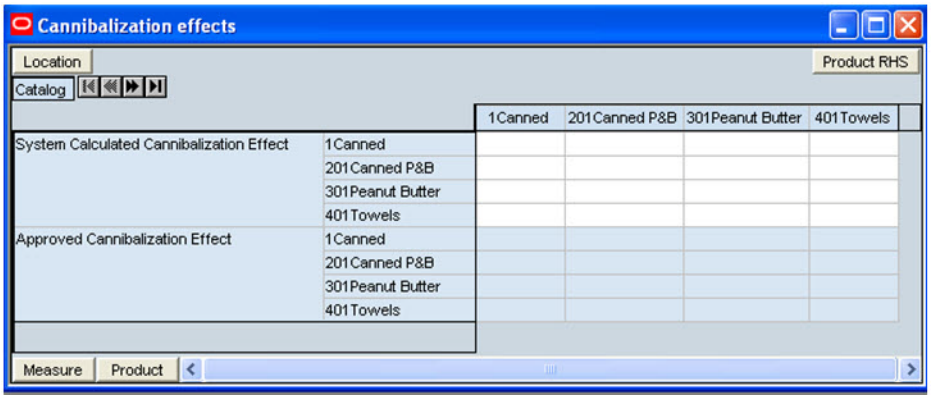

Measures: Cannibalization Effects Worksheet

The Cannibalization Effects worksheet contains the following measures:

Approved Cannibalization Effect

This measure stores the Cannibalization value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

System Calculated Cannibalization T Statistic

This measure displays the system calculated T-Statistics for each cannibalization effect at product/product/location level.

System Calculated Cannibalization Coefficient Standard Error

This measure displays the system calculated standard error for each cannibalization effect at product/product/location level.

System Calculated Cannibalization P Value

This measure displays the system calculated P-value for each cannibalization effect at product/product/location level.

Selected Cannibalization P Value

This measure displays the user selected P-value for each cannibalization effect at product/product/location level.

System Calculated Cannibalization Effect

This measure displays the system calculated Cannibalization effects at product/product/location level.

Cannibalization Effects Percentage Worksheet

The Cannibalization Effects Percentage worksheet allows you to review the reactions of the impacted item's demand change given a change in the driver item's demand. The percentage changes due to cannibalization effects will be displayed for a series of self lifts of the driver item.

Measures: Cannibalization Effects Percentage Worksheet

The Cannibalization Effects Percentage worksheet contains the following measures:

Approved Cannibalization Effect

This measure stores the approved cannibalization effects.

RHS Product Sales Increase 10%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 10%.

RHS Product Sales Increase 20%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 20%.

RHS Product Sales Increase 30%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 30%.

RHS Product Sales Increase 40%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 40%.

RHS Product Sales Increase 50%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 50%.

RHS Product Sales Increase 60%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 60%.

RHS Product Sales Increase 70%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 70%.

RHS Product Sales Increase 80%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 80%.

RHS Product Sales Increase 90%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 90%.

RHS Product Sales Increase 100%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 100%.

Cannibalization Effects Drive View Worksheet

The Cannibalization Effects Drive View worksheet contains the same information as the Cannibalization Effects Worksheet, but has a different layout and a position query defined.

Figure 15–6 Cannibalization Effects Drive View Worksheet

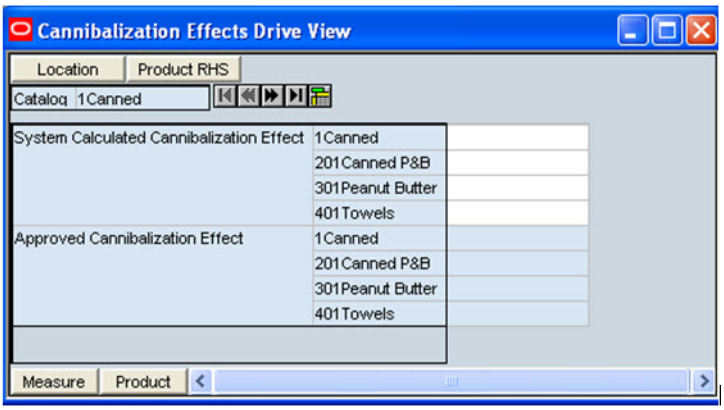

Measures: Cannibalization Effects Drive View Worksheet

The Cannibalization Effects Drive View worksheet contains the following measures:

System Calculated Cannibalization Effect

This measure displays the system calculated Cannibalization effects at product/product/location level.

Approved Cannibalization Effect

This measure stores the Cannibalization value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

Cannibalization History Data Worksheet

The Cannibalization History Data worksheet allows you to view the historical facts used to estimate the Cannibalization effects.

Figure 15–7 Cannibalization History Data Worksheet

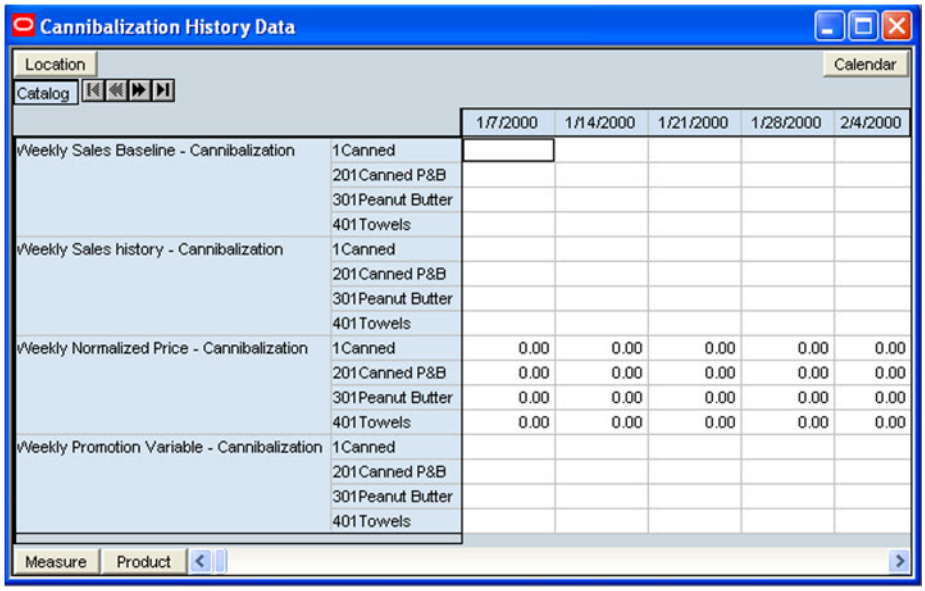

Measures: Cannibalization History Data Worksheet

The Cannibalization History Data worksheet contains the following measures:

Weekly Sales Baseline - Cannibalization

This measure represents the sales history cleansed of Promo Self Effects. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week.

Weekly Sales History - Cannibalization

This measure represents the total sales history. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week.

Weekly Normalized Price - Cannibalization

This measure stores the normalized price. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week. There is no guarantee that all items in a subclass (or region) have the same price, hence the need to normalize them.

Weekly Promotion Variable - Cannibalization

This measure stores the information if any promotion was active for a given time period. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week.

Halo Effects Worksheet

The Halo Effects worksheet allows you to review and adjust the Halo cross elasticities.

Figure 15–8 Halo Effects Worksheet

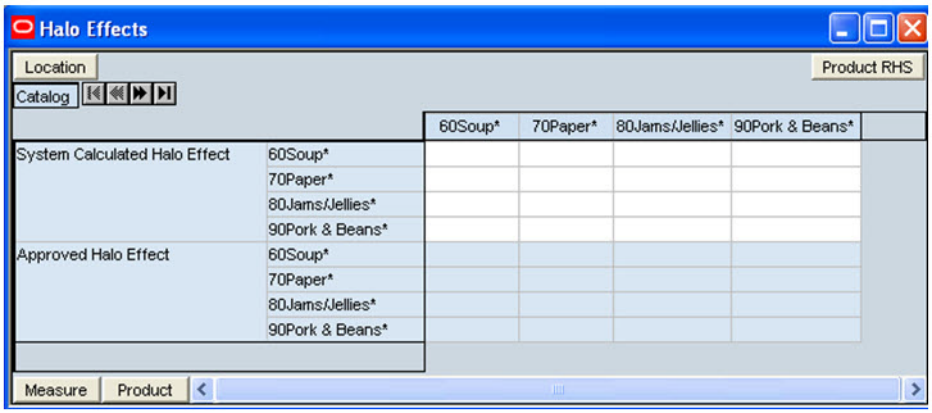

Measures: Halo Effects Worksheet

The Halo Effects worksheet contains the following measures:

System Calculated Halo Effect

This measure displays the system calculated Halo effects at product/product/location level.

Approved Halo Effect

This measure stores the Halo value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

System Calculated Halo T-Statistic

This measure displays the system calculated T-Statistics for each Halo effect at product/product/location level.

System Calculated Halo Coefficient Standard Error

This measure displays the system calculated standard error for each Halo effect at product/product/location level.

System Calculated Halo P-Value

This measure displays the system calculated P-value for each Halo effect at product/product/location level.

Selected Halo P-Value

This measure displays the user selected P-value for each Halo effect at product/product/location level.

Halo Effects Percentage Worksheet

The Halo Effects Percentage worksheet allows you to review the reactions of the impacted item's demand change given a change in the driver item's demand. The percentage change due to Halo effects will be displayed for a series of self lifts of the driver item.

Measures: Halo Effects Percentage Worksheet

The Halo Effects Percentage worksheet contains the following measures:

Approved Halo Effect

This measure stores the approved Halo effects.

RHS Product Sales Increase 10%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 10%.

RHS Product Sales Increase 20%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 20%.

RHS Product Sales Increase 30%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 30%.

RHS Product Sales Increase 40%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 40%.

RHS Product Sales Increase 50%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 50%.

RHS Product Sales Increase 60%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 60%.

RHS Product Sales Increase 70%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 70%.

RHS Product Sales Increase 80%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 80%.

RHS Product Sales Increase 90%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 90%.

RHS Product Sales Increase 100%

This measure displays the impacted item's sales percentage change when the driver item's (the RHS) sales increase by 100%.

Halo Effects Drive View Worksheet

The Halo Effects Drive View worksheet contains the same information as the Halo Effects Worksheet, but has a different layout and a position query defined.

Figure 15–9 Halo Effects Drive View Worksheet

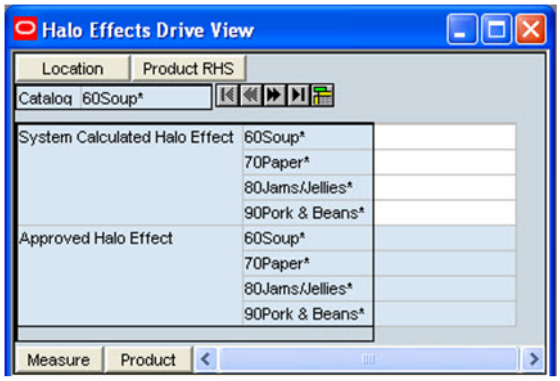

Measures: Halo Effects Drive View Worksheet

The Halo Effects Drive View worksheet contains the following measures:

System Calculated Halo Effect

This measure displays the system calculated Halo effects at product/product/location level.

Approved Halo Effect

This measure stores the Halo value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

Halo History Data Worksheet

The Halo History Data worksheet allows you to view the historical facts used to estimate the Halo effects.

Figure 15–10 Halo History Data Worksheet

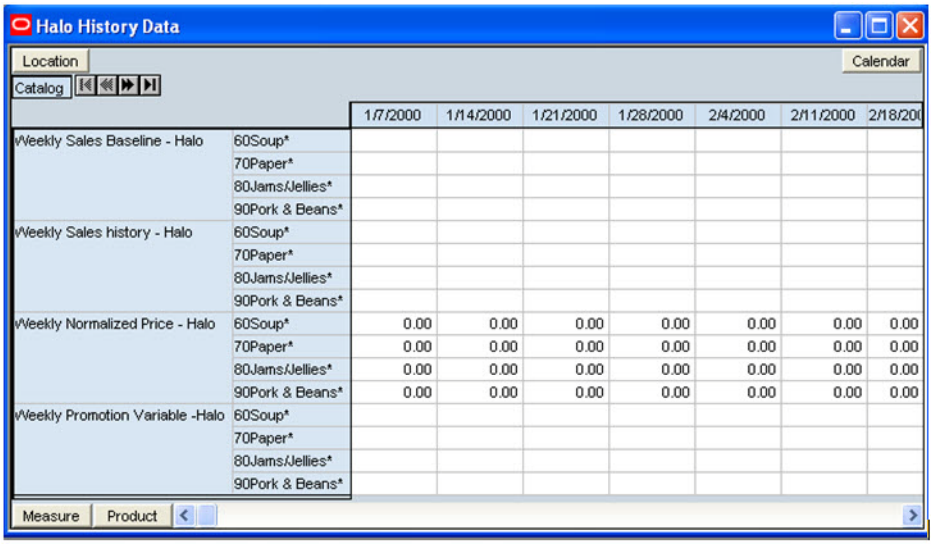

Measures: Halo History Data Worksheet

The Halo History Data worksheet contains the following measures:

Weekly Sales Baseline - Halo

This measure represents the sales history cleansed of Promo Self Effects. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week.

Weekly Sales History - Halo

This measure represents the total sales history. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week.

Weekly Normalized Price - Halo

This measure stores the normalized price. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week. There is no guarantee that all items in a subclass (or region) have the same price, hence the need to normalize them.

Weekly Promotion Variable - Halo

This measure stores the information if any promotion was active for a given time period. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week.

Cannibalization Statistic Report Worksheet

The Cannibalization Statistic Report worksheet allows you to view Cannibalization statistics on the cross effects.

Figure 15–11 Cannibalization Statistic Report Worksheet

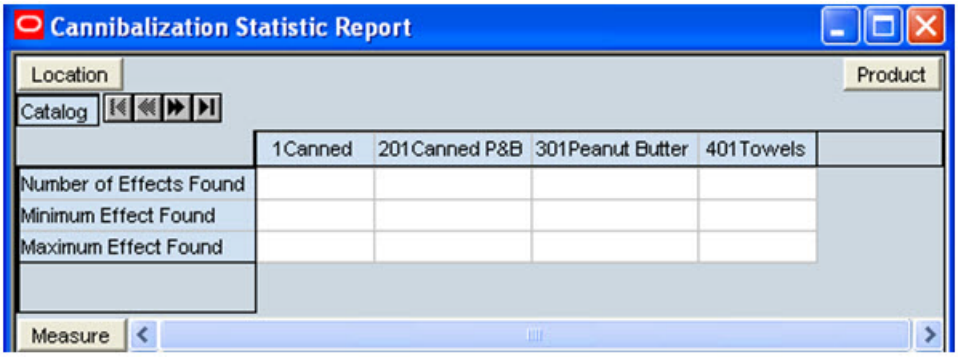

Measures: Cannibalization Statistic Report Worksheet

The Cannibalization Statistic Report worksheet contains the following measures:

Number of Effects Found

This measure displays the number of Cannibalization cross effects found for a given intersection. A typical example for an intersection would be subclass/region.

Minimum Effect Found

This measure displays the minimum value of the Cannibalization effects found for a given intersection. A typical example for an intersection would be subclass/region.

Maximum Effect Found

This measure displays the maximum value of the Cannibalization effects found for a given intersection. A typical example for an intersection would be subclass/region.

Halo Statistic Report Worksheet

The Halo Statistic Report worksheet allows you to view Halo statistics on the cross effects.

Figure 15–12 Halo Statistic Report Worksheet

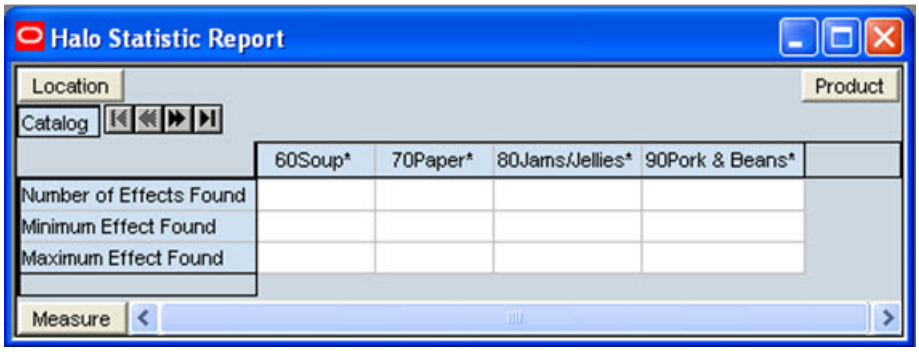

Measures: Halo Statistic Report Worksheet

The Halo Statistic Report worksheet contains the following measures:

Number of Effects Found

This measure displays the number of Halo cross effects found for a given intersection. A typical example for an intersection would be class/region.

Minimum Effect Found

This measure displays the minimum value of the Halo effects found for a given intersection. A typical example for an intersection would be class/region.

Maximum Effect Found

This measure displays the maximum value of the Halo effects found for a given intersection. A typical example for an intersection would be class/region.

16

Curve

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 lifecycle profiles.

Curve meets the need of operational systems (such as Oracle Retail Demand Forecasting (RDF) and Oracle Retail Merchandising System (RMS)) to have sales unit predictions at a more detailed level than those provided by 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 RDF require data to be at the lowest level of execution (that is, item/store/week or item/store/day) because these systems are responsible for ensuring that the right quantity of each product is in the right store at the right time.

In the most basic sense, a profile represents the ratio of an aggregate dimension to the dimension for execution. For example, you may have a forecast generated at the item/store/week level, but for execution purposes the data must be spread down to the item/store/day level. It is the point of aggregation (source-level) and the desired destination intersection (final profile) that are the unique identifiers of each profile. Using this example, the point of aggregation of the data (where the data equals 100%) is item/store/week, and the desired destination intersection (where all data ratios sum to 100%) is item/store/day.

There are several parameters within RDF that may take a Curve-generated profile as an input. These are: Causal Aggregation Profile, Causal Spread Profile, Seasonal Profile, and Spreading Profile. The most common input from Curve that is used in RDF is the Spreading Profile. This profile can be manually generated and approved by you. It can also be dynamically generated as part of the RDF batch forecast process. For more information on the different parameters in RDF, refer to *Chapter 3, "Setting Forecast Parameters."*.

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 positions for which the values set in Profile Administration differ from the defaults. This step is necessary in those situations where it is not efficient to use the same parameters for all positions in the hierarchy data.

The Profile Maintenance Workbook allows edits to intersections that vary from the default values set in the Profile Administration Workbook.

The Profile Maintenance workbook includes these worksheets:

- Final Approval and Sourcing Worksheet
- Final Training Window Worksheet

Creating a Profile Maintenance Workbook

As in Profile Administration, the first Profile Maintenance wizard window prompts you to select a final profile level.

1. Within the Master or Local domain; select **New** from the **File** menu.
2. Select the **Curve t**ab to display a list of workbooks for Profiling.
3. Select **Profile Maintenance**. Click **OK**.
4. The Profile Maintenance Wizard opens and prompts you to select the final profile. Make the appropriate selection, and click **Next.**
5. If the Location hierarchy is defined in the profile, select the locations to include in the workbook, and click **Next**.
6. If the Merchandise hierarchy is defined in the profile, select the products to include in the workbook, and click **Next**.
7. If the Calendar hierarchy is defined in the profile, select the time periods to include in the workbook, and click **Next**.
8. An additional wizard window prompts you to select any additional measures (that is, measures not standard in the Profile Maintenance Workbook) that you would like included. The measure options available in this window are set in the RPAS Security Administration Workbook/Workbook Template Measure Rights worksheet. Make the appropriate selections (if any). Click **Finish** to display the workbook.

Final Approval and Sourcing Worksheet

Figure 16–1 Final Approval and Sourcing Worksheet

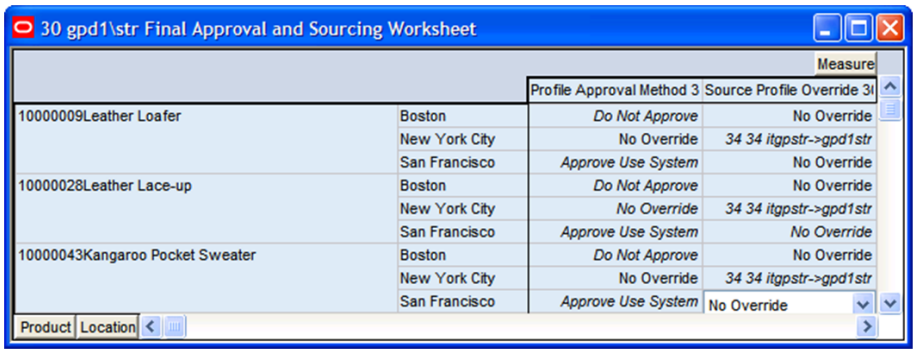

Measures: Final Approval and Sourcing Worksheet

The following is a description of the measures contained in the Final Approval and Sourcing Worksheet:

Profile Approval Method

The Profile Approval Method displays the primary approval policy that is used for the profile. No Override is displayed in this field if the Default Approval Method is used. **Select Approve Use System** if profile results are to be automatically approved during the batch profile generation. Select **Do Not Approve** if profile results are to be manually approved by the user. If a profile is being generated dynamically (to support Source Level Forecasting) as part of the RDF batch forecast process, the Default Profile Approval Method should be set to Approve Use System.

> **Note:** The Profile Approval Method is only used when the Use Advanced Profile Features are enabled.

Source Profile Override

Make edits to the Source Profile Override if the source-level for an intersection varies from the Default Source Profile. No Override is displayed in this field if the value set in the Default Source Profile is to be used. When only a single source is configured for a profile, all profile results are calculated using the same intersections. Therefore, edits to this parameter are not required. When multiple sources are defined for a profile, some profile results are calculated using different intersections. Select the appropriate source-level for an intersection.

Final Training Window Worksheet

Figure 16–2 Final Training Window Worksheet

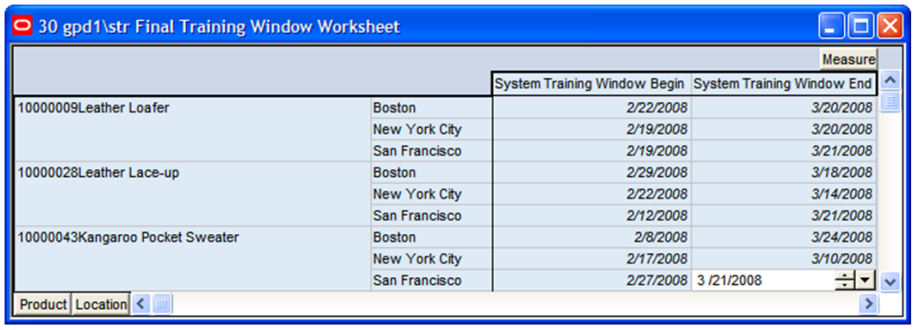

Measures: Final Training Window Worksheet

The following is a description of the measures contained in the Final Approval and Sourcing Worksheet:

System Training Window Begin

Select a date in this field if this date is different from the default value set in the Training Window Start date in Profile Administration. Click the pop-up calendar to change the value in this field.

System Training Window End

Select a date in this field if this date is different from the default value set in the Training Window End date in Profile Administration. Click the pop-up calendar to change the value in this field.

Profile Administration Workbook

The Profile Administration Workbook allows you to set default parameters for profile generation, which is 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.

The Profile Administration workbook includes these worksheets:

- Profile Parameter Worksheet
- Profile and Source Level Intersection Worksheet

Selecting a Final Profile to Edit

The Profile Administration wizard requires you to select the final profile that you want to edit. These profiles are determined during the system implementation/configuration.

1. Within the Master or Local domain, select **New** from the **File** menu.
2. On the **Curve** tab, select **Profile Administration**. Click **OK**.
3. Select the final profile to analyze. Click **Finish**.

Profile Parameter Worksheet

The Profile Parameter Worksheet allows you to specify default values for parameters affecting profile generation. The following image provides an example of the Profile Parameters worksheet in a Master Domain with three partitions/Local Domains, partitioned on Group.

Figure 16–3 Profile Parameter Worksheet

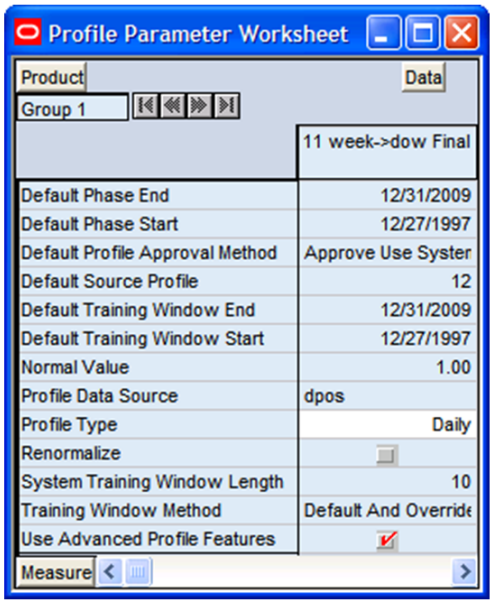

Measures: Profile Parameter Worksheet

The following is a description of the measures that are contained in the Profile Parameter Worksheet:

Default Phase End

The Default Phase End defines the end date of the period in which profile results will be applied. Click the pop-up calendar to change the value in this field. If phase definitions are unavailable, the default phase end date will be used. When calculating time profiles, default dates are used for intermediate computations. For computational efficiency, use the most common phase definition as the default value.

Default Phase Start

The Default Phase Start defines the first date of the period in which profile results will be applied. Click the pop-up calendar to change the value in this field. If phase definitions are unavailable, the default phase start date will be used. When calculating time profiles, default dates are used for intermediate computations. For computational efficiency, use the most common phase definition as the default value.

Default Profile Approval Method

The Default Profile Approval Method displays the primary approval policy to be used for the profile. Select Approve Use System if profile results are to be automatically approved during the batch profile generation. Select Do Not Approve if profile results are to be manually approved by the user. If a profile is being generated dynamically (to support Source Level Forecasting) as part of the RDF batch forecast process, the Default Profile Approval Method should be set to Approve Use System.

> **Note:** The Default Profile Approval Method is unavailable unless Use Advanced Profile Features is selected.

Default Source Profile

The Default Source Profile determines the primary source-level that will be used to generate the profile. When only a single source is configured for a profile, all profile results will be calculated using the same intersections. When multiple sources are defined for a profile, some profile results will be calculated using different intersections. The Profile Maintenance Workbook may be used to define exceptions to the Default Source Profile.

Default Training Window End

The Default Training Window End defines the last date in history that will be used to calculate profile results. Click the pop-up calendar to change the value in this field. The default date will be used only if the training window method is set to Defaults and Overrides.

Default Training Window Start

The Default Training Window Start defines the first date in history that will be used to calculate profile results. Click the pop-up calendar to change the value in this field. The default date will only be used if the training window method is set to Defaults and Overrides.

Normal Value

If the profile is aggregated to the Aggregation Intersection, this will be the value in all the cells (or zero (0) outside of the phase window). The desired setting for Normal Value is usually 1.00 (100%); however, there may be instances where it is desired to have the profile normalize to a value different than 1.00. For example, when generating seasonal profiles, such as by Week of Year, the Normal Value needs to be set to 52.

Profile Data Source

The Profile Data Source displays the name of the measure that contains the data to be used to generate the profile.

Note: For Dynamic profiles used by RDF, the profile Data Source can be left blank. Curve will use the Interim Forecast generated by RDF as the Data Source for generating the profile, and the value specified here will be ignored.

Profile Type

The Profile Type is a read-only measure that displays the value for this measure set in the Curve Plug-In. The Profile Type is used to determine the profile algorithm and validation required by the profile level. Profile Types are represented with pre-defined configuration information.

The following Profile Types share the same profile algorithm. The rationale for providing different types that have the same behavior is strictly to remind the user of the intent of the profile while using this workbook:

- **Store Contribution Profile** - The Store Contribution Profile is used to determine the data relationship between stores to aggregate dimensions in the location hierarchy.
- **Daily Profile** - The Daily Profile is used to determine the data relationship between a given day to the week in which it belongs.
- **Hourly Profile** - The Hourly Profile is used to determine the spreading ratios from aggregate dimensions to the hour, hour of day, or hour of week dimensions.
- **Product Profile** - The Product Profile is used to determine the data relationship between any two dimensions along the product hierarchy.
- **Size Profile** - Size profiles are used to determine the data relationship between any size in the product hierarchy and another dimension in the product hierarchy. Note that size can be configured as a dimension along an alternate rollup along the product hierarchy, such as itemParent-size.
- **User Defined Profile** - The User Defined Profile may be used to support any basic profile configuration.

The following Profile Types have unique behavior:

- **Diff Profile** - Diff Profiles are used to determine spreading ratios from aggregate dimensions in the Product hierarchy to diff dimensions. Used to support the spreading of data in RMS Allocation, Diff Profiles exhibit the same behavior as the previous profile types. However, unique to Diff Profiles is special validation of the relationship between the defined diff dimensions to dimensions along the main branch of the Product hierarchy. Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on validation criteria.

- **Daily Seasonal Profile** - The Daily Seasonal Profile is used to determine the data relationship between a given day of the week to aggregate dimensions in the calendar hierarchy. This profile type uses training window data to compute the profile. The resulting profile is then clipped to fit within the defined phase window.
- **Life Cycle Profile** - The Life Cycle Profile uses data along a user-defined training window and then stretches or shrinks data to fit a user-defined phase window.

Renormalize

Renormalize is a Boolean measure. When set to True, it automatically renormalizes the calculated profile result at the corresponding final-level. Typically, the renormalization is not necessary. For example, if you have a source profile at week of season and its final profile is at day of season, you would need to renormalize the final-level because going from week to day replicates the level. At day level, the profile will sum up to greater than 1 for a season. Since it was a week to day, it will probably sum up to seven. The renormalize will force the final profile to sum to 1.00 (100%).

System Training Window Length

The System Training Window Length is necessary when Use Training Window is set as the Training Window Method for the profile. This field specifies the number of weeks of the most recent data to use as the training window for calculating the profile. The System Training Window Length defaults to 10 weeks.

Training Window Method

The Training Window Method is used to determine the default method that is used to define the training window. The difference between these options is how the default Training Window start and default Training Window end is determined. The Training Window start/end measure in Profile Maintenance are overrides on item/store. If populated, they will take precedence over the default.

The options are:

- **Default and Overrides** - Uses the default dates as set in the Training Window Start Date and Training Window End Date measures.
- **Phase Definitions and Overrides** - Calculates the Training Window Start Date and Training Window End Date based on the Phase Start Date and Phase End Date measures.
- **Use Training Window** - Used with System Training Window Length to specify the number of weeks of the most recent data to use for calculating the profile. Today is used as default end. Today minus Training Window Length is used as default start.

Use Advanced Profile Features

This parameter is always selected by default.

Profile and Source Level Intersection Worksheet

The Profile and Source Level Intersection Worksheet is a read-only view to the different intersections defined for the Profile and Source level configured in the Curve Plug-in. Refer to the *Oracle Retail Predictive Application Server Solution Extension Configuration Guide* for more information on defining these intersections during configuration of the Curve Solution.

Figure 16–4 Profile and Source Level Intersection Worksheet

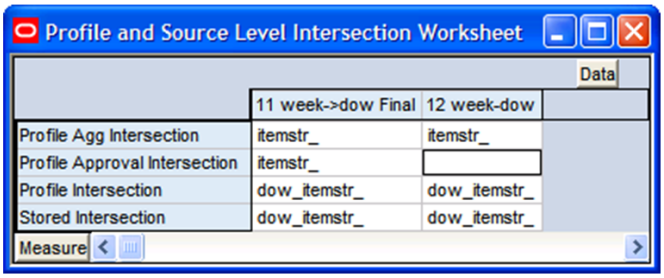

Measures: Profile and Source Level Intersection Worksheet

Profile Agg Intersection

The Profile Agg Intersection is the intersection where the profile will sum to one (or 100%). If the profile is being used as the Spreading Profile in RDF, this Aggregation Intersection should be the same as the Source Forecast Level.

Profile Approval Intersection

Assigned only at the Final Profile, the Approval Intersection is the intersection where the profile is approved. Approval Intersection should be higher than or equal to the Aggregation Intersection. If the profile is being used as the Spreading Profile in RDF, this Approval Intersection should be the same as the Aggregation Intersection.

Profile Intersection

The Profile Intersection is the intersection where an intermediate profile is calculated. This intermediate profile is then replicated down or aggregated up to the Stored Intersection. If the Store Intersection is the same as the Profile Intersection, the values in intermediate profile are copied to the Stored Intersection. The Profile Intersection must be lower than the Aggregation Intersection. If the profile is being used as the Spreading Profile in RDF, this Profile Intersection should be the same as the Final Forecast Level.

Stored Intersection

The Stored Intersection is the destination intersection of the profile. The intermediate profile produce at the Profile Intersection is either replicated down to or aggregated up to the Stored Intersection. If the Store Intersection is the same as the Profile Intersection, the values in intermediate profile are copied to the Stored Intersection. The Stored Intersection should not be greater than the Aggregation Intersection. If the profile is being used as the Spreading Profile in RDF, this Stored Intersection should be the same as the Profile Intersection.

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.

The Profile Approval workbook includes these worksheets:

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

Creating a New Profile Approval Workbook

1. With the Local domain, select **New** from the **File** menu.
2. Select the **Curve** tab to display a list of workbooks.
3. Select Profile Approval. Click **OK**.
4. The Profile Approval Wizard opens and prompts you to select the final profile. Make the appropriate selection. Click **Next**.
5. If the Location hierarchy is defined in the profile, select the locations to include in the workbook. Click **Next**.
6. If the Merchandise hierarchy is defined in the profile, select the products to include in the workbook. Click **Next**.
7. If the Calendar hierarchy is defined in the profile, select the time periods to include in the workbook. Click **Next**.
8. An additional wizard window prompts you to select any additional measures (that is, measures not standard in the Profile Approval Workbook) that you want to include. The measure options available in this window are set in the RPAS Security Administration Workbook/Workbook Template Measure Rights worksheet. Make the appropriate selections (if any). Click **Finish** to display the workbook.

Final Profile Worksheet

Through this worksheet, you can view the system calculated final profile and make adjustments to this profile. Following is an example Final Profile Worksheet for a day-of-week profile.

Figure 16–5 Final Profile Worksheet

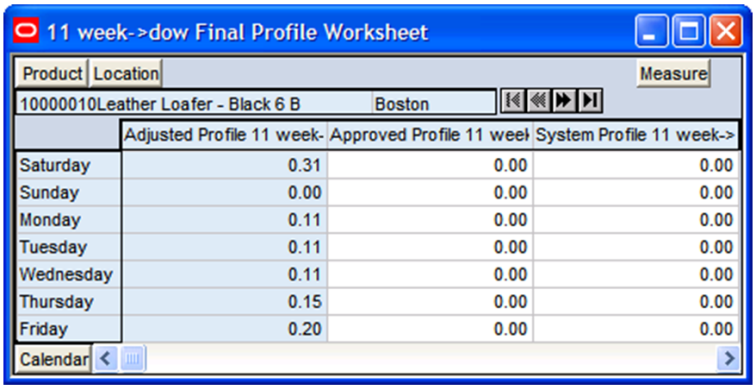

Measures: Final Profile Worksheet

Adjusted Profile

This is the user-adjusted profile. If edits are necessary to the Adjusted Profile, it is first required to lock the Adjust Profile at the Aggregation Intersection. This will prevent the Normal Value from recalculating to a value different than 1.00 (100%) when the adjustments are made. To determine the Aggregation Intersection, view the intersection displayed on the Approval worksheet. Once adjustments are made to the profile, the user must go to the Final Profile Worksheet and set the approval method to *Do Not Approve* flag for the adjusted profile intersection. The Approved Profile measure will update with these changes.

Approved Profile

The Approve Profile displays the approved profile values. If a profile intersection is set to **Do Not Approve**, no value will be displayed in this field. The system will automatically approve all profile intersections set to Approve Use System. If changes are made to the Adjusted Profile, the values in this measure updates once the Manually Approved flag is set to True in the Final Profile worksheet.

System Profile

A read-only measure that displays the system-generated profiles calculated at the final profile's profile intersection.

Profile Approval Worksheet

The Profile Approval Worksheet allows you to review and approve final profiles.

Figure 16–6 Final Profile Approval Worksheet

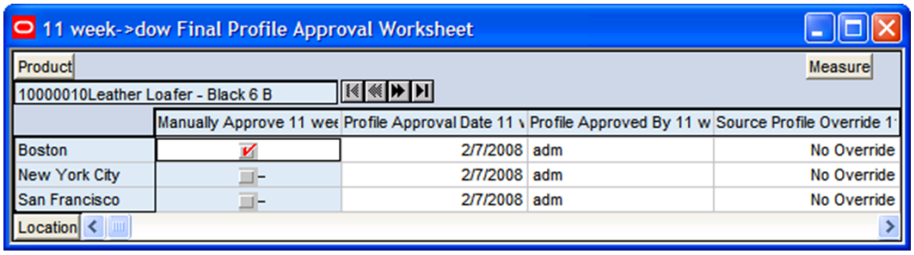

Measures: Profile Approval Worksheet

The Profile Approval Worksheet contains the following measures.

Profile Approved By

This measure displays who approved the profile for a given product/location combination. For all profile intersections with an Approval Method set to Approve Use System and no adjustment occurs, the Approved By measure will contain System.

Profile Approval Date

Displays the date on which a profile was approved, whether it is automatically approved by the system or manually approved by the user.

Manually Approve

A Boolean measure that must be activated (checked) for all profile intersections that are set to **Do not approve** or for intersections in which the user makes changes to the Adjusted Profile on the Final Profile worksheet.

Source Profile Override

This field displays the source-level that was used to generate the profile. If the Default Source Level was used, this measure will display No Override.

Source Profile Worksheet

The Source Profile Worksheet displays the profiles generated at the source-level for all product/location/calendar combinations selected to appear in the workbook. This worksheet displays the source system profiles (that is, the profiles that are calculated by the system during the profile generation process).

Figure 16–7 Source Profile Worksheet

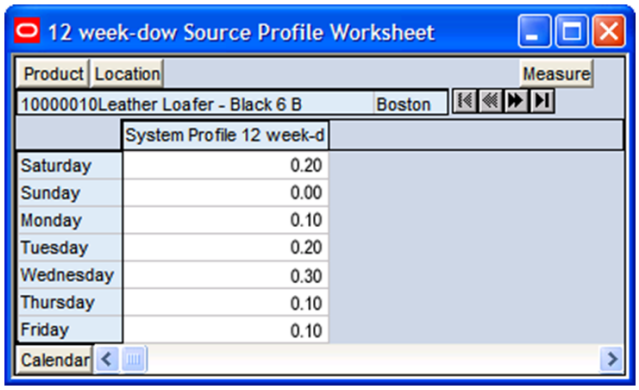

Measures: Source Profile Worksheet

The following measure is contained in the Source Level worksheet:

System Profile

A read-only measure that displays the system-generated profiles calculated at the source profile's profile intersection for each product/location combination displayed.

Generate Profiles

The batch profile generation process creates profile results for all hierarchy positions set in the Profile Intersection. Profiles may be run from the backend of the domain using the curvebatch executable or run manually using the Run Batch Profile workbook. For more information on curvebatch, refer to the *Oracle Retail Predictive Application Server Administrators Guide*.

Generating a Profile Manually

1. With the Local domain, select **New** from the **File** menu.
2. Select the **Curve** tab to display a list of workbooks. Select **Run Batch Profile** and click **OK**.

3. The Run Batch Profile Wizard opens and prompts you to select the profiles to generate. Select **Next** or **Finish**.

The Run Batch Profile wizard automatically executes curvebatch within the Local Domain. If Next is selected from the last wizard window, the wizard will not advance to the completion message until the profiles have been generated. Depending on the data set, this process may take a several minutes before the system advances to the final window.

17

Grade

Grade is a clustering tool that provides insight into how various parts of a retailer's operations can be grouped together. Typically, a retailer may cluster stores over item sales to create logical groupings of stores based upon sales of particular products. This provides increased visibility to where products are selling, and it allows the retailer to make more accurate decisions in merchandising.

Beyond this traditional use of clusters, Grade is flexible enough to cluster any business measure based on products, locations, time, promotions, customers, or any hierarchy configured in the solution.

Key Grade functionality includes:

- Two methods of creating Grades/Clusters:
  - Breakpoints–the sorting of data points into groups based on user-defined indexes
  - Clustering (or the BaNG) Algorithm–the optimization of data points into clusters based on the user-defined number of clusters
- Group By capabilities: support the segmentation of clusters for more detailed and focused cluster generation
- Clustering statistics: provide insight into the relationship of members within a cluster and how all clusters relate to one another
- Cluster What-if: allows user changes to members assigned to clusters and the review of recalculated clustering statistics

Regardless of the method employed to create clusters, Grade is designed to support the decision-making process necessary to create effective and actionable groupings of data. The following chapters describe the process to generate Grades/clusters and analyze results. All of Grade's functionality exists in the following workbooks and wizards:

- Breakpoints Administration Workbook
- Generate Breakpoint Grades Wizard
- Generate Clusters Wizard
- Cluster Review Workbook
- Delete Cluster Run Wizard

Information on the functionality of these workbooks is provided in the following sections.

Cluster Methods in Grade

Grade supports two cluster methods: Breakpoint and Batch Neural Gas Algorithm (BaNG).

Index to Average

Index to Average is the average sales of the stores in each grade divided by the average sales of all stores. This value provides a relative indication of how well a grade performed compared to the total store average. A value of 1.00 indicates that average sales in the grade were the same as average sales across the chain. A value of 3.74, for example, indicates that average sales in the grade were 3.74 times the chain average.

Breakpoint

Breakpoint method clusters, or groups, data points based on user-defined thresholds, or breakpoints. Breakpoints are properly defined by entering a decreasing sequence of real numbers that terminate in a value of zero (0). Breakpoints are entered in the Breakpoints Administration Workbook. Breakpoints are then used to cluster data points based on the index to average of each data point.

Example 17–1 Breakpoint

While clustering stores, if the user has entered the breakpoints 2.0, 1.5, 1.0, 0.5, and 0.0, the system will generate five grades. Stores that sell more than 2.0 times the average sales (over all stores) will be assigned to the highest grade, and stores that sell less than 0.5 times the average will be assigned to the lowest grade.

Batch Neural Gas Algorithm (BaNG)

The BaNG algorithm automatically generates optimal clusters based on user-specified number of clusters and clustering criteria. The algorithm provides a means for clustering data based on data distributions. For example, while clustering on weekly store sales data, the BaNG algorithm considers the Euclidean distance of the individual Store/week level data points from a cluster center to determine the clusters. This is different from the Breakpoint method, where clustering is performed based on average sales.

The BaNG algorithm iteratively updates cluster centers while considering the distance of each data vector from the cluster centers and its contribution to each cluster center. For every data point, cluster centers are ranked based on their distance from the data point within each iteration.

Additionally, the cluster centers are guided, using a control parameter, to gradually spread from the center of the distribution to their optimal locations.

The BaNG algorithm is a non-trivial extension of the K-means clustering approach. It is usually faster than the K-means and is guaranteed to converge.

BaNG versus Breakpoint

The BaNG algorithm generates statistically optimal clusters based on the number of clusters specified by the user. Breakpoint generates clusters based on user input breakpoints, and the number of clusters generated depends on the breakpoints.

In order to generate store clusters that vary by Dept, users need to specify a Group By option of Dept. Breakpoint will cluster stores based on the ratio of total store sales versus the average total store sales within each Group By intersection. BaNG can

consider an additional dimension for generating the clusters. For example, BaNG can cluster based on weekly sales of each stores within the Dept. Here, weekly sales are the coordinates over which the clustering is performed.

Grade in a Global Domain Environment

When implemented in a global domain environment, the following workbooks are available to be accessed from the Master domain:

- Breakpoint Administration

> **Note:** The Breakpoint Administration workbook is accessible from the Master domain to allow for the centralized administration of breakpoints.

The remaining Grade workbooks can only be accessed from local domains. These include:

- Cluster Review
- Delete Cluster Run
- Generate Breakpoint Grades
- Generate Clusters

Breakpoints Administration Workbook

The Breakpoints Administration workbook is used with the breakpoints method of grading. In this workbook, the user sets the index to average for each breakpoint. This includes the ability to set multiple breakpoint configurations to allow for Grades to be produced and compared using different breakpoint settings.

The Breakpoints Administration workbook includes this worksheet:

- Breakpoint Administration Worksheet

Creating a Breakpoints Administration Workbook

1. In the Local Domain, select **New** from the **File** menu.
2. Select the **Grade** tab to display a list of workbooks.
3. Select **Breakpoints Administration**. Click **OK**.

Breakpoint Administration Worksheet

If the breakpoints method is to be used to generate grades, you must go through the process of setting of the Index to Average for each grade's breakpoint range. Figure 17–1 is an example of the Breakpoint Administration worksheet.

Figure 17–1 Breakpoint Administration Worksheet

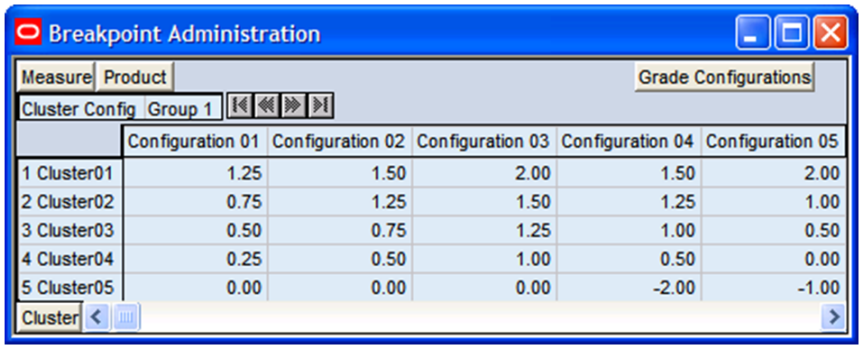

On the Breakpoint Administration worksheet, for a Configuration/Cluster intersection, set the Index to Average to be used by the Breakpoints algorithm for sorting data points into grades. The number of configurations available in this worksheet is based on the Maximum Number of Clusters configured in the Grade Plug-In. Only one configuration is required for use with the breakpoints method.

Breakpoints should be set from high Index to Average to low Index to Average, starting with the first Cluster ordered in the list of available clusters.

Note: The Cluster Labels may vary based on the configuration.

In Figure 17–1, the Index to Average is set as the following for Configuration 01:

- 1 Cluster01: 1.25
- 2 Cluster02: 0.75
- 3 Cluster03: 0.50
- 4 Cluster04: 0.25
- 5 Cluster05: 0.00

Using this example, the Breakpoints algorithm groups data based on the following:

Data	Description
1 Cluster01	Sort all data with an Index to Average at or above 1.25 into 1 Cluster01.
2 Cluster02	Sort all data with an Index to Average from 0.75 to 1.24 into 2 Cluster02.
3 Cluster03	Sort all data with an Index to Average from 0.50 to 0.74 into 3 Cluster03.
4 Cluster04	Sort all data with an Index to Average from 0.25 to 0.49 into 4 Cluster04.
5 Cluster05	Sort all data with an Index to Average from 0.00 to 0.24 into 5 Cluster05.
Junk Cluster	Using the previous example, all data points with an Index to Average that is less than 0.00 will be sorted into the Junk Cluster.
No Cluster	Any data points will null values in history (no loaded history) will not be graded. Refer to the Cluster Review Workbook for more information on Cluster Membership results.

Generate Breakpoint Grades Wizard

The Generate Breakpoint Grades wizard allows you to range the input data and hierarchies that are used to produce grades based on the Breakpoints method.

Opening the Generate Breakpoint Grades Wizard

1. Within the Local Domain, select **New** from the **File** menu.
2. Select the **Grade** tab to display a list of workbooks.
3. Select **Generate Breakpoints Grades**. Click **OK**.

Using the Generate Breakpoint Grades Wizard

The following steps outline the wizard process required to use the Generate Breakpoint Grades wizard:

1. Select **Source Measure for Cluster Analysis**.

   This wizard window prompts the user to select the measure that will be used to generate clusters. The options include all source measures defined at the time of the configuration.

2. Select **Hierarchy and Dimension to Cluster On**.

   The Cluster On dimension is the dimension that will be clustered. For example, while generating store clusters, store would be the Cluster On dimension. All positions in this dimension are sorted into grades based on the Breakpoints Configuration that will be selected later in this wizard process.

   Note: The dimension chosen should be at or higher than the base intersection of the selected Source Measure.

3. **Optional**: Select **Hierarchy and Dimension to Group By**.

   By Group is an optional setting that is used to partition data along a hierarchy dimension. This functionality allows for grades to be further segmented along multiple hierarchy dimensions. Grades are generated independently for each position within the selected dimensions. The dimensions displayed are equal to or higher than the Cluster On dimension.

   For example, while clustering stores by Department, that is, cluster stores by a Department at a time, the Cluster By Group is Department.

4. Select **Products**.

   Select the products for which you want to generate clusters.

5. Select **Locations**.

   Select the locations for which you want to generate clusters.

6. Select **Calendar**.

   Select the time periods that you want to be considered when generating clusters.

7. Select **Configuration Name**.

   Select the Breakpoint Configuration that will be used to produce the grades. These configurations must be set prior to generating grades using the Breakpoint Administration workbook.

8. Set **Grade Run Name**.

   Assign a label that will be used to identify the clustering run.

9. Select **Next** or **Finish**.

   If **Next** is selected, Break Point Run Succeeded is displayed once the grading process is completed. If **Finish** is selected this message will be skipped.

Generate Clusters Wizard

The Generate Clusters wizard allows you to range the input data and hierarchies that is used to produce clusters based on the Clustering (BaNG) method.

Opening the Generate Clusters Wizard

1. Within the Local Domain, select **New** from the **File** menu.
2. Select the **Grade** tab to display a list of workbooks.
3. Select **Generate Clusters**. Click **OK**.

Using the Generate Clusters Wizard

The following steps outline the wizard process required to use the Generate Clusters wizard:

1. Select **Source Measure for Cluster Analysis**.

   This wizard window prompts the user to select the measure that will be used to generate clusters. The options include all source measures defined at the time of the configuration.

2. Select **Hierarchy and Dimension to Cluster On**.

   The Cluster On dimension is the dimension that will be clustered. For example, while generating store clusters, store would be the Cluster On dimension.

   Note: The dimension chosen should be at or higher than the base intersection of the selected Source Measure.

3. Select **Hierarchy and Dimension to Cluster Over**.

   The Cluster Over dimension allows the user to define the dimension that will be used for clustering. The algorithm uses the positions in this dimension as the co-ordinates when clustering. Internally, data will be aggregated to this level before performing the clustering. For example, when using item/store/week level sales data to generate store clusters, the user has the option of specifying a dimension, such as Week, to which sales data will be aggregated. In this case, the algorithm considers Store/Week positions and clusters stores based on their weekly sales.

   Note: The dimension chosen should be at or higher than the base intersection of the selected Source Measure.

Note: The hierarchies presented in this wizard window include all hierarchies associated with the base intersection of the source measure except the hierarchy and dimension chosen to Cluster On.

4. **Optional**: Select **Hierarchy and Dimension to Group By**.

   By Group is an optional setting that is used to partition data along a hierarchy dimension. This functionality allows for clusters to be further segmented along multiple hierarchy dimensions. Clusters are generated independently for each position within the selected dimensions. The dimensions displayed are equal to or higher than the Cluster On and Cluster Over dimensions.

   For example, while clustering stores by Department, that is, cluster stores by a Department at a time, the Cluster By Group is Department.

5. Select **Products**.

   Select the products for which you want to generate clusters.

6. Select **Locations**.

   Select the locations for which you want to generate clusters.

7. Select **Calendar.**

   Select the time periods that you want to be considered when generating clusters.

8. Set **Name of Clusters**.

   Assign a label that will be used to identify the clustering run.

9. Set **Number of Clusters**.

   Define the number of clusters that are generated during the cluster generation process. If a Group By dimension was selected, this is the number of clusters that will be generated for each data partition within the dimension. The number of clusters that may be generated is based on the maximum number of clusters configured in the Grade configuration. Refer to the *Oracle Retail Grade Configuration Guide* for more information on Grade solution configuration.

10. Select **Next** or **Finish**.

    If **Next** is selected, Cluster Run Succeeded displays once the grading process is completed. If **Finish** is selected, this message is skipped.

Cluster and Breakpoint Grade Review

The Cluster Review workbook is a view of grade or cluster results and statistics. This workbook supports Grade/Cluster What-if, the re-assignment of members to Grades/Clusters on the fly, and the recalculation of clustering statistics. In addition, measures from different cluster runs may be inserted into the workbook to compare results.

Opening the Cluster Review Workbook

1. Within the Local Domain, select **New** from the **File** menu.
2. Select the **Grade** tab to display a list of workbooks.
3. Select **Cluster Review**. Click **OK**.

Cluster Review Wizard

The following sections outline the wizard process that is required to review Grade and Cluster results using the Cluster Review wizard:

Select Grade Birth to Review

Select the Grade Birth you want to review.

Select Additional Measures to Include in the Workbook

This wizard window displays all measures in the domain with the insertable measure property set to True. All measures generated as part of the clustering process for the birth date selected in the previous step are included in the base workbooks. These measures should not be selected in this wizard; however, measures associated with other cluster runs may be selected for comparison (for example, to compare results using Breakpoints vs. the Clustering method). To compare cluster results in this workbook, it is required that the cluster runs measures are generated with the same Cluster On and Group By dimensions.

Optional Additional Measures

The following additional measures can be added to the workbook by selecting them from the Additional Measures wizard window, in the Cluster Review Wizard.

Note: Once a cluster has been generated, all elements that have clustered together are referred to as members of the cluster in the following documentation. For example, while clustering stores, all stores that have been assigned to a cluster are its members.

Measure	Description
Aggregation Method	This is a read only measure that displays the Aggregation method used to aggregate Source data to the cluster over dimension for generating the clusters.
Cluster By Group Intersection	This is chosen by the user while generating the cluster run. For example, while clustering stores by time period, such as quarter, the Cluster By Group is Qrtr (within the Calendar hierarchy). This means that clusters are generated one quarter at a time.
Cluster Method	The method used for clustering.
Cluster Run Name	The name associated with the cluster run, entered by the user at cluster generation time.
Dimension to Cluster	The Cluster Over dimension is the dimension that was used for clustering. The algorithm uses the positions in this dimension as the co-ordinates when clustering. Internally, data will be aggregated to this level before performing the clustering. For example, when using item/store/week level sales data to generate store clusters, the user has the option of specifying a dimension, such as Week, that will be used to aggregate the sales data to. In this case, the algorithm considers Store/Week positions and clusters stores based on their weekly sales. **Note**: The dimension chosen should be at or higher than the base intersection of the selected Source Measure.
Mask Array for Clustering	A Boolean mask array indicating whether or not the element was used in the clustering run.

Cluster Review Workbook

The Cluster Review Workbook contains the following workbooks and worksheets:

Workbook	Worksheet
Cluster Results workbook tab	▪ Cluster Membership Worksheet ▪ Cluster Statistics Worksheet ▪ Cluster Centroid Statistics Worksheet ▪ Source Data Worksheet
Cluster Input Summary workbook tab	▪ Cluster Input Summary Worksheet

Gen ID

All measures that are specific to a given cluster run will have a Gen ID appended to them. This will allow users to store and review results from multiple cluster runs. Since the Cluster Review workbook allows users to review cluster results for a given cluster, most of the measures in this workbook will be appended by a Gen ID.

Cluster Membership Worksheet

Figure 17–2 Cluster Membership Worksheet

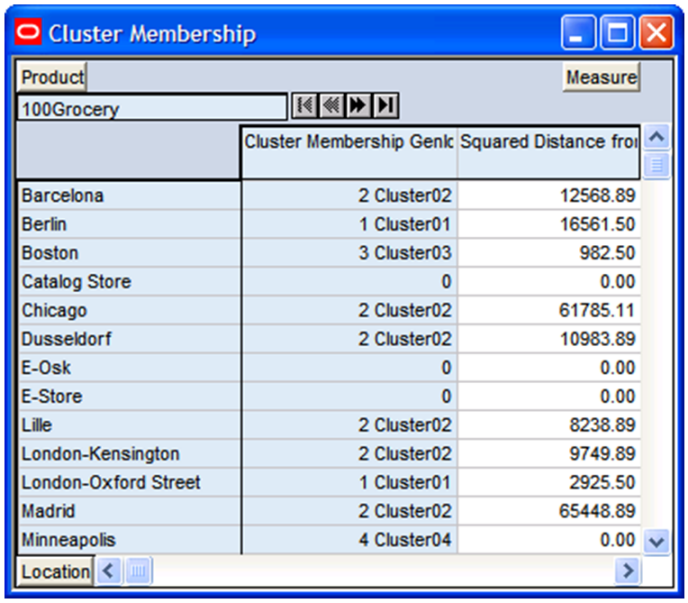

Measures	Description
Cluster Membership	Displays the positions that are assigned to a cluster or grade. Positions may be reassigned to another cluster and the cluster statistics will recalculate based on these user changes. If changes are made to Cluster Membership, the data must be committed for changes to be stored.

Measures	Description
Squared Distance from Centroid	A measure that quantifies how close an element is to its statistical centroid. It is calculated as the square of the Euclidian distance of all each elements being clustered from its centroid. One example of a situation where the Squared Distance from Centroid measure would be used is if the user were to cluster stores using the item/store/week level data, with a cluster over dimension as week (that is, the user clusters stores based on weekly store level sales). In this example, the Squared Distance from Centroid is the square of the Euclidian distance of each store from its centroid, where weekly sales form the coordinates, or dimensions, over which this distance is computed. **Note**: Euclidian distance is a metric that quantifies the straight line distance between two data points. It is the square root of the differences between coordinates of the two data points.

Cluster Statistics Worksheet

Figure 17–3 Cluster Statistics Worksheet

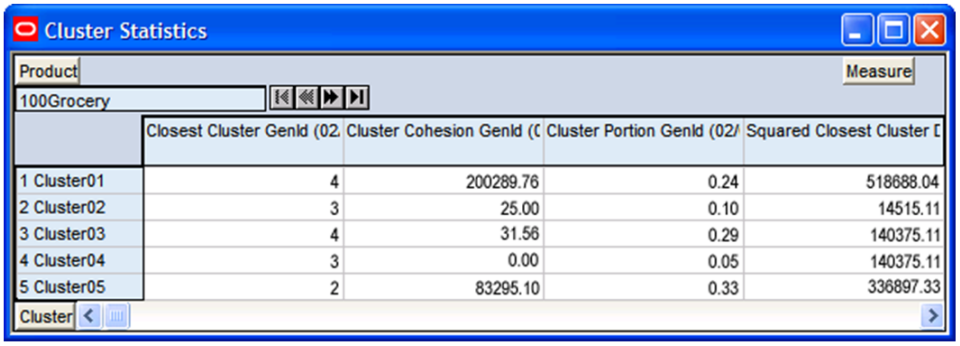

Measures	Description
Closest Cluster	This a read only measure that displays for each cluster the cluster ID that is statistically closest to its centroid. Valid only for clusters generated using the Clustering (BaNG) method.
Cluster Cohesion	A measure that quantifies how tight the cluster is, that is, how similar elements in the cluster are. For example, clustering stores using the item/store/week level, with a cluster over dimension as week, that is, cluster Stores based on weekly store level sales. First, calculate the Store/week level sales. Next, calculate the Squared Euclidian distance between each Store and its cluster centroid. Finally, calculate the average of this Squared Euclidian distance across all stores that belong to this cluster. Valid only for clusters generated using the Clustering (BaNG) method.
Cluster Portion	The ratio of the number of members in a cluster compared to all members being clustered. For example, while clustering stores, this is the ratio of the number of stores in a cluster to the total number of stores being clustered. Valid only for clusters generated using the Clustering (BaNG) method.
Squared Closest Cluster Distance	The squared Euclidian distance of all data points used for clustering, from the centroid of its nearest cluster. Valid only for clusters generated using the Clustering (BaNG) method.

Cluster Centroid Statistics Worksheet

Figure 17–4 Cluster Centroid Statistics Worksheet

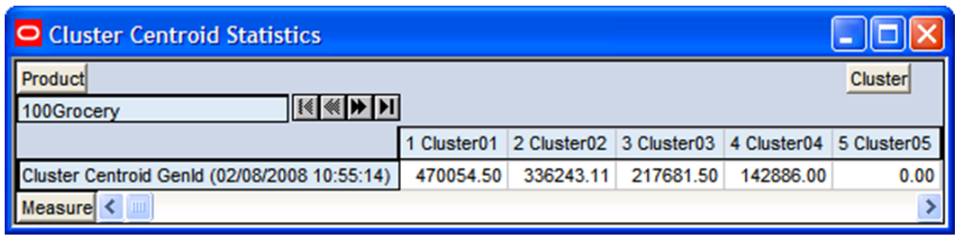

Measures	Description
Cluster Centroid	The average of the source data measure calculated across all cluster members; calculated by first aggregating the source data to the cluster on dimension for each member. For example, while clustering stores using the item/store/week level sales history, the centroid is calculated as the average of the total historic sales of each store in the cluster.
Cluster Centroid to Average Ratio	The ratio of the cluster centroid, to the average of all data points used for the cluster run; calculated by first aggregating the source data to the cluster on dimension for each member. For example, while clustering stores using the item/store/week level sales history, the Cluster to Average Ratio is calculated as the ratio of the Cluster Centroid to the average of the total historic sales for all stores being clustered. Valid only for clusters generated using the Clustering (BaNG) method.

Source Data Worksheet

Figure 17–5 Source Data Worksheet

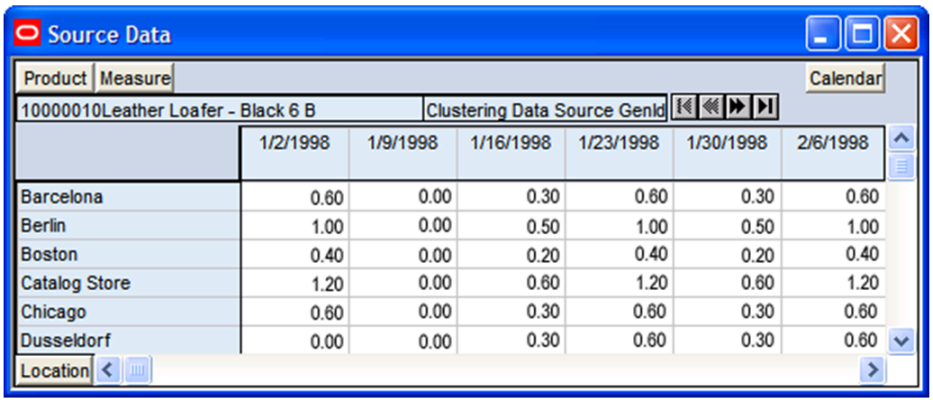

Measure	Description
Clustering Data Source	The data used to generate the clusters or grades.

Cluster Input Summary Worksheet

The Cluster Input Summary worksheet is a view to the settings used to generate the grades/clusters being reviewed in the workbook.

Figure 17–6 Summary Worksheet

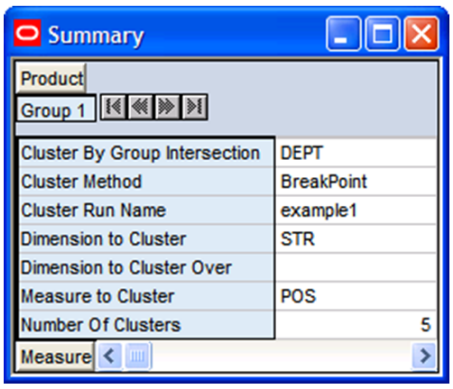

Measures	Description
Cluster By Group Intersection	Displays Group By dimensions if selected during the cluster generation process. For example, while clustering stores by time period, such as quarter, the Cluster By Group is Qrtr (within the Calendar hierarchy). This means that clusters are generated one quarter at a time.
Cluster Method	Displays the method used to generate the grades or clusters. Breakpoint or BaNG are valid methods.
Cluster Run Name	Displays the name assigned by the user to the cluster run at the time the clusters were generated.
Dimension to Cluster	Displays the Cluster Over dimension selected by the user during the cluster generation process. The algorithm uses the positions in this dimension as the coordinates when clustering. Internally, data will be aggregated to this level before performing the clustering. For example, when using the item/store/week level sales data to generate store clusters, the user has the option of specifying a dimension, such as Week, that will be used to aggregate the sales data to. In this case, the algorithm considers Store/Week positions and clusters stores based on their weekly sales. A value will only be displayed in this field if the BaNG method was used.
Measure to Cluster	Displays the measure selected as the clustering Data Source during the cluster generation process.
Number of Clusters	Displays the number of clusters used during the cluster generation process.

Delete Cluster Run Wizard

The Delete Cluster Run wizard allows you to delete clusters from the system based on Cluster Run Label and generation date (birth date).

Opening the Delete Cluster Run Wizard

1. Within the Local Domain, select **New** from the File menu.
2. Select the Grade tab to display a list of workbooks.
3. Select Delete Cluster Run. Click **OK**.

Using the Delete Cluster Run Wizard

The following steps outline the wizard process required to use the Delete Cluster Run wizard:

1. Select **Cluster Run to Delete**.

   This wizard window is a single select pick-list that includes the Cluster Run Label and the date/time stamp (birth date) for all cluster runs currently stored in the system. Select the Cluster Run Label to be deleted.

2. Select **Next** or **Finish**.

   If **Next** is selected, *Delete Cluster Run Succeeded* displays once the cluster deletion process is completed.

A

Appendix: Preprocessing

Preprocessing is a filtering module that automatically adjusts historical data to correct data points that do not represent general demand pattern. Essentially, it smoothes out spikes and dips in historical sales data, replacing stock-out data and data from short term events, such as promotions and temporary price changes, with data points that more accurately represent typical sales for that period. By adjusting the historical sales, Preprocessing can provide smarter data to the RDF Causal Engine, thus creating a smarter baseline forecast.

> **Note:** There are no workbooks associated with Preprocessing - it is available as a configuration option.

Common Preprocessing corrections are:

- Out of stock - Interfaced from RMS, weekly or daily
- Outliers - Indicator not required, depends on method
- Short term events - Promotions, temporary price changes

For example, Figure A–1 illustrates how Preprocessing adjusts for stock-outs.

Figure A–1 Preprocessing for Stock-outs

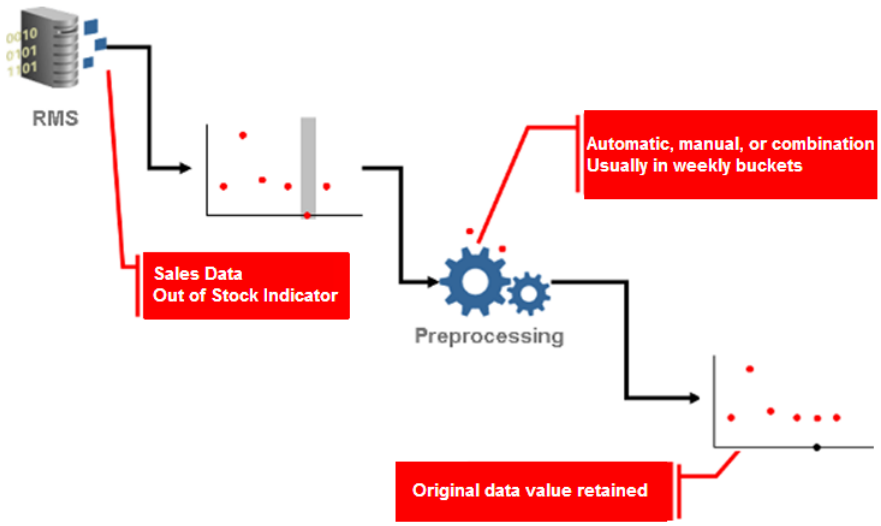

In Figure A–1, RMS sends historical sales data to the Preprocessing module of RDF. In that sales data, RMS has flagged out-of-stock instances with indicators (the gray portion of the first data set). Preprocessing takes note of that out-of-stock indicator and adjusts the sales for that time period to reflect a more typical sales quantity, taking into account trending and seasonality. Note in Figure A–1 that Preprocessing has removed the dip in sales in the second data set and has replaced it with a new data point.

Preprocessing Methods

Note: In order to run any preprocessing method, there needs to be at least three periods with non-zero data in the preprocessing window. If there are less than three periods with non-zero data, then the time series is skipped.

Preprocessing uses several methods to massage historical data. The following sections detail these methods:

- Standard Median
- Retail Median
- Standard Exponential Smoothing
- Lost Sales Standard Exponential Smoothing
- Override
- Increment
- Forecast Sigma
- Forecast Sigma Event
- DePrice
- Clear
- No Filtering

Standard Median

Standard Median calculates baselines on long time ranges.

Input: none

Optional parameter: window length

Figure A–2 Standard Median Formula

$$LSOVER(t) = median\ value\ over\left(t - \frac{window}{2},\ t + \frac{window}{2}\right)$$

Figure A–3 Standard Median Example

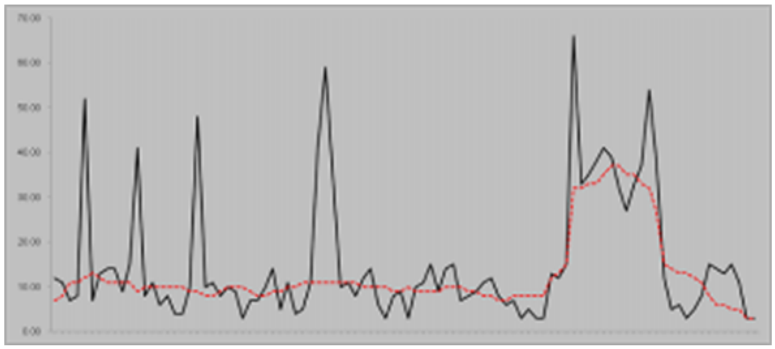

When data points for the full window are not available, Preprocessing pads the beginning and end of the time series with the first and the last data points, respectively, so that there are values for the full window.

Retail Median

Retail Median calculates baselines on long time ranges and improves side effects by making five standard median filter passes.

Input: none

Optional parameter: window length

Figure A–4 Retail Median Formula

$$Median1 = StdMedian(Src)$$
$$Median2 = StdMedian(Median1)$$
$$Diff1(t) = Median2(t) - Median2(t\text{-}1)$$
$$Median\_Diff1 = StdMedian(Diff1)$$
$$Smooth1(t) = Smooth1(t\text{-}1) + Median\_Diff1(t)$$
$$Diff2(t) = Smooth1(t) - Smooth1(t\text{-}1)$$
$$Avg\_Diff2 = StdMedian(Diff2)$$
$$Smooth2(t) = Smooth2(t\text{-}1) + Avg\_Diff2(t)$$

Figure A–5 Retail Median Example

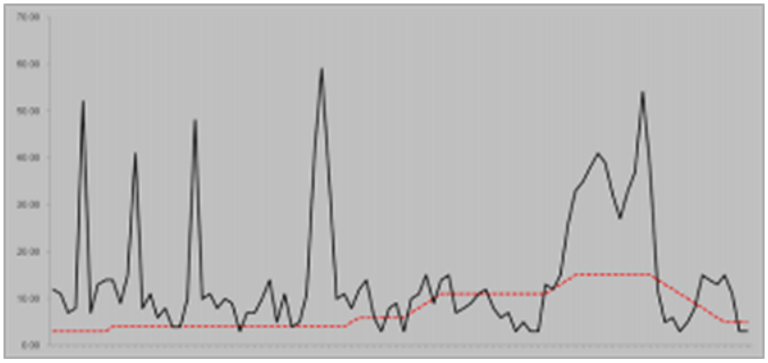

Standard Exponential Smoothing

Standard Exponential Smoothing removes spikes (such as promotional promo, temporary price changes, and so on), as well as filling the gaps (out-of-stock, unusual events such as a fire or hurricane).

Input: event indicator

Optional parameters:

The following table details the optional parameters for Standard Exponential Smoothing.

Optional Parameters	Description
ES (Exponential Smoothing)	The alpha parameter that determines the weight put on observations of periods included in the calculations.
Number of future periods (nfut)	The number of periods after an outage periods that are considered in the calculation of the future velocity. Note that if during these periods an event flag or a event indicator is on, the particular period is excluded from the calculation.
Number of past periods (npast)	The number of periods before an outage periods that are considered in the calculation of the past velocity. Note: When calculating the past velocity and the first period in the preprocessing window is flagged, then the past velocity is calculated using earlier periods outside the preprocessing window. Note that if during these periods an event flag or a event indicator is on, the particular period is excluded from the calculation.
Event flag	This parameter indicates if a period should be excluded from the calculation of past/future velocities.
Stop at event flag	This parameter determines which periods are included in the calculation of past/future velocities. If the flag is set to True, then the algorithm only includes periods before the first event flag or event indicator. If the flag is False, then all available, non-flagged periods, within the windows defined by nfut and npast, are used in the calculation of the past and future velocities. The default setting for the flag is False.

Figure A–6 Standard Exponential Smoothing Formula

$$past\ velocity = \frac{\sum_{i=1}^{np}(1-\alpha)^{i-1} * src(t_f - i)}{\sum_{i=1}^{np}(1-\alpha)^{i-1}}$$

$$future\ velocity = \frac{\sum_{i=1}^{nf}(1-\alpha)^{i-1} * src(t_l + i)}{\sum_{i=1}^{nf}(1-\alpha)^{i-1}}$$

$$LSOVER(t) = past\ velocity + \frac{future\ velocity - past\ velocity}{t_l - t_f + 2} * (t - t_f + 1),\ where\ \ t \in [t_f, t_l]$$

Figure A–7 Standard Exponential Smoothing Example

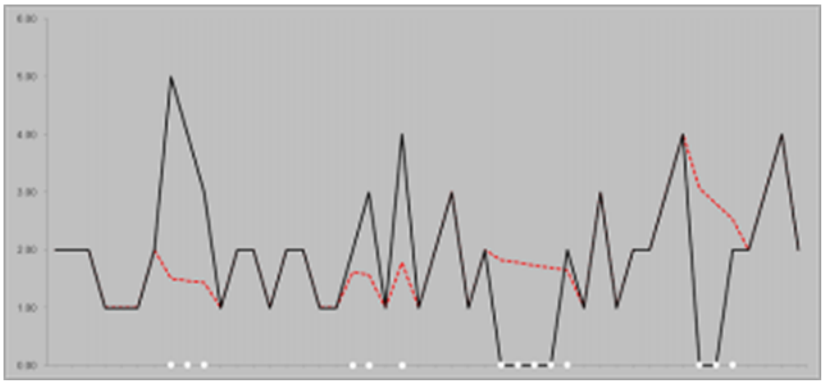

When event flags exist within the future and past velocity windows, rather than consider the entire window, Preprocessing only considers unflagged data points after the last event flag in the history window to compute the past velocity. It does a similar process for the future window by using the unflagged data points prior to the first event flag in the future window to compute the future velocity. Consecutive events are smoothed using the same velocities. A data point becomes flagged, and hence not part of the future/past velocity calculation, if either the event indicator or the optional event flag are on.

If future velocities cannot be calculated, then the past velocities, if they exist, are used as future and past velocities, and vice versa. When neither of the velocities can be calculated, there is no adjustment.

If the velocity window contains all zero values, then the calculated velocity is zero. A velocity of zero is a legitimate value if it occurs within the selling window. A velocity of zero is not acceptable if it is calculated based on values outside of the selling window.

Lost Sales Standard Exponential Smoothing

Lost Sales Standard Exponential Smoothing calculates baselines on long time ranges. Lost Sales Standard Exponential Smoothing makes positive adjustments to the flagged periods and to the period immediately following the flagged period.

Input: out-of-stock indicator

Optional parameters: ES (Exponential Smoothing) parameter and a number of future and past periods used to calculate the future/past velocities, an event flag, and a scalar parameter indicating if the period immediately following an out-of-stock period should be adjusted.

Figure A–8 Lost Sales Standard Exponential Smoothing Formula

$$past\ velocity = \frac{\sum_{i=1}^{np} (1-\alpha)^{i-1} * src(t_f - i)}{\sum_{i=1}^{np} (1-\alpha)^{i-1}}$$

$$future\ velocity = \frac{\sum_{i=1}^{nf} (1-\alpha)^{i-1} * src(t_l + i)}{\sum_{i=1}^{nf} (1-\alpha)^{i-1}}$$

$$LSOVER(t) = \max\{src(t),\ past\ velocity + \frac{future\ velocity - past\ velocity}{t_l - t_f + 2} * (t - t_f + 1)\},\ where\ t \in [t_f, t_l]$$

Figure A–9 Lost Sales Standard Exponential Smoothing Example

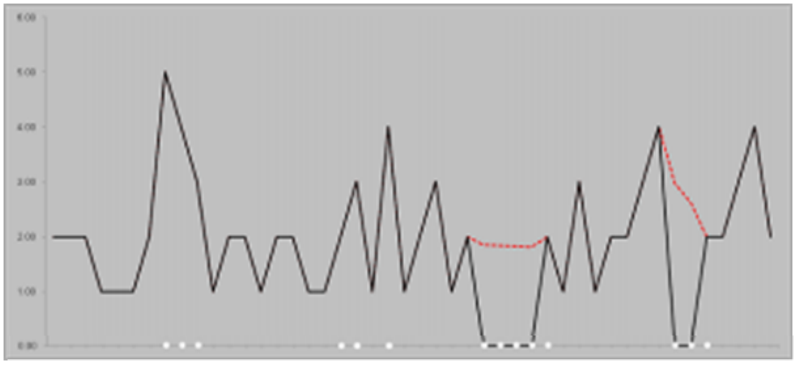

When event flags exist within the future and past velocity windows, rather than consider the entire window, Preprocessing only considers unflagged data points after the last event flag in the history window to compute the past velocity. It does a similar process for the future window by using the unflagged data points prior to the first event flag in the future window to compute the future velocity. Consecutive events are smoothed using the same velocities. A data point becomes flagged, and hence not part of the future/past velocity calculation, if either the event indicator or the optional event flag is on.

If future velocities cannot be calculated, then the past velocities, if they exist, are used as future and past velocities, and vice versa. When neither of the velocities can be calculated, there is no adjustment.

If the velocity window contains all zero values, then the calculated velocity is zero. A velocity of zero is a legitimate value if it occurs within the selling window. A velocity of zero is not acceptable if it is calculated based on values outside of the selling window. Note that by default, the periods being adjusted are the periods flagged by an out-of-stock indicator and the period immediately following any such period. If the optional scalar parameter POA (Partial Outage Allowed) is set to False, then this extra period will not be adjusted, and only the out-of-stock periods will be adjusted.

Override

Override fills gaps in data when a reference measure exists.

Input: reference measure (R(t)) to copy data from

Optional parameter: outage/mask (M(i)), adjustment ratio (a)

Formula: Overrides *LSOVER* with the Src adjusted by the adjustment ratio according to the mask:

Figure A–10 Override Formula

$$LSOVER(t) = a*R(t) \quad \text{if } M(i) \text{ is true}$$

$$LSOVER(t) = Src(t) \quad \text{if } M(i) \text{ is false}$$

Figure A–11 Override Example

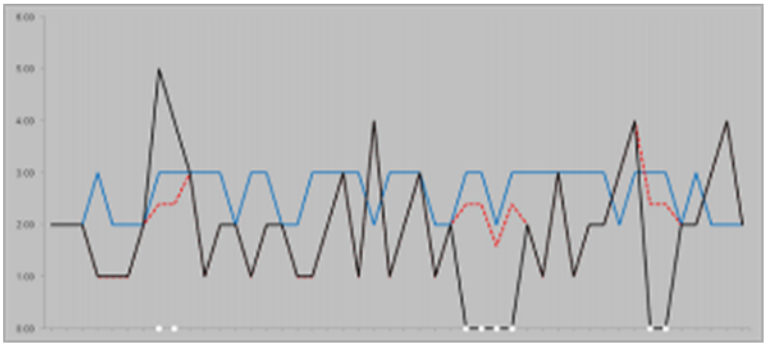

Increment

Increment updates gaps or outliers in data when a reference measure exists.

Input: reference measure (R(t)) to copy data from

Optional parameter: outage/mask (M(i)), adjustment ratio (a)

Increments the Src with the reference adjusted by the adjustment ratio according to the mask:

Figure A–12 Increment Formula

$$LSOVER(t) = SRC(t) + a*R(t) \quad \text{if } M(i) \text{ is true}$$

$$LSOVER(t) = Src(t) \quad \text{if } M(i) \text{ is false}$$

Figure A–13 Increment Example

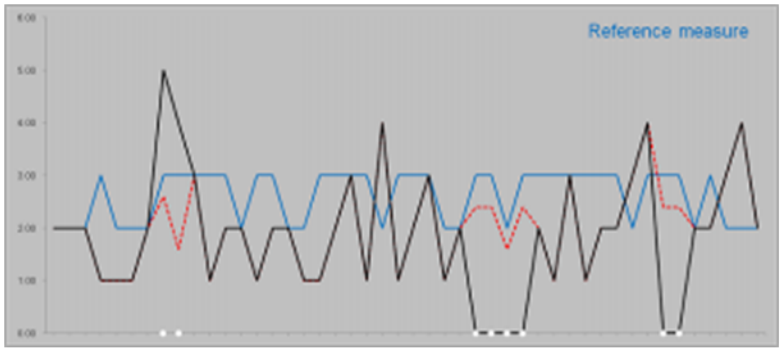

Forecast Sigma

Forecast Sigma removes spikes in recent data when no indicators are available.

Inputs: forecast and confidence intervals

Optional parameters:

- number of stddev for upper bound
- number of stddev for lower bound
- forecast lower bound
- minimum history required for filtering

Formula: If the difference in the sales and forecast is larger than a threshold, the override value is brought within some bounds of the forecast.

Forecast Sigma Event

Like Forecast Sigma, Forecast Sigma Event removes spikes in recent data when no indicators are available but also takes outage as input.

Inputs: outage, forecast, confidence intervals

Optional parameters:

- number of stddev for upper bound
- number of stddev for lower bound
- forecast lower bound
- minimum history required for filtering

Formula:

If outage is on:

LSOVER = forecast

Otherwise, if the difference in the sales and forecast is larger than a threshold, the override value is brought within some bounds of the forecast.

DePrice

DePrice removes the pricing effects.

Inputs: price, maximum price

Optional parameters: none

Formula:

*Smoothed = original * (price/maxprice)* ^2

Clear

This Preprocessing method clears the Preprocessing adjustments from previous runs and also clears the lsover measure.

LSOVER(t) = 0

LS(t) = 0

No Filtering

This Preprocessing method does not filter the source data.The preprocessing adjustments are cleared and lsover is set to the source data.

LS(t) = 0

LSOVER(t) = SRC(t)

Preprocessing for Stock-outs

When using Preprocessing to correct for stock-outs, the system expects out-of-stock indicators from a merchandising system like RMS. The system can be set up for automatic adjustment of sales history to correct for stock-outs as well as for manual user overrides under exception cases.

When set to automatically adjust sales history to correct stock-outs, Preprocessing takes into account trending and seasonality and adjusts the sales that were flagged by the out-of-stock indicator to reflect a more typical sales quantity.

Preprocessing for Outliers

When using Preprocessing to correct for outliers, the system expects outlier indicators. These are typically loaded.

Preprocessing for Promotional Forecasting

Preprocessing adjusts promotional data in a similar way that it does stock-outs. Typically, historical data shows a higher rate of sales during promotional periods. Were these spikes in sales to be left in historical sales data and loaded in the RDF Causal Engine, the baseline forecast created from this data would reflect similar spikes in future sales.

Glossary

Note: With a few exceptions, this glossary contains definitions of terms specific to RDF. For further definitions of terms and concepts relating to the RPAS user interface, refer to the *Oracle Retail Predictive Application Server Online Help* or *Oracle Retail Predictive Application Server User Guide*.

Additive Seasonal Method

Also referred to as Additive Winters Model, this model is similar to the Multiplicative Winters model, but it 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 forecast value is higher than or lower than 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 or Automatic Exponential Smoothing Method

RDF 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:

- SimpleES
- IntermittentES
- TrendES
- Multiplicative Seasonal
- Additive Seasonal and SeasonalES

The final selection between the models is made according to a performance criterion (Bayesian Information Criterion) that involves a trade-off between the model's fit over the historic data and its complexity.

Bayesian Method

Useful for short-lifecycle 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 lifecycle) 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. Typically, the Causal method is used at the Final Levels (that is, item/week/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.

Croston's Model of Exponential Smoothing

See IntermittentES or 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. Its weight declines in data exponentially. The most recent data is weighted more heavily. It requires a smoothing constant ( ). It ranges from 0 to 1 and is 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

In RDF, Forecast refers to RDF's statistical forecasting capabilities.

Forecast-Driven Planning

Planning that keys off of forecasts fed directly into a planning system. Connection to RDF is built directly into the business process supported by Oracle Retail 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 Oracle Retail Predictive Planning.

Halo

Used to explain the bias shown by customers towards certain products because of a favorable experience with other products.

Holt's Model of Exponential Smoothing

See Trend Exponential Smoothing or TrendES.

Interactive Forecasting

A workbook in RDF that is used to simulate forecast by modifying parameters such as Forecast Method and History Start Date.

IntermittentES or Intermittent Exponential Smoothing

RDF 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 transitions from non-zero sales to zero points (that is, intermittent demand data). The original time series is split into a Magnitude and Frequency series, and then the SimpleES 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 or Like SKU

An item that is 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.

Moving Average

For each period *t*, the moving average method takes the average of the periods from *t*-2 to *t*+2 as the smoothed baseline.

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.

Preprocessing

In 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

RDF 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 deseasonalized and then fed to the SimpleES 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 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 used in the context of promotional forecasting. This workbook allows you to analyze the effects of promotions on items 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.

SeasonalES 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. Today is used as default end. Today minus Training Window Length is used as default start. Even if there is too little data to create a seasonal forecast (in general less than 52 weeks, or a full sales cycle) the system still attempts to generate a seasonal forecast if at all possible. The sales of last year are time shifted by 52 weeks, and become the forecast for the corresponding periods. If the third week in 2020 is in the forecast horizon, and sales for the third week in 2019 are available, they will become the forecast. If not, the forecast will be generated by running SimpleES with the available (less than 52 weeks) history. For this particular case, the method picked is Seasonal Regression.

Seasonal Regression

Seasonal Regression cannot be selected as a forecasting method, but it is a candidate model used when the SeasonalES 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/IntermittentES Method

A combination of the SimpleES and IntermittentES methods. This method applies the SimpleES model unless a large number of zero transitions from non-zero sales to zero points are present. In this case, the Croston's model is applied.

SimpleES or Simple Exponential Smoothing Method

RDF uses a simple exponential smoothing model to generate forecasts. SimpleES ignores seasonality and trend features in the demand data, and it 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.

Simple Moving Average

See Moving Average.

Sister Store

A store that is 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 continues influence in future.

Training Window

The number of weeks of historical sales data to use in generating a forecast.

Trend Exponential Smoothing or TrendES

Also referred to as Holt's Model, RDF 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, and then provides forecast point estimates by combining an estimated trend and the smoothed level at the end of the series.

Wizard

A set of windows that guide you through the process of creating a new workbook or performing other actions in a solution 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. Also see 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 Oracle Retail 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.