

ORACLE® CRYSTAL BALL, FUSION EDITION

RELEASE 11.1.2



Crystal Ball User's Guide, 11.1.2

Copyright © 1988, 2010, Oracle and/or its affiliates. All rights reserved.

Authors: EPM Information Development Team

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited. The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS:

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Chapter 1. Welcome	3
Introduction	3
Who Should Use This Program	4
What You Will Need	4
About the Crystal Ball Documentation Set	4
Screen Capture Notes	6
Getting Help	6
Technical Support and More	7
Chapter 2. Crystal Ball Overview	9
Introduction	9
Model Building and Risk Analysis Overview	9
What is a Model?	0
Risk and Certainty	0
How Crystal Ball Differs from Traditional Analysis Tools	1
Monte Carlo Simulation and Crystal Ball	2
Crystal Ball Feature Overview	4
Charts and Analysis Tools	4
Other Crystal Ball Tools	1
Process Capability Features	3
Trend Analysis with Predictor	3
Optimizing Decision Variable Values with OptQuest	3
Steps for Using Crystal Ball	4
Resources for Learning Crystal Ball	4
Starting and Closing Crystal Ball	4
Starting Crystal Ball Manually	5
Starting Crystal Ball Automatically	5
Crystal Ball Welcome Screen	5
Closing Crystal Ball	6
Crystal Ball Menus and Toolbar	6
The Crystal Ball Menus	6
The Crystal Ball Toolbar	7

	Crystal Ball Tutorials	38
Chapter 3. De	Defining Model Assumptions	39
	Introduction	39
	Types of Data Cells	39
	About Assumptions and Probability Distributions	40
	Defining Assumptions	40
	Entering Assumptions	40
	Additional Assumption Features	43
	Entering Cell References and Formulas	44
	Alternate Parameter Sets	45
	Fitting Distributions to Data	46
	Specifying Correlations Between Assumptions	49
	Setting Assumption Preferences	52
	Using the Crystal Ball Distribution Gallery	53
	Displaying the Distribution Gallery	53
	Distribution Gallery Window	53
	Managing Distributions	55
	Managing Categories	59
Chanter 4 De	Defining Other Model Elements	65
Oliaptei 4. De	Introduction	
	Defining Decision Variable Cells	
	Defining Forecasts	
	Setting Forecast Preferences	
	Working with Crystal Ball Data	
	Editing Crystal Ball Data	
	Selecting and Reviewing Data	
	Setting Cell Preferences	
	Saving and Restoring Models	
	0.1.1.0	
Chapter 5. Ru	Running Simulations	
	Introduction	
	About Crystal Ball Simulations	
	Setting Run Preferences	
	Setting Trials Preferences	
	Setting Sampling Preferences	
	Setting Speed Preferences	
	Setting Options Preferences	
	Setting Statistics Preferences	82

	Freezing Crystal Ball Data Cells
	Running Simulations
	Starting Simulations
	Stopping Simulations
	Continuing Simulations
	Resetting and Rerunning Simulations
	Single-Stepping
	The Crystal Ball Control Panel
	Managing Chart Windows
	Single Windows
	Multiple Windows
	Saving and Restoring Simulation Results
	Saving Crystal Ball Simulation Results
	Restoring Crystal Ball Simulation Results
	Using Restored Results
	Restored Results with Capability Metrics
	Functions for Use in Microsoft Excel Models
	Running User-Defined Macros
	User-Defined Macro Interfaces
	Priority Rules
	Global Macros
	Toolbar Macros
Chapter 6	. Analyzing Forecast Charts
	Introduction
	Guidelines for Analyzing Simulation Results
	Understanding and Using Forecast Charts
	Determining the Certainty Level
	Focusing On the Display Range
	Formatting Chart Numbers
	Changing the Distribution View and Interpreting Statistics
	Setting Forecast Preferences
	Setting Forecast Chart Preferences
	Using Additional Forecast Features
	Fitting a Distribution to a Forecast
	Defining Assumptions from Forecasts
	Setting Chart Preferences
	Setting Preferences with Shortcut Keys
	Basic Customization Instructions

	Specific Customization Instructions
	Managing Existing Charts
	Opening Charts
	Copying and Pasting Charts to Other Applications
	Printing Charts
	Closing Charts
	Deleting Charts
Chapter 7. Ai	nalyzing Other Charts
	Overview
	Selecting Assumptions, Forecasts, and other Data Types
	Understanding and Using Overlay Charts
	Creating Overlay Charts
	Customizing Overlay Charts
	Using Distribution Fitting with Overlay Charts
	Understanding and Using Trend Charts
	Creating Trend Charts
	Customizing Trend Charts
	Understanding and Using Sensitivity Charts
	About Sensitivity Charts
	Creating Sensitivity Charts
	How Crystal Ball Calculates Sensitivity
	Limitations of Sensitivity Charts
	Customizing Sensitivity Charts
	Understanding and Using Assumption Charts
	Customizing Assumption Charts
	Understanding and Using Scatter Charts
	Creating Scatter Charts
	Customizing Scatter Charts
Chapter 8. Ci	reating Reports and Extracting Data
	Introduction
	Creating Reports
	Basic Steps for Creating Reports
	Setting Report Options
	Defining Custom Reports
	Report Processing Notes
	Extracting Data
	Data Extraction Examples 168

Chapter 9.	Crystal Ball Tools
	Introduction
	Overview
	Setup Tools
	Analysis Tools
	Integration Tools
	Tools and Run Preferences
	Batch Fit Tool
	Using the Batch Fit Tool
	Batch Fit Example
	Correlation Matrix Tool
	About Correlations
	About the Correlation Matrix Tool
	Using the Correlation Matrix Tool
	Correlation Matrix Example
	Tornado Chart Tool
	Tornado Chart
	Spider Chart
	Using the Tornado Chart Tool
	Tornado Chart Example
	Limitations of the Tornado Chart Tool
	Bootstrap Tool
	Using the Bootstrap Tool
	Bootstrap Example
	Decision Table Tool
	Using the Decision Table Tool
	Decision Table Example
	Scenario Analysis Tool
	Using the Scenario Analysis Tool
	Scenario Analysis Example
	2D Simulation Tool
	Using the 2D Simulation Tool
	2D Simulation Example
	Second-Order Assumptions
	Data Analysis Tool
	Using the Data Analysis Tool
	Data Analysis Example

Appendix A.	Selecting and Using Probability Distributions
	Introduction
	Understanding Probability Distributions
	A Probability Example
	Continuous and Discrete Probability Distributions
	Selecting Probability Distributions
	Probability Distribution Descriptions
	Beta Distribution
	BetaPERT Distribution
	Binomial Distribution
	Custom Distribution
	Discrete Uniform Distribution
	Exponential Distribution
	Gamma Distribution
	Geometric Distribution
	Hypergeometric Distribution
	Logistic Distribution
	Lognormal Distribution
	Maximum Extreme Distribution
	Minimum Extreme Distribution
	Negative Binomial Distribution
	Normal Distribution
	Pareto Distribution
	Poisson Distribution
	Student's t Distribution
	Triangular Distribution
	Uniform Distribution
	Weibull Distribution
	Yes-No Distribution
	Using the Custom Distribution
	Custom Distribution Example 1
	Custom Distribution Example 2
	Custom Distribution Example 3 — Loading Data
	Entering Tables of Data into Custom Distributions
	Changes from Crystal Ball 2000.x (5.x)
	Other Important Custom Distribution Notes
	Truncating Distributions
	Be Aware
	Comparing Distributions

Ε	Distribution Parameter Summary
U	Using Probability Functions
	Limitations of Probability Functions
	Probability Functions and Random Seeds
S	equential Sampling with Custom Distributions
	Creating Custom SIP Distributions
	Running Simulations with SIPs
Appendix B. Using	the Extreme Speed Feature
C	Overview 293
C	Compatibility Issues
	Multiple-Workbook Models
	Circular References
	Crystal Ball Microsoft Excel Functions
	User-Defined Functions
	Running User-Defined Macros
	Special Functions
	Undocumented Behavior of Standard Functions
	Incompatible Range Constructs
	Data Tables
N	Numerical Differences and the Compare Run Modes Tool
C	Other Important Differences
	OptQuest and other Crystal Ball Tools
	Precision Control and Cell Error Checking
	Spreadsheet Updating
	Very Large Models
	Memory Usage
	Spreadsheets with No Crystal Ball Data
N	Maximizing the Benefits of Extreme Speed
	String Intermediate Results in Formulas
	Calls to User-Defined Functions
	Dynamic Assumptions
	Microsoft Excel Functions
Appendix C. Crysta	al Ball Tutorials
I	ntroduction
Т	Sutorial 1 — Futura Apartments 307
	Starting Crystal Ball
	Opening the Example Model
	Running Simulations

	Results Analysis — Determining Profit	311
	Take a Look Behind the Scenes	312
	Resetting and Single-Stepping	313
	Closing Crystal Ball	314
	Tutorial Review	314
	Tutorial 2 — Vision Research	315
	Starting Crystal Ball	315
	Opening the Example Model	315
	Reviewing the Vision Research Scenario	316
	Defining Assumptions	316
	Improving Process Quality — Tutorial 3	334
	Overall Approach	335
	Using the Loan Processing Model	337
	Packaging Pump Design — Tutorial 4	348
	Overall Approach	348
	Using the DFSS Liquid Pump model	351
	Using OptQuest to Optimize Quality and Cost	359
Appendix D	D. Using the Process Capability Features	365
	Introduction	365
	About Crystal Ball's Process Capability Features	365
	About Quality Improvement Methodologies	366
	Methodologies for Improving Processes — Six Sigma	367
	Methodologies for Improving Design — DFSS	367
	Methodologies for Adding Value — Lean Principles	367
	How Crystal Ball Supports Quality	368
	Working Through Process Capability Tutorials	370
	Preparing to Use Process Capability Features	370
	Activating the Process Capability Features	370
	Setting Capability Calculation Options	370
	Setting Specification Limits and Targets	372
	Analyzing Process Capability Results	372
	Viewing Capability Metrics	372
	Extracting Capability Metrics	375
	Including Capability Metrics in Reports	377
	Capability Metrics List	378
Appendix E	E. Maximizing the Use of Crystal Ball	381
	Introduction	381
	Simulation Accuracy	381

	Precision Control	82
	Sampling Method	83
	Simulation Speed	83
	Sample Size	84
	Correlated Assumptions	84
	Crystal Ball and Multi-processor Computers	85
	Crystal Ball and Multiple Processors	85
	Crystal Ball and Multi-threading	86
Appendix F	F. Accessibility	87
	Documentation Accessibility	
	Access to Oracle Support for Hearing-Impaired Customers	87
	Enabling Accessibility for Crystal Ball	88
	Activating and Using Accessibility Mode	88
	Activating Accessibility Mode	89
	A Note About Accessible Default Cell Preferences	89
	Using the Tab and Arrow Keys in the Crystal Ball User Interface	90
	Accessing Charts without a Mouse	90
	Tabbing Through the Crystal Ball Control Panel	91
	Using Microsoft Windows Large Fonts in Crystal Ball	91
	Controlling Screen Flicker when Resizing Charts	92
	Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2003 or Earlier 3	92
	Crystal Ball Toolbar 3	92
	Crystal Ball Menus	93
	Shortcut key Combinations in Microsoft Excel 2003 or Earlier	93
	Distribution Gallery Shortcut Keys	00
	Chart Preference Shortcut Keys	01
	Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2007 or Later 4	02
	Introduction	02
	Crystal Ball Ribbon in Microsoft Excel 2007 or Later	02
	Shortcut Key Combinations in Microsoft Excel 2007 or later	04
	Distribution Gallery Shortcut Keys	09
	Chart Preference Shortcut Keys	09
	Compatibility and File Conversion Issues	09
Appendix (G. Bibliography 4	11
	Bootstrap	11
	Monte Carlo Simulation	11
	Probability Theory and Statistics	12
	Random Variate Generation Methods	12

	Specific Distributions
	Extreme Value Distribution
	Lognormal Distribution
	Weibull Distribution
	Tornado Charts and Sensitivity Analysis
	Two-Dimensional Simulation
	Uncertainty Analysis
	Sequential Sampling with SIPs
Glossary	415
ndex	419

1

Welcome

In This Chapter

Introduction	13
Who Should Use This Program	14
What You Will Need	
About the Crystal Ball Documentation Set	14
Getting Help	16
Technical Support and More	17

Introduction

This Guide describes how to use all three current releases of Oracle Crystal Ball:

- Oracle Crystal Ball, Fusion Edition
- Oracle Crystal Ball Decision Optimizer, Fusion Edition
- Oracle Crystal Ball Enterprise Performance Management, Fusion Edition

Unless otherwise noted, when this Guide refers to Crystal Ball, the information applies to all three versions.

Crystal Ball is a graphically oriented forecasting and risk analysis program that takes the uncertainty out of decision-making.

Through the power of simulation, Crystal Ball becomes an effective tool in the hands of decision-makers. You can answer questions such as, "Will we stay under budget if we build this facility?" or, "What are the chances this project will finish on time?" or, "How likely are we to achieve this level of profitability?" With Crystal Ball, you will become a more confident, efficient, and accurate decision-maker.

Crystal Ball is easy to learn and easy to use. Unlike other forecasting and risk analysis programs, you do not have to learn unfamiliar formats or special modeling languages. To get started, all you have to do is create a spreadsheet. From there, this manual guides you step by step, explaining Crystal Ball terms, procedures, and results.

And you do get results from Crystal Ball. Through a technique known as Monte Carlo simulation, Crystal Ball forecasts the entire range of results possible for a given situation. It also shows you confidence levels, so you will know the likelihood of any specific event taking place.

The following sections explain more about Crystal Ball and how it works:

- "Who Should Use This Program" on page 14
- "What You Will Need" on page 14
- "About the Crystal Ball Documentation Set" on page 14
- "Getting Help" on page 16
- "Technical Support and More" on page 17

Who Should Use This Program

Crystal Ball is for decision-makers, from the analyst exploring the potential for new markets to the scientist evaluating experiments and hypotheses. Crystal Ball has been developed with a wide range of spreadsheet uses and users in mind.

You do not need highly advanced statistical or computer knowledge to use Crystal Ball to its full potential. All you need is a basic working knowledge of the personal computer and the ability to create a spreadsheet model.

What You Will Need

Crystal Ball runs on several versions of Microsoft Windows and Microsoft Excel. For a complete list of required hardware and software, see the system requirements list in the *Oracle Crystal Ball Installation and Licensing Guide*.

About the Crystal Ball Documentation Set

The Oracle Crystal Ball User's Guide is intended for students, analysts, engineers, executives, and others who want to learn how to use the main features of Crystal Ball. As mentioned earlier, unless otherwise noted, the Crystal Ball documentation pertains to all current Crystal Ball releases:

- Crystal Ball
- Crystal Ball Decision Optimizer
- Crystal Ball EPM

The Oracle Crystal Ball Enterprise Performance Management Integration Guide contains special Crystal Ball integration information for users of Crystal Ball EPM.

The Oracle Crystal Ball Installation and Licensing Guide describes how to install and license Crystal Ball.

For information about distribution defaults and formulas plus other statistical information, see the *Oracle Crystal Ball Statistical Guide*.

The Oracle Crystal Ball Predictor User's Guide, Oracle Crystal Ball Decision Optimizer OptQuest User's Guide, Oracle Crystal Ball Developer's Guide and Oracle Crystal Ball API for .NET Developer's Guide offer additional information about those Crystal Ball products. Note that the

Oracle Crystal Ball Decision Optimizer OptQuest User's Guide is only for users of Crystal Ball Decision Optimizer.

This Oracle Crystal Ball User's Guide includes the following additional chapters and appendices:

• Chapter 2, "Crystal Ball Overview"

Introduces Crystal Ball and explains how it uses spreadsheet models to help with risk analysis and many types of decision-making.

• Chapter 3, "Defining Model Assumptions"

Describes how to define assumption cells in models and how to use the Crystal Balll Distribution Gallery.

• Chapter 4, "Defining Other Model Elements"

Describes how to define decision variable cells and forecast cells in models. It also explains how to set cell preferences.

• Chapter 5, "Running Simulations"

Provides step-by-step instructions for setting up and running a simulation in Crystal Ball.

• Chapter 6, "Analyzing Forecast Charts"

Explains how to use Crystal Ball's powerful analytical features to interpret the results of a simulation, focusing on forecast charts.

• Chapter 7, "Analyzing Other Charts"

Provides additional information to help you analyze and present the results of the simulations using advanced charting features.

• Chapter 8, "Creating Reports and Extracting Data"

Provides information to help you share Crystal Ball data and graphics with other applications, and describes how to prepare reports with charts and data.

Chapter 9, "Crystal Ball Tools"

Describes tools that extend the functionality of Crystal Ball, such as the Tornado Chart and Decision Table tools.

Appendix A, "Selecting and Using Probability Distributions"

Describes all the pre-defined probability distributions used to define assumptions in Crystal Ball, and includes suggestions on how to choose and use them.

• Appendix B, "Using the Extreme Speed Feature"

Discusses the optional Extreme Speed feature available with Crystal Ball and describes its benefits and compatibility issues.

Appendix C, "Crystal Ball Tutorials"

Demonstrates Crystal Ball basics and shows how to use more advanced features in a variety of settings.

• Appendix D, "Using the Process Capability Features"

Discusses the process capability features that can be activated to support Six Sigma, DFSS, Lean principles, and similar quality programs.

• Appendix E, "Maximizing the Use of Crystal Ball"

Describes how to maximize the accuracy and speed of Crystal Ball simulations.

Appendix F, "Accessibility"

Describes Crystal Ball accessibility features including command equivalents and descriptions for each Crystal Ball toolbar button (Microsoft Excel 2003 and earlier) or Crystal Ball ribbon icon (Microsoft Excel 2007 or later) and provides more information about using Crystal Ball with Microsoft Excel 2007 or later and Windows Vista.

• Appendix G, "Bibliography"

Lists related publications, including statistics textbooks.

Glossary

Defines terms specific to Crystal Ball and other statistical terms used in this manual.

Screen Capture Notes

All the screen captures in this document were taken in Microsoft Excel 2003 for Windows XP Professional and Microsoft Excel 2003 for Windows XP, using a Crystal Ball Run Preferences random seed setting of 999, unless otherwise noted.

Due to rounding differences between various system configurations, you might obtain slightly different calculated results than those shown in the examples.

Getting Help

- To display online help in a variety of ways as you work in Crystal Ball:
- Click the Help button, , in a dialog.
- Click the Help tool in the Crystal Ball toolbar or ribbon in Microsoft Excel.
- In the Microsoft Excel menubar, choose Help, then Crystal Ball, then Crystal Ball Help.
- In the Distribution Gallery and other dialogs, press F1.

Note: In Microsoft Excel 2007 or later, click Help at end of the Crystal Ball ribbon. Note that if you press F1 in Microsoft Excel 2007 or later, Microsoft Excel help opens unless you are viewing the Distribution Gallery or another Crystal Ball dialog.

Tip: When help opens, the Search tab is selected. Click the Contents tab to view a table of contents for help.

Technical Support and More

Oracle offers a variety of resources to help you use Crystal Ball, such as technical support, training, and other services. For information, see:

http://www.oracle.com/crystalball

2

Crystal Ball Overview

In This Chapter

Introduction	19
Model Building and Risk Analysis Overview	19
Crystal Ball Feature Overview	24
Steps for Using Crystal Ball	34
Starting and Closing Crystal Ball	34
Crystal Ball Menus and Toolbar	36
Crystal Ball Tutorials	38

Introduction

This chapter presents the basics you need to understand, start, review the menus and toolbars, and close Crystal Ball. Now, spend a few moments learning how Crystal Ball can help you make better decisions under conditions of uncertainty.

Chapters 2 through 7 of this *Oracle Crystal Ball User's Guide* describe how to build, run, analyze, and present the own Crystal Ball model simulations.

Model Building and Risk Analysis Overview

Crystal Ball is an analytical tool that helps executives, analysts, and others make decisions by performing simulations on spreadsheet models. The forecasts that result from these simulations help quantify areas of risk so decision-makers can have as much information as possible to support wise decisions.

The basic process for using Crystal Ball is to:

- 1. Build a spreadsheet model that describes an uncertain situation.
- Run a simulation on it.
- 3. Analyze the results.

The following sections give a brief overview of risk analysis and modeling. They build a foundation for understanding the many ways Crystal Ball and related products can help you minimize risk and maximize success in virtually any decision-making environment.

What is a Model?

Crystal Ball works with spreadsheet models, specifically Microsoft Excel spreadsheet models. the spreadsheet might already contain a model, depending on what type of information you have in the spreadsheet and how you use it.

If you only use spreadsheets to hold data — sales data, inventory data, account data, and so on — then you do not have a model. Even if you have formulas that total or subtotal the data, you might not have a model that is useful for simulation.

A model is a spreadsheet that has moved from being a data organizer to an analysis tool. A model represents the relationships between input and output variables using a combination of functions, formulas, and data. As you add more cells to the model, the spreadsheet begins to portray the behavior of a real-world system or situation.

Risk and Certainty

The following sections describe risk and how Crystal Ball helps handle it:

- "About Risk" on page 20
- "Quantifying Risks with Spreadsheet Models" on page 20
- "Assumption Ranges" on page 21
- "Forecast Ranges" on page 21
- "Analyzing Certainty" on page 21

About Risk

Risk is usually associated with uncertainty, where risk includes the possibility of an undesirable event coupled with severity. For example, if sales for next month are above a certain amount (a desirable event), then orders will reduce the inventory. If the reduction in inventory is large enough, there will be a delay in shipping orders (an undesirable event). If a delay in shipping means losing orders (severity), then that possibility presents a risk. As uncertainty and risk increase, decision-making becomes more difficult.

There are two points to keep in mind when analyzing risk:

- Where is the risk?
- How significant is the risk?

Quantifying Risks with Spreadsheet Models

After you identify risks, a model can help you quantify them. Quantifying a risk means determining the chances that the risk will occur and the cost if it does, to help you decide whether a risk is worth taking. For example, if there is a 25% chance of running over schedule, costing you \$100 out of the own pocket, that might be a risk you are willing to take. But if you have a 5% chance of running over schedule, knowing that there is a \$10,000 penalty, you might be less willing to take that risk.

Finding the certainty of achieving a particular result is often the goal of model analysis. Risk analysis takes a model and sees what effect changing different values has on the bottom line. Risk analysis can:

- Help end "analysis paralysis" and contribute to better decision-making by quickly examining all possible scenarios
- Identify which variables most affect the bottom-line forecast
- Expose the uncertainty in a model, leading to a better communication of risk

Assumption Ranges

For each uncertain variable in a simulation, you can define the possible values with a probability distribution. A simulation calculates numerous scenarios of a model by repeatedly picking values from the probability distribution for the uncertain variables and using those values for the cell. In Crystal Ball, distributions and associated scenario input values are called assumptions. They are entered and stored in assumption cells. For more information on assumptions and probability distributions, see "About Assumptions and Probability Distributions" on page 40.

Forecast Ranges

Since all those scenarios produce associated results, Crystal Ball also keeps track of the forecasts for each scenario. These are important outputs of the model, such as totals, net profit, or gross expenses. For each forecast, Crystal Ball remembers the cell value for all the trials (scenarios). After hundreds or thousands of trials, you can view sets of values, the statistics of the results (such as the mean forecast value), and the certainty of any particular value. Chapter 6 provides more information about charts of forecast results and how to interpret them.

Analyzing Certainty

The forecast results show the different result values for each forecast and also the probability of obtaining any value. Crystal Ball normalizes these probabilities to calculate another important number: the certainty.

The chance of any forecast value falling between —Infinity and +Infinity is always 100%. However, the chance — or certainty — of that same forecast being at least zero (which you might want to calculate to make sure that you make a profit) might be only 45%. For any range you define, Crystal Ball calculates the resulting certainty. This way, not only do you know that the company has a chance to make a profit, but you can also quantify that chance by saying that the company has a 45% chance of making a profit on a venture (a venture you might, therefore, decide to skip).

How Crystal Ball Differs from Traditional Analysis Tools

When risk and uncertainty exist, traditional spreadsheet analysis tries to analyze this uncertainty using:

- "Point Estimates" on page 22
- "Range Estimates" on page 22
- "What-if Scenarios" on page 22

Point Estimates

Point estimates are when you use what you think are the most likely values (technically referred to as the mode) for the uncertain variables. These estimates are the easiest, but can return misleading results. For example, try crossing a river with an average depth of three feet. Or, if it takes you an average of 25 minutes to get to the airport, leave 25 minutes before the flight takes off. You will miss the plane 50% of the time.

Range Estimates

Range estimates typically calculate three scenarios: the best case, the worst case, and the most likely case. These types of estimates can show you the range of outcomes, but not the probability of any of these outcomes.

What-if Scenarios

What-if scenarios are usually based on range estimates, and are often constructed informally. What is the worst case for sales? What if sales are best case but expenses are the worst case? What if sales are average, but expenses are the best case? What if sales are average, expenses are average, but sales for the next month are flat?

These traditional methods have limitations:

- Changing only one spreadsheet cell at a time makes it virtually impossible to explore the entire range of possible outcomes.
- "What-if" analysis always results in single-point estimates that do not indicate the likelihood of achieving any particular outcome. While single-point estimates might tell you what is possible, they do not tell you what is probable.

This is where simulation with Crystal Ball comes in. Crystal Ball uses Monte Carlo simulation to generate a range of values for assumptions you define. These inputs feed into formulas defined in forecast cells.

You can use this process to explore ranges of outcomes, expressed as graphical forecasts. You can view and use forecast charts to estimate the probability, or certainty, of a particular outcome.

Monte Carlo Simulation and Crystal Ball

Spreadsheet risk analysis uses both a spreadsheet model and simulation to analyze the effect of varying inputs on outputs of the modeled system. One type of spreadsheet simulation is Monte Carlo simulation, which randomly generates values for uncertain variables over and over to simulate a model.

Monte Carlo simulation was named for Monte Carlo, Monaco, where the primary attractions are casinos containing games of chance. Games of chance such as roulette wheels, dice, and slot machines exhibit random behavior.

The random behavior in games of chance is similar to how Monte Carlo simulation selects variable values at random to simulate a model. When you roll a die, you know that either a 1, 2, 3, 4, 5, or 6 will come up, but you do not know which for any particular trial. It is the same with the variables that have a known range of values but an uncertain value for any particular time or event (for example, interest rates, staffing needs, stock prices, inventory, phone calls per minute).

The following sections describe the benefits of Monte Carlo simulation and how it works in Crystal Ball:

- "Benefits of Monte Carlo Simulation" on page 23
- "How Crystal Ball Uses Monte Carlo Simulation" on page 23

Benefits of Monte Carlo Simulation

Crystal Ball uses Monte Carlo simulation to overcome both of the spreadsheet limitations encountered with traditional spreadsheet analysis ("How Crystal Ball Differs from Traditional Analysis Tools" on page 21):

- You can describe a range of possible values for each uncertain cell in a spreadsheet. Everything you know about each assumption is expressed all at once. For example, you can define the business phone bill for future months as any value between \$2500 and \$3750, instead of using a single-point estimate of \$3000. Crystal Ball then uses the defined range in a simulation.
- With Monte Carlo simulation, Crystal Ball displays results in a forecast chart that shows the entire range of possible outcomes and the likelihood of achieving each of them. In addition, Crystal Ball keeps track of the results of each scenario for you.

How Crystal Ball Uses Monte Carlo Simulation

Crystal Ball implements Monte Carlo simulation in a repetitive three-step process. For each trial of a simulation, Crystal Ball repeats the following three steps:

- 1. For every assumption cell, Crystal Ball generates a random number according to the probability distribution you defined and places it into the spreadsheet.
- 2. Crystal Ball recalculates the spreadsheet.
- 3. Crystal Ball then retrieves a value from every forecast cell and adds it to the chart in the forecast windows.

This is an iterative process that continues until either:

- The simulation reaches a stopping criterion
- You stop the simulation manually

The final forecast chart reflects the combined uncertainty of the assumption cells on the model's output. Keep in mind that Monte Carlo simulation can only approximate a real-world situation. When you build and simulate the own spreadsheet models, you need to carefully examine the nature of the problem and continually refine the models until they approximate the situation as closely as possible.

Crystal Ball Feature Overview

The following sections introduce the main features of Crystal Ball:

- "Charts and Analysis Tools" on page 24
- "Other Crystal Ball Tools" on page 31
- "Process Capability Features" on page 33
- "Trend Analysis with Predictor" on page 33
- "Optimizing Decision Variable Values with OptQuest" on page 33

Charts and Analysis Tools

Crystal Ball offers several types of charts and reports, introduced in these sections:


- "Forecast Charts" on page 24
- "Overlay Charts" on page 25
- "Trend Charts" on page 26
- "Sensitivity Charts" on page 26
- "Scatter Charts" on page 27
- "Assumption Charts" on page 28
- "OptQuest Charts" on page 29
- "Reports" on page 29
- "Extracting and Pasting Data" on page 30

These graphical analysis tools are all accessed through the Analyze menu. They are discussed in Chapter 6, "Analyzing Forecast Charts," Chapter 7, "Analyzing Other Charts," and Chapter 8, "Creating Reports and Extracting Data."

Forecast Charts

Forecast charts are the basic tool for Crystal Ball results analysis. They use frequency distributions to show the number of values occurring in a given interval. The result is a range of values representing possible and probable values for a given forecast formula based on input assumption definitions. You can use forecast charts to evaluate the certainty of obtaining a particular value or range of forecast values. For more information on forecast charts, see Chapter 6, "Analyzing Forecast Charts."

Figure 1 Forecast Chart Showing Gross Profit



Overlay Charts

After completing a simulation with multiple related forecasts, you can create an overlay chart to view the frequency data from selected forecasts in one location. Then, you can compare differences or similarities that otherwise might not be apparent. You can customize the overlay chart to accentuate these characteristics. You can also use the overlay chart to fit standard distributions to forecasts.

The chart in Figure 2 shows reliabilities of three kinds of manufacturing material with lines of best fit for each type of distribution.

Overlay Chart 1 Edit View Overlay Preferences Help ✓ Enable Rotation 5,000 Trials Frequency View Overlay Chart 1 280 0.06 240 Fit: Beta 0.05 200 Material 1 Reliability Probability Fit: Gamma 160 Material 2 Reliability 0.03 120 Fit: Beta 80 0.02 Material 3 Reliability 40 0.01 0.00 1.00 1.10 1.30

Figure 2 Overlay Chart with Forecasts and Lines of Best Fit

Note: This overlay chart is shown in 3D, rotated view with legend to the right.

Trend Charts

After completing a simulation with multiple related forecasts, you can create a trend chart to view the certainty ranges of all the forecasts on a single chart. The ranges are displayed as a series of patterned bands. Each band represents the certainty range into which the actual values of the forecasts fall. For example, the band which represents the 90% certainty range shows the range of values into which a forecast has a 90% chance of falling.

The chart in Figure 3 shows increases in sales over time.




Figure 3 Upward trending sales figures, by quarter


If you have cyclical or seasonal time-series data, you can use Predictor for data trend analysis. For more information, see "Trend Analysis with Predictor" on page 33.

Sensitivity Charts

Sensitivity charts show the influence of each assumption cell on a particular forecast cell. During a simulation, Crystal Ball ranks the assumptions according to their correlation (or sensitivity) to each forecast cell. The sensitivity chart displays these rankings as a bar chart, indicating which assumptions are the most or least important in the model. You can print the sensitivity chart or copy it to the clipboard.

The chart in Figure 4, following, shows the effects of several assumptions on forecasted net profits for the pharmaceutical company discussed in "Tutorial 2 — Vision Research" on page 315. Market penetration accounts for about 84% of variation of net profits for a new product under consideration.

Figure 4 Effects of Assumptions on Net Profit




The Tornado Chart tool provides alternate ways to measure and chart sensitivity. For more information, see "Tornado Chart" on page 32.

Scatter Charts

Scatter charts show correlations, dependencies, and other relationships between pairs of forecasts and assumptions plotted against each other. You can plot scatter charts directly through the Analyze menu, or you can create a sensitivity chart and choose Sensitivity, then Open Scatter Chart to create a chart showing how the assumptions with the greatest impact relate to the target forecast.

In its basic form, a scatter chart (Figure 5) contains one or more plots of a target variable mapped against a set of secondary variables. Each plot is displayed as a cloud of points or symbols aligned in a grid within the scatter chart window. Optional correlation coefficients indicate the strength of the relationship.


Figure 5 Scatter Chart with Correlations Displayed



In another form of scatter chart, called the Matrix view, each selected variable is plotted against the other selected variables to show the relationships among them, as shown in Figure 6.

Figure 6 Scatter Chart in Matrix View


Scatter Chart 1



Assumption Charts

Assumption charts show trial values (that is, random numbers) for the current simulation layered over the assumption's ideal probability distribution. They are generated automatically each time a simulation runs.

Figure 7 Assumption chart for Body Weight



OptQuest Charts

You can use OptQuest, included in Crystal Ball Decision Optimizer, to optimize certain forecast statistics or process capability metrics using variable inputs provided by Crystal Ball decision variables. Optimization results are displayed in OptQuest charts. For more information, see "Optimizing Decision Variable Values with OptQuest" on page 33,

Predictor Charts

You can use Predictor to perform time series forecasts based on seasonal and non-seasonal historical data. Time series forecasts are displayed in Predictor charts. For more information, see "Trend Analysis with Predictor" on page 33.

Reports


Crystal Ball has powerful reporting capabilities. You can customize reports to include the following charts and data:

- Assumption, forecast, overlay, trend, sensitivity, scatter, and (optionally) OptQuest charts
- Forecast summaries, statistics, percentiles, and frequency counts
- Assumption parameters
- Decision variables

Reports are created as Microsoft Excel workbooks. You can modify, print, or save the report in the same way as any other workbook.

Figure 8 shows part of a forecast report for the Vision Research model.

Figure 8 Sample Forecast Report



Extracting and Pasting Data

You can select Analyze, then Extract Data to extract forecast information generated by a simulation and place it in a new Microsoft Excel workbook. You can extract these types of data: statistics, percentiles, chart bins, sensitivity data, and trial values.

Figure 9 shows statistics data extracted from a sales spreadsheet.

Figure 9 Extracted Statistics Data

	Α	В	С
1	Statistics	Ending Sales Year 3 - Q1	Ending Sales Year 3 - Q2
2	Trials	1000	1000
3	Mean	\$17,081,883	\$17,950,914
4	Median	\$17,042,665	\$17,930,927
5	Mode		
6	Standard Deviation	\$1,117,986	\$1,297,577
7	Variance	\$1,249,891,924,272	\$1,683,706,187,431
8	Skewness	0.19	0.15
9	Kurtosis	3.22	3.19
10	Coeff. of Variability	0.07	0.07
11	Minimum	\$13,695,983	\$14,060,365
12	Maximum	\$21,289,239	\$22,981,379
13	Range Width	\$7,593,255	\$8,921,013
14	Mean Std. Error	\$35,354	\$41,033

Other Crystal Ball Tools

The Run menu offers a variety of special tools for analyzing data and displaying results in more detail. Select Run, then More Tools to choose from tools discussed in the following sections (in Microsoft Excel 2007 or later, select More Tools in the Tools group of the Crystal Ball ribbon):

- "Batch Fit" on page 31
- "Bootstrap" on page 31
- "Correlation Matrix" on page 32
- "Data Analysis" on page 32
- "Decision Table" on page 32
- "Scenario Analysis" on page 32
- "Tornado Chart" on page 32
- "2D Simulation" on page 32
- "Integration Tools" on page 33
- "Compare Run Modes" on page 33

All of these except the integration tools are discussed in Chapter 9, "Crystal Ball Tools.". The integration tools are only available in Crystal Ball EPM and are described in the Oracle Crystal Ball Enterprise Performance Management Integration Guide.

Additional tools, Predictor and OptQuest, may also appear on the Run menu (or the Tools group in Microsoft Excel 2007 or later) in certain editions of Crystal Ball. For a description of these features, see "Trend Analysis with Predictor" on page 33 and "Optimizing Decision Variable Values with OptQuest" on page 33.

Batch Fit

The Batch Fit tool fits probability distributions to multiple data series. It helps create assumptions when you have historical data for several variables. Inputs are rows or columns of data. Outputs include fitted assumptions (probability distributions), tables of goodness-of-fit statistics and correlation coefficients calculated from the data series.

Bootstrap

The Bootstrap too estimates the reliability or accuracy of statistics or percentiles for forecasts or other sample data. This tool doesn't assume that the statistics or percentiles are normally distributed. The main input is the forecast to be analyzed. Outputs are a forecast chart of the distributions for each statistic or percentile.

Correlation Matrix

The Correlation Matrix tool defines a matrix of correlations between assumptions to more accurately model the interdependencies between variables. Inputs are the assumptions to correlate. The output is a correlation matrix, loaded into the model.

Data Analysis

The Data Analysis tool imports data directly into Crystal Ball forecasts for analysis in Crystal Ball.

Decision Table

Decision variables are values you can control, such as such as how much to charge for a product or how many wells to drill. The Decision Table tool runs multiple simulations to test different values for one or two decision variables. Inputs are the decision variables to test. The output is a table of results you can analyze further using forecast, trend, or overlay charts.

If you want to optimize decision variable values to reach a specific objective and you have OptQuest, you can use it to find solutions. For more information, see "Optimizing Decision Variable Values with OptQuest" on page 33.

Scenario Analysis

The Scenario Analysis tool runs a simulation and then matches all of the resulting values of a target forecast with the corresponding assumption values. Then, you can see which combination of assumption values gives a particular result. The input is the forecast to be analyzed. The output is a table of all the forecast values matched with the corresponding value of each assumption.

Tornado Chart

The Tornado Chart tool measures the input of each model variable one at a time, independently, on a target forecast. The inputs are the target forecast and the assumptions, decision variables, and precedent cells to test against. The output is a tornado chart, which shows the sensitivity of the variables using range bars, or a spider chart, which shows the sensitivity of the variables using sloping lines.

2D Simulation

The 2D Simulation tool helps determine how much of the variation within a model is caused by uncertainty and how much by true variability. The inputs are the target forecast and the assumptions to analyze. The output is a table which includes the forecast means, the uncertainty assumption values, and the statistics, including percentiles, of the forecast distribution for each simulation. Results are also graphed on an overlay chart and a trend chart.

Integration Tools

The optional integration tools link Crystal Ball EPM with Oracle Hyperion Strategic Finance, Fusion Edition and other Enterprise Performance Management applications using the Strategic Finance Setup wizard and Oracle Hyperion Smart View for Office, Fusion Edition. The Strategic Finance tool creates a workbook of Strategic Finance information that can then be analyzed with Crystal Ball. The Enterprise Performance Management tool connects to Oracle Enterprise Performance Management applications such as Oracle Essbase using Smart View so you can run Crystal Ball simulations on data in the underlying database. The integration tools are only available if you have Strategic Finance and Crystal Ball EPM. They are described in the *Oracle Crystal Ball Enterprise Performance Management Integration Guide* and a help file that is available after you start the Strategic Finance Setup wizard.

Compare Run Modes

If you have Crystal Ball Decision Optimizer including the Extreme Speed feature, you can use the Compare Run Modes tool to compare simulation run time in Normal and Extreme speed. For more information about this tool, see Appendix B, "Using the Extreme Speed Feature".

Process Capability Features

If you use Six Sigma or other quality methodologies, Crystal Ball's process capability features can help you improve quality in the organization. For a brief description of these features and how to use them, see Appendix D, "Using the Process Capability Features."

Trend Analysis with Predictor

You can use Predictor to project trends based on time-series data, such as seasonal trends.

For example, you can look at home heating fuel sales for previous years and estimate sales for the current year. You can also run regression analysis on related time-series data.

For more information about Predictor, see the Oracle Crystal Ball Predictor User's Guide.

Optimizing Decision Variable Values with OptQuest

Decision variables are variables you can control, such as product pricing or investment levels. If you have OptQuest, an optional feature available in Crystal Ball Decision Optimizer, you can use it to find the best values for the decision variables to obtain the desired outcomes.

For example, you can find the optimal investment mix that maximizes the probability of a portfolio's return being above a certain threshold.

For more information about OptQuest, see the Oracle Crystal Ball Decision Optimizer OptQuest User's Guide.

Steps for Using Crystal Ball

Follow these general steps to create and interpret simulations with Crystal Ball. The remaining chapters provide detailed instructions:

- 1. Create a spreadsheet model in Microsoft Excel format with data and formula cells that represent the situation to analyze.
 - "What is a Model?" on page 20 discusses spreadsheet models. Also see the references in the Bibliography.
- 2. Start Crystal Ball.
 - If you have not set up Crystal Ball to load automatically with Microsoft Excel, start Crystal Ball as described in "Starting and Closing Crystal Ball" on page 34.
- 3. Load a spreadsheet model.
- 4. Using Crystal Ball, define assumption cells and forecast cells. If appropriate for the situation, you can also define decision variable cells.
 - For more information, see "Entering Assumptions" on page 40 and continue on with Chapter 4.
- 5. Set run preferences for the simulation, as described in "Setting Run Preferences" on page 78.
- 6. Run the simulation, following the instructions in "Starting Simulations" on page 83.
- 7. Analyze the results. See "Understanding and Using Forecast Charts" on page 97 for suggestions.
- 8. Consider using Predictor or OptQuest, if available, for further analysis.
- 9. Take advantage of the many resources available to help you get the most out of Crystal Ball.

Resources for Learning Crystal Ball

If you are new to Crystal Ball, Appendix C offers tutorials to quickly introduce Crystal Ball's features and workflow. Consider completing the tutorials before you continue on with the more detailed instructions that follow.

For information on support, training, and referral services, see the Crystal BallWeb site at:

http://www.oracle.com/crystalball

Starting and Closing Crystal Ball

You can start Crystal Ball manually or you can set up Crystal Ball to start automatically whenever you start Microsoft Excel.

Starting Crystal Ball Manually

To start Crystal Ball manually, in Windows, choose **Start**, then **All Programs**, then **Oracle** Crystal Ball, then Crystal Ball.

Microsoft Excel opens with the Crystal Ball menus and toolbar. If Microsoft Excel is already running when you give this command, Crystal Ball opens a new instance of Microsoft Excel.

Starting Crystal Ball Automatically

- To set Crystal Ball to start automatically each time you start Microsoft Excel:
- 1 In Windows, choose Start, then All Programs, then Oracle Crystal Ball, then Application Manager.
- 2 Check Automatically Launch Crystal Ball When Microsoft Excel Starts.
- 3 Click OK.

Note: You can also use the Microsoft Excel Add-ins Manager to open Crystal Ball within Microsoft Excel when Microsoft Excel is already open and to close Crystal Ball without closing Microsoft Excel. For instructions, see the *Oracle Crystal Ball Installation and Licensing Guide*.

Crystal Ball Welcome Screen

The first time you start Crystal Ball, the Welcome screen opens, similar to Figure 10, following. Depending on the Crystal Ball version, the license features, and whether you are using a purchased or a trial version, the screen might differ somewhat from this illustration.

Figure 10 Crystal Ball Welcome Screen



You can use the Welcome screen to:

- Set certain preferences according to how you use Crystal Ball
- Enable accessibility preferences (see "Activating and Using Accessibility Mode" on page 388.
- View the Crystal Ball Web site
- Display the Oracle Technology Network, where you can download applications and documentation
- Display the Crystal Ball online documentation in HTML format
- Display instructions for licensing Crystal Ball
- Close the screen and start using Crystal Ball
- Display the dialog for opening workbook files
- Display the Oracle Crystal Ball Examples Guide and open example workbooks

For an explanation of the "primary application type" and "accessibility" settings, click the button.

Closing Crystal Ball

- To close Crystal Ball, use one of the following options:
- Right-click the Crystal Ball icon, Crystal Ball, in the Windows taskbar and choose Close, or
- Close Microsoft Excel.

If you want, you can choose Run, then Reset Simulation to reset the model and then choose File, then Save to save it before you close Crystal Ball.

Crystal Ball Menus and Toolbar

The following sections describe the main elements of the Crystal Ball user interface:

- "The Crystal Ball Menus" on page 36
- "The Crystal Ball Toolbar" on page 37

The Crystal Ball Menus

When you load Crystal Ball with Microsoft Excel, some new menus are displayed in the Microsoft Excel menubar:

- Define—contains commands that define and select assumption, decision variable, and forecast cells; perform Crystal Ball copy data, paste data, and clear data operations; and set cell preferences.
- Run—contains commands that start, stop, reset, and single-step through simulations; freeze
 variables; launch the Crystal Ball tools and Predictor as well as OptQuest, if available; and
 set run preferences.
- Analyze—contains commands that create a variety of charts and reports, extract data, and save or restore results.

For specific information about commands, see Appendix F, "Accessibility."

Note: If you are using Crystal Ball in Microsoft Excel 2007 or later, you will see a single Crystal Ball menu and then a ribbon instead of a toolbar. For more information about using Crystal Ball with Office 2007, see Appendix F.

The Crystal Ball Toolbar

To help you set up spreadsheet models and run simulations, Crystal Ball has a customized toolbar that provides instant access to the most commonly used menu commands.

Figure 11 The Crystal Ball Toolbar



The tools in the first three groups are from the Define menu. The tools from the next two groups are from the Run menu. The remaining tool groups are for creating charts, running the OptQuest and Predictor tools (if available to you), creating reports and extracting data from results, and displaying Crystal Ball online help. When you point to a tool icon, the tooltip is displayed and identifies it.

To hide or display the Crystal Ball toolbar for the current session, choose View, then Toolbars, then Crystal Ball.

Note: If you are using Crystal Ball in Microsoft Excel 2007 or later, you will see a single ribbon instead of a toolbar. For more information about using Crystal Ball with Office 2007, see "Crystal Ball Ribbon in Microsoft Excel 2007 or Later" on page 402.

Crystal Ball Tutorials

The easiest way to learn Crystal Ball is to work through the tutorials in Appendix C. Tutorial 1 is basic and will help you understand the material in the following chapters. Tutorial 2 shows more about how to create models and run Crystal Ball simulations as described in the following chapters.

3

Defining Model Assumptions

In This Chapter

Introduction	39
Types of Data Cells	39
About Assumptions and Probability Distributions	40
Defining Assumptions	40
Entering Assumptions	40
Additional Assumption Features	43
Using the Crystal Ball Distribution Gallery	

Introduction

This chapter provides step-by-step instructions for setting up assumption cells in Crystal Ball models so simulations can be run against them. This chapter also describes all the ways you can use the Distribution Gallery to organize the favorite distributions and define categories of distributions to share with others. The next chapter describes how to define decision variable and forecast cells and to cut, copy, and paste data.

If you are a new user, you should start by working through the tutorials in Appendix C, and then read this chapter. After you complete this chapter and Chapter 4, read Chapter 5 for information on setting preferences and running simulations.

Types of Data Cells

Crystal Ball uses three types of data cells as inputs and outputs:

- Assumption cells are input cells that contain values that you are unsure of: the uncertain
 independent variables in the problem you are trying to solve. The assumption cells must
 contain simple numeric values, not formulas or text.
- Decision variable cells are input cells that contain values that are within the control to change. The decision variable cells must contain simple numeric values, not formulas or text. These are used by some of the Crystal Ball tools and by OptQuest.
- Forecast cells (dependent variables) are output cells that contain formulas that refer to one or more assumption and decision variable cells. The forecast cells combine the values in the assumption, decision variable, and other cells to calculate a result. A forecast cell, for example, might contain the formula =C17*C20*C21.


About Assumptions and Probability Distributions


For each uncertain variable in a simulation, or assumption, you define the possible values with a probability distribution. The type of distribution you select depends on the conditions surrounding the variable. Common distribution types are normal, triangular, uniform, and lognormal, as shown in Figure 12.

Figure 12 Common Distribution Types









During a simulation, Crystal Ball calculates numerous scenarios of a model by repeatedly picking values from the probability distribution for the uncertain variables and using those values for each assumption cell. Commonly, a Crystal Ball simulation calculates hundreds or thousands of scenarios, or trials, in just a few seconds. The value to use for each assumption for each trial is selected randomly from the defined possibilities.

Because distributions for independent variables are so important to simulations, selecting and applying the appropriate distribution is the main part of defining an assumption cell. For more information on probability distributions, see "Understanding Probability Distributions" on page 229.

Defining Assumptions

- To define an assumption:
- 1 Determine the most appropriate probability distribution for each uncertain variable:
 - List everything you know about the conditions surrounding this variable.
 - Review the descriptions of the probability distributions in "Selecting Probability Distributions" on page 233.
 - Consider using Crystal Ball's distribution fitting feature, described in "Fitting Distributions to Data" on page 46.
 - Select the distribution that characterizes the variable.
- 2 Enter the assumption as described in the next section, "Entering Assumptions" on page 40.

Entering Assumptions


- To enter an assumption:
- Select a cell or a range of cells. The cells can be blank or have numeric values, but they cannot have formulas or text.

2 Choose **Define**, then Define Assumption,




(In Microsoft Excel 2007 or later, click the upper half of the Define Assumption icon.) For each selected cell or cells in the selected range, Crystal Ball displays the Distribution Gallery dialog (Figure 13).

Figure 13 The Distribution Gallery with Basic category Selected



- 3 In the Distribution Gallery, select the distribution you want from the gallery. The Basic category contains several frequently used distributions. Click All to view all distributions shipped with Crystal Ball. For details, see "Using the Crystal Ball Distribution Gallery" on page 53.
 - Alternately, click the Fit button to fit a distribution to historical data, as described in "Fitting Distributions to Data" on page 46.
 - For more information on the Distribution Gallery, see "Distribution Gallery Window" on page 53.
- 4 When the Define Assumption dialog opens (Figure 14), type in a title and the parameters for the distribution. The parameters can be either numeric values or cell references.


Figure 14 Normal Distribution



To change the distribution type, click Gallery to return to the Distribution Gallery and then select another distribution.

To see more information, click the More button, , near the Name field.More information is displayed in the Define Assumption dialog as shown in Figure 15.

Figure 15 The Define Assumption Dialog, Expanded



In the expanded Define Assumption dialog, you can:

- Enter truncation minimum and maximum values in the minimum and maximum truncation fields (just below the distribution, accessible by tabbing).
- Use the truncation grabbers to truncate the value range.
- Use numeric spinners (arrows beside the field) to adjust parameter settings.

• Click the Less button, (a), to hide the minimum and maximum value fields and truncation grabbers. (For more information about truncating distributions, see "Truncating Distributions" on page 285.)

You can perform the following activities in both the standard and expanded Define Assumption dialog:

- Click the Gallery button to display the Distribution Gallery window and choose another distribution.
- Click the Correlate button to define correlations as described in "Specifying Correlations Between Assumptions" on page 49.
- Choose **Edit**, then **Add** in the menubar to add the currently defined assumption distribution to the Favorites category or a user-defined category in the Distribution Gallery.
- Use other menu commands to copy the chart, paste it into Microsoft Excel or another application, print data, change the view, use alternate parameters, set assumption and chart preferences, and display help as described in "Additional Assumption Features" on page 43.
- 6 When you have finished entering parameters to define the assumption, click **Enter**.

The distribution changes to reflect the values you entered. If you click OK instead of Enter, Crystal Ball accepts the parameters and closes the dialog.

7 Click OK.

If you selected a range of cells, repeat these steps to define the assumption for each cell.

Additional Assumption Features

As you enter assumption parameters, you can use cell references and alternate parameters. If you have historical data available, you can use Crystal Ball's distribution fitting feature to help simplify the process of selecting a probability distribution. You can also specify correlations between assumptions or freeze assumptions to exclude them from a simulation.

The following sections discuss advanced features that help you refine assumption definitions and use assumptions more effectively:

- "Entering Cell References and Formulas" on page 44
- "Alternate Parameter Sets" on page 45
- "Fitting Distributions to Data" on page 46
- "Specifying Correlations Between Assumptions" on page 49
- "Setting Assumption Preferences" on page 52
- "Freezing Crystal Ball Data Cells" on page 82

Entering Cell References and Formulas

In addition to numeric values, you can enter a reference to a specific cell in a parameter field. Cell references must be preceded by an equal sign (=). Cell references can be either absolute or relative. You can also enter formulas and range names.

If necessary, you can press F4 to change references from relative to absolute or back to relative. This also applies to cell references in fields other than assumption parameters.

Note: All cell references in parameters are treated like absolute references when cutting and pasting Crystal Ball data. Crystal Ball always stores the cell reference in A1 format even if the Microsoft Excel preference is set to R1C1 format. The global R1C1 format preference is not affected by running Crystal Ball, but the name ranges are, in fact, changed to A1 format since that is the way Crystal Ball stores them.

To show cell references instead of current values when you enter them in parameter fields, choose **Parameters**, then **Show Cell References** in the Define Assumption dialog.

The following topics contain details about cell references, range names, and formulas:

- "Dynamic vs. Static Cell References" on page 44
- "Relative References" on page 44
- "Absolute References" on page 45
- "Range Names" on page 45
- "Formulas" on page 45

Dynamic vs. Static Cell References

Cell references in assumption parameters are dynamic and are updated each time the workbook is recalculated. Dynamic cell referencing gives you more flexibility in setting up models because you can change an assumption's distribution during a simulation.

Other types of cell references are static, such as the assumption name field and correlation coefficients. These cell references are calculated once at the beginning of a simulation.

Relative References

Relative references describe the position of a cell relative to the cell containing the assumption. For example, suppose an assumption in cell C6 refers to cell C5. If the assumption in C6 is copied to cell C9, the relative reference to C5 will then refer to the value in cell C8. With relative references, you can easily set up a whole row or column of assumptions, each having similar distributions but slightly different parameters, by performing just a few steps. An absolute reference, on the other hand, always refers back to the originally referenced cell, in this case C5.

Absolute References

To indicate an absolute reference, use a dollar sign (\$) before the row and the column. For example, to copy the exact contents of cell C5 into an assumption parameter field, enter the cell reference =\$C\$5. This causes the value in cell C5 to be used in the assumption cell parameter field. Later, if you decide to copy and paste this assumption in the worksheet, the cell references in the parameter field will refer to the contents of cell C5.

Range Names

You can also enter cell references in the form of range names, such as =*cellname*. Then, the referenced cell can be located anywhere within a worksheet as long as its name does not change.

Formulas

You can enter Microsoft Excel formulas to calculate parameter values as long as the formula resolves to the type of data acceptable for that parameter. For example, if a formula returns a string, it would not be acceptable in a parameter that requires a numeric value, such as Minimum or Maximum.

Alternate Parameter Sets

For all the continuous probability distributions except uniform, you can use percentiles for parameters when defining a distribution. This option gives you added flexibility to set up assumptions when only percentile information is available or when specific attributes (such as the mean and standard deviation) of the variables in a model are unknown.

For example, if you are defining a triangular distribution, but are unsure of the absolute minimum and maximum values of the variable, you can instead define the distribution using the 10th and 90th percentiles along with the likeliest value. This gives you a distribution that has 80%, or four-fifths of the values, occurring between the two specified percentiles.

To change the parameter sets for the continuous distributions, use the Parameters menu in the menubar of the Define Assumption dialog. The currently selected parameter set has a check mark next to it. If you choose Custom in the Parameters menu, you can replace any or all of the standard parameters with any percentile.

To select a parameter set to use as the default when defining new assumptions of this type, choose Set Default from the Parameters menu.

Several special parameter sets are available with the lognormal distribution, including geometric and logarithmic sets. For more information, see the "Equations and Methods" chapter in the online *Oracle Crystal Ball Statistical Guide*.

Note: Alternate parameters cannot always be used with highly skewed distributions and extremely large or small parameter values.

Fitting Distributions to Data

If you have historical data available, Crystal Ball's distribution fitting feature can substantially simplify the process of selecting a probability distribution. Not only is the process simplified, but the resulting distribution more accurately reflects the nature of the data than if the shape and parameters of the distribution were estimated.

Distribution fitting automatically matches historical data against probability distributions. A mathematical fit determines the set of parameters for each distribution that best describe the characteristics of the data. Then, the closeness of each fit is judged using one of several standard goodness-of-fit tests. The highest ranking fit is chosen to represent the data. You can choose from among all distributions supported by Crystal Ballexcept the yes-no distribution.

Using Distribution Fitting

- To use distribution fitting:
- Select the cell where you want to create an assumption.

It can be blank or contain a simple value, not a formula.

Choose **Define**, then **Define Assumption**.



(In Microsoft Excel 2007 or later, click the upper half of the Define Assumption icon.)

In the Distribution Gallery, click Fit to select the source of the fitted data.

The Fit Distribution dialog opens.

- Choose a data location.
 - If the historical data is in a worksheet in the active workbook, choose Range, and then enter the data's cell range. If the range has a name, you can enter the name, preceded by an = sign.
 - If the historical data is in a separate text file, click Text File, and then either enter the path and name of the file or click Browse to search for the file. If you want, you can check Column and enter the number of columns in the text file.

When you use a file as the source of data, each data value in the file must be separated by either a comma, a tab character, a space character, or a list separator defined in Windows' Regional and Language Options panel. If actual values in the file contain commas or the designated list separator, those values must be enclosed in quotes. Allowable formats for values are identical to those allowed within the assumption parameter dialog, including date, time, currency, and numbers.

- Specify which distributions are to be fitted: 5
 - AutoSelect performs a basic analysis of the data to choose a distribution fitting option and ranking method. If the data includes only integers, fitting to all discrete distributions (with the exception of Yes-No) is completed using the Chi-square ranking statistic choice.

- All Continuous fits the data to all of the built-in continuous distributions (these distributions are displayed as solid shapes on the Distribution Gallery).
- All Discrete fits to all discrete distributions except yes-no and uses the Chi-square ranking statistic.
- Choose displays another dialog where you can select a subset of the distributions to include in the fitting.
- The final setting selects the distribution that was highlighted on the Distribution Gallery when you clicked the Fit button.
 - If you try to fit negative data to a distribution that can only accept positive data, that distribution will not be fitted to the data.
- 6 Specify how the distributions should be ranked.

In ranking the distributions, you can use any one of three standard goodness-of-fit tests:

- Anderson-Darling. This method closely resembles the Kolmogorov-Smirnov method, except that it weights the differences between the two distributions at their tails greater than at their mid-ranges. This weighting of the tails helps to correct the Kolmogorov-Smirnov method's tendency to over-emphasize discrepancies in the central region.
- Chi-Square. This test is the oldest and most common of the goodness-of-fit tests. It gauges the general accuracy of the fit. The test breaks down the distribution into areas of equal probability and compares the data points within each area to the number of expected data points. The chi-square test in Crystal Ball does not use the associated p-value the way other statistical tests (e.g., *t* or *F*) do.
- Kolmogorov-Smirnov. The result of this test is essentially the largest vertical distance between the two cumulative distributions.
 - The first setting, **AutoSelect**,, selects the ranking statistic automatically based on several factors. If all data values are integers, Chi-Square is selected.
- 7 If you know the data corresponds to certain shape, location, or other special parameter values for some distributions, you can choose to lock parameters to those values.
 - Check Lock Parameters and enter appropriate values in the Lock Parameters dialog. For details, see "Locking Parameters When Fitting Distributions" on page 49.
- 8 By default, only values for the selected ranking statistic are displayed in the Comparison Chart dialog. To show values for all three statistics, check Show All Goodness-of-fit Statistics at the bottom of the Fit Distribution dialog.
- 9 Click OK.

The Comparison Chart opens.

Confirming the Fitted Distribution

When the Comparison Chart opens (Figure 16), the fitted distributions are displayed in the Comparison Chart dialog, starting with the highest-ranked distribution (best fit) down through to the lowest (worst fit).




Figure 16 Comparison Chart with Goodness of Fit View, Chi-Square Ranking Statistic

- To confirm which of the selected distributions to use for an assumption:
- 1 Use the Comparison Chart dialog to visually compare the quality of the fits or to view the goodness-offit statistics. You can:
 - Use the Next and Previous buttons to scroll through the fitted probability distributions. Each probability distribution is shown superimposed over the data.
 - Choose Preferences, then Chart to change chart features so that similarities or differences are more clearly accentuated.
 - Click Cancel to return to the Fit Distribution dialog.
- 2 To use the currently displayed distribution, either the best fit or another of the choice, click Accept.

The Assumption dialog opens with the parameter entries taken from the chosen distribution. You can change the distribution parameters before you click OK.

Distribution Fitting Notes

p-Values

When goodness-of-fit values are displayed, as in the distribution fitting comparison chart, *p*-values are displayed for some combinations of ranking methods and fitted distributions. These express the degree to which the actual fit conforms to a theoretical fit for that fitting test and distribution (see "Goodness Of Fit" on page 108 for more information). When the Chi-square method is used, *p*-values are displayed for all continuous and discrete distributions. *P*-values are also displayed for the following continuous distributions when the Anderson-Darling or Kolmogorov-Smirnov methods are used: normal, exponential, logistic, maximum extreme, minimum extreme, uniform, gamma, Weibull, and lognormal. *P*-values for the other distributions are under development.

Since p-values for Anderson-Darling and Kolmogorov-Smirnov statistics are influenced by the number of data points being fitted, an adjustment formula is used to arrive at the asymptotic Anderson-Darling and Kolmogorov-Smirnov statistic for a given sample size. The quality of fitted parameters and the calculated p-value deteriorates as the sample size decreases. Currently, Crystal Ball needs at least 15 data points for fitting all the distributions.

Multiple Fittings

To run fittings on multiple data sets, use the Batch Fit tool.

Locking Parameters When Fitting Distributions

Some distributions can fit the data more accurately if you are able to enter and lock parameter values so the shape, location, or certain other parameters of a distribution better match the data. In most cases where you can fit data to distributions within Crystal Ball, you can also choose to lock parameters.

- To lock parameters:
- 1 Check Lock Parameters in a distribution fitting dialog box. For example, you can make that setting in the Fit Distribution dialog for assumptions.
 - The Lock Parameters dialog opens.
- 2 Select one of the available distributions and enter a value for one or more of its parameters.

Specifying Correlations Between Assumptions

In Crystal Ball, assumption values are usually calculated independently of each other. Crystal Ball generates random numbers for each assumption without regard to how random numbers are generated for other assumptions. However, dependencies often do exist between variables in a system being modeled.

You can define correlations between pairs of assumptions. These relationships are described in mathematical terms using a correlation coefficient, a number between -1.0 and +1.0 that measures the strength of the relationship. A positive value means that when one assumption is high, the other is likely to be high. A negative value means that the assumptions are inversely related; when one is high, the other is likely to be low.

Crystal Ball uses rank correlation (Spearman) for all correlation computations to relate assumptions with different distribution types. For more information, see the "Statistical Definitions" chapter of the *Oracle Crystal Ball Statistical Guide*.

- To correlate an assumption to one or more assumptions:
- 1 Select an assumption cell.
- 2 Choose Define, then Define Assumption, Land Choose Define Assumption dialog opens.

- 3 Click the More button beside the Assumption name in the Define Assumption dialog.
- 4 Click Correlate.

The Define Correlation dialog opens.


In the Define Correlation dialog, click Choose to select the assumption to correlate from the Choose Assumptions dialog. (You can only correlate assumptions in the same workbook.)

The Choose Assumptions dialog provides a list of the names of all the assumptions defined in the workbook.

By default, the dialog opens in a hierarchical Tree view. If you prefer, you can click the List icon, , to change it to List view.

6 Choose one or more of the assumption names on the list in the Define Correlation dialog and click OK (Figure 17).

Figure 17 Define Correlation Dialog with Second Assumption



After you select the second assumption, the cursor moves to the Coefficient Entry field (below the Choose button).

The chosen assumptions are displayed in the list of correlations. The currently selected assumption is also displayed immediately next to the Choose button.

- 7 Enter a correlation coefficient using one of the following methods and then click OK to accept the correlation and close the dialog box:
 - Enter a value between -1 and 1 (inclusive) in the Coefficient Entry field.
 After you type the number, the slider control on the correlation coefficient scale moves to the selected value.
 - Choose a cell that contains the correlation coefficient.
 If you choose a cell with values that change during the simulation, it is the initial value of the cell that is used for the coefficient.
 - Drag the slider control along the correlation coefficient scale.

The value you select is displayed in the Coefficient Entry field.

- Type the desired correlation coefficient in the Coefficient field in the correlation list.
- Click Calc.

A small dialog opens at the bottom of the first dialog. Enter the range of cells on the spreadsheet that contains the empirical values that Crystal Ball should use to calculate a correlation coefficient.

Enter the range of cells in the standard A1:A2 format, where A designates the column and 1 and 2 designate the first and last cell rows, respectively. For example, if one set of values is in column Q, rows 10 through 15 and the second set of values is in column R, rows 10 through 15, enter the range in the first field as Q10:Q15 and the range in the second field as R10:R15.

Crystal Ball calculates the correlation coefficient, enters it in Coefficient Entry field, and moves the slider control to the correct position.

The two cell ranges do not necessarily have to have the same dimensions, but they must contain the same number of value cells and must be in the same workbook. The cell ranges are read in a row-by-row fashion.

Each time you select a new assumption or correlation coefficient, Crystal Ball displays a sample correlation of the correlated assumption values in the chart.




Figure 18 Correlation Chart

The points on the chart represent the pairing of assumption values as they would actually occur when running a simulation. The solid line running through the middle of the chart indicates the location where values of a perfect correlation (+1.0 or -1.0) would fall. The closer the points are to the solid line, the stronger the correlation.

In the example, an Inflation Rate assumption and an Oil Price/Barrel assumption have been correlated using a coefficient of 0.8, a strong positive correlation. As the points on the chart show, higher inflation values tend to be associated with higher oil prices and vice versa. This chart can help you begin to understand how the two assumptions are related.

You can specify as many of these paired correlations as you want for each assumption, up to the total number of assumptions defined in a workbook.

You can generally ignore correlations between variables if one or both variables do not impact the output or are not highly correlated.

Setting Assumption Preferences

The Define Assumption dialog has a Preferences menu in the menubar. This menu has the following main options:

Table 1 Preferences Menu, Define Assumptions Dialog Box

Setting	Effect
Assumption Preferences	Manage window display during simulations
Chart Preferences	Determine the appearance of the assumption chart

The Chart Preferences settings are discussed in "Setting Chart Preferences" on page 115.

If you choose Assumption Preferences, the Assumption Preferences dialog opens.

You can use this dialog to:

- Choose a view for the assumption chart:
 - o Probability shows a graph of all possible values for the assumption variable and the probability of their occurrence.
 - O Cumulative Probability shows a graph of the probability that the assumption variable will fall at or below a given value.
 - Reverse Cumulative Probability shows a graph of the probability that the assumption variable will fall at or above a given value.
 - Statistics shows a table of measures of central tendency, variability, minimum and maximum values, and other statistics for the assumption variable.
 - Percentiles shows a table of percentiles and their associated values for the assumption variable.

Note: For examples of each view, see "Changing the Distribution View and Interpreting Statistics" on page 104.

Determine if and when the assumption chart window opens when a simulation runs.

To show the generated values in the window, turn on the **Store Assumption Values For Sensitivity Analysis** run preference. To do so, click the Run Preferences button, then click the Options tab.

You can click Apply To to copy these settings to other assumptions. If necessary, you can click Defaults to restore original default settings. When the settings are complete, click OK.

Using the Crystal Ball Distribution Gallery

You can use the Distribution Gallery to add, manage, and share libraries of distributions. Workgroups can use this powerful feature to modify and share customized distributions over their local networks when collaborating on custom models. They can also email these to other Crystal Ball users for use with their models.

Displaying the Distribution Gallery

- To display the Distribution Gallery:
- 1 With Crystal Ball open within Microsoft Excel, click in a cell.
- 2 Click the Define Assumption tool, or choose Define, then Define Assumption.

(In Microsoft Excel 2007 or later, click the upper half of the Define Assumption icon. Or, click the bottom half of the Define Assumption icon and choose Distribution Gallery at the bottom of the distribution list.)

The Distribution Gallery opens, as shown in Figure 19.

Distribution Gallery Window

As shown in Figure 19, the Distribution Gallery has a menubar, a category pane with folders containing distributions, a distribution pane that displays all distributions in the selected category, and a description pane that describes the selected distribution.




Figure 19 The Distribution Gallery Window

The following sections explain each part of the Distribution Gallery:

- "The Distribution Gallery Menubar and Buttons" on page 54
- "The Category Pane" on page 54

- "The Distribution Pane" on page 54
- "Description Pane" on page 55

The Distribution Gallery Menubar and Buttons

The Distribution Gallery menubar has the menus summarized in Table 2.

Table 2 Distribution Gallery Menus

Menu	Command Summary
Edit	Commands that copy, paste, modify, and delete distributions. You can copy from any category you are allowed to use but you can only paste, modify, and delete within the Favorites category or new categories you or others have created. You cannot modify or delete distributions in the Basic and All categories; these categories are reserved only for unmodified distributions shipped with Crystal Ball.
Categories	Commands that create, delete, view and modify properties of, and rearrange category folders in the Category pane. You can use two additional commands to share categories with others (Publish) and use categories others have shared (Subscribe).
View	Commands that change how distributions are displayed in the Distribution pane (as Thumbnails, Large Icons, or Small Icons) and hide or show distribution details and descriptions.
Help	Commands that display online help for the Distribution Gallery and the selected distribution.

The Fit button at the bottom of the Distribution Gallery opens the Crystal Ball distribution fitting feature. This feature can help you select an appropriate distribution for the assumption you are defining. For more information, see "Fitting Distributions to Data" on page 46.

The Help button displays online help for the currently selected distribution.

The Category Pane

Categories are groups of distributions contained in folders.

- Basic is the default category. It contains several of the most common distributions: normal, triangular, uniform, lognormal, yes-no, and discrete uniform.
- All contains all distributions shipped with Crystal Ball, in unmodified form, including those
 also supplied in the Basic category.
- Favorites is the default category for distributions that are copied or modified by users. For
 example, if you want to copy a triangular distribution from Basic and modify it, you could
 paste it to Favorites and change it there.

You can use the Categories menu to create new category folders for holding distributions. Then, you can use the Edit commands to add distributions to the new categories and modify them.

The Distribution Pane

The Distribution pane shows all distributions in the selected category. You can use the View menu to change how they are displayed, as shown in Table 3.

Table 3 Distribution View Examples

View Command	Example
Thumbnails	Normal Triangular
Large Icons	Normal Triangular Uniform
Small Icons	Normal Triangular Uniform Lognormal Yes-No Discrete Uniform
Details	Name Parameters Summary Normal Mean, Standard Deviation Familiar bell curve used to describe natural phenom Triangular Minimum, Likeliest, Maximum Used for rough estimation when data is limited

Description Pane

The Description pane is displayed at the bottom of the Distribution Gallery and provides a detailed description of the selected distribution.

You can turn off Show Description in the View menu to show more distributions in the Distribution pane.

Managing Distributions

Working with distributions is a key part of defining a model in Crystal Ball. Assumptions are the main source of simulation input and distributions are the core component of assumption definitions.

Several Crystal Ball features help you modify and save distributions for future use and share them with other Crystal Ball users:

- To create new distributions, see "Creating New Distributions" on page 56.
- To copy distributions, see "Copying and Pasting Distributions" on page 56.
- To paste distributions, see "Copying and Pasting Distributions" on page 56.
- To modify distributions, see "Modifying Distributions" on page 57.
- To modify distribution summaries and descriptions, see "Modifying Distribution Summaries and Descriptions" on page 57.
- To delete distributions, see "Deleting Distributions" on page 58.
- To set up pages for printing, see "Setting Up Distributions for Printing" on page 58.
- To print distribution information, see "Printing Distribution Information" on page 58.

Creating New Distributions

You can create new distributions by:

- Copying and renaming a distribution within the Distribution Gallery ("Creating Distributions by Copying" on page 56)
- Defining a distribution in the Define Assumption dialog and adding it to the Distribution Gallery ("Creating Distributions by Saving Assumptions" on page 56)

Creating Distributions by Copying

To create a distribution by copying another distribution, follow the steps in the next section, "Copying and Pasting Distributions" on page 56

Then, you can rename and modify the new distribution as needed.

Creating Distributions by Saving Assumptions

- To create a distribution by saving an assumption:
- With Crystal Ball running, click a cell in the spreadsheet.
 It can be a value cell where you want to define an assumption or a blank cell.
- 2 Click the Define Assumption tool or choose Define, then Define Assumption.
- 3 Follow the steps in "Defining Assumptions" on page 40 to select a distribution and enter parameters.
- 4 Choose Edit, then Add To Gallery in the Define Assumption dialog.
- 5 The Add To Gallery dialog opens, where you can name the new distribution and select a category for it.
- 6 Click OK to save the distribution in the selected category.
 - If you created correlation data, it is not saved although the distribution type and any parameter settings are saved.

Now, the new distribution is available for use just as another other distribution, either supplied with Crystal Ball or user-defined.

Copying and Pasting Distributions

- To copy and paste a distribution within Favorites or another user-defined category or from one category to another:
- 1 In the Distribution Gallery, click a category in the Category pane.
- 2 Click a distribution in the Distribution pane.
- 3 Choose Edit, then Copy or right-click and choose Copy from the menu.
- 4 Click Favorites or another user-defined category in the Category pane.
- 5 Choose Edit, then Paste or right-click and choose Paste from the menu.

The distribution is pasted into the selected category. If another distribution in that category had the same name, the new distribution is displayed with the next larger number at the end. For example, if the category already had Normal1, a new normal distribution would be named Normal2.

- 6 To rename the copied distribution, click its name and wait a moment. The name is highlighted for editing.
- 7 Type the new name and press Enter.
- 8 Modify the new distribution as needed.

Note: To easily copy a distribution from one category to another, select the category with the distribution to copy. Then, drag the distribution into the closed category folder where you want to copy it. When you select that category, the copied distribution is displayed and you can rename it or modify it as described in the next section. However, the distribution you copied remains in its original folder, unchanged.

Modifying Distributions

Distributions in the Basic or All categories are read-only and cannot be changed. However, after you have copied a distribution or used it to define an assumption, you can modify it in a number of ways.

- To modify a distribution:
- 1 Create a distribution as described in "Creating New Distributions" on page 56 and "Copying and Pasting Distributions" on page 56.
- 2 Right-click the user-defined distribution and choose Modify (or click the distribution and choose Edit, then Modify in the Distribution Gallery menubar).
 - A prompt opens as a reminder that you are modifying a distribution and not defining an assumption.
- 3 Click OK to display the distribution in a window that looks like the Define Assumption dialog.
- 4 Modify the distribution as if it were an assumption, following the instructions in "Defining Assumptions" on page 40.
- 5 When you are finished, click **OK**.

The modifications are saved with the distribution in the category and no assumption is created in the spreadsheet.

Modifying Distribution Summaries and Descriptions

You can edit the distribution Summary and Description text for any distribution in Favorites or another user-defined category.

- To modify Summary text for a distribution:
- 1 In the Distribution Gallery, set the View menu to **Details**.

- 2 In the Distribution pane, right-click the Summary text to edit.
- 3 In the menu, choose **Modify Summary For** distribution name.
- 4 When the Modify Summary dialog opens, enter the new text in the edit field and click **OK**.

The new text immediately replaces the original Summary text.

- To modify Description text for a distribution:
- 1 Select the distribution in the Distribution Gallery.
- 2 Be sure View, then Show Description is checked.
- 3 Right-click the text in the Description pane.
- 4 In the menu, choose **Modify Description For** distribution name.
- 5 When the Modify Description dialog opens, edit the description in the edit field and click OK.

(Alternately, you can just select the description text to modify and edit it. Be careful to avoid modifying descriptions accidentally.)

The new text immediately replaces the original description.

Deleting Distributions

You can delete a distribution in the Favorites category or other user-defined category. However, the original distribution of that type remains in Basic or All.

To delete a distribution, choose Edit, then Delete in the Distribution Gallery menubar, or rightclick the distribution and choose Delete.

Setting Up Distributions for Printing

- To select paper, print orientation, and margins before printing:
- 1 Double-click a distribution in the Distribution Gallery to display it in the Define Assumption dialog.
- 2 Choose Edit, then Page Setup in the Define Assumption menubar.
- 3 Set paper size and source, orientation, and margins.
- 4 Click OK to accept the settings.

Printing Distribution Information

- To print a distribution:
- 1 Double-click a distribution in the Distribution Gallery to display it in the Define Assumption dialog.
- 2 Choose **Edit**, then **Print** in the Define Assumption menubar.
- 3 Confirm the printer and other settings.
- 4 Click **OK** to print the distribution.

Printed information includes the chart with current parameter settings. The distribution name and parameters are displayed in a box beneath the graphic.

You can preview before printing. In step 2, choose Edit, then Print Preview before you choose Edit, then Print.

Managing Categories

You can use distribution categories to organize distributions and share them with other users in the organization:

- To create new categories, see "Creating New Categories" on page 59.
- To view and edit category properties, see "Viewing and Editing Category Properties" on page 60.
- To delete categories, see "Deleting Categories" on page 60.
- To rearrange category order, see "Rearranging Category Order" on page 60.
- To share categories over networks, see "Sharing Categories over Networks" on page 61.
- To use categories shared by others, see "Using Shared Categories" on page 61.
- To send and receive categories, see "Sharing Categories Through Email" on page 62.
- To edit shared categories, see "Editing Shared Categories" on page 62.
- To unpublish shared categories, see "Unpublishing Shared Categories" on page 63.
- To change distribution display within categories, see "The Distribution Pane" on page 54.

Creating New Categories

You can create a category of distributions to help with model development or to share with other users.

- To create a category:
- 1 Choose Categories, then New.

The New Category dialog box opens.

Enter the name of the category. Optionally, you can enter a description, the name, and a version number (helpful for shared categories).

Note: If you enter an ampersand (&) before a letter in the category name, that letter becomes a shortcut key. You can then use it to select the category from the keyboard by holding down the Alt key and pressing the shortcut key. Shortcut keys are displayed underlined when you press the Alt key. The new category should use a different shortcut key from other categories.

3 Click OK.

The new folder is displayed in the Category pane and the category can be selected and used like Favorites or any other user-defined category.

Viewing and Editing Category Properties

Category properties include the information entered in the New Category dialog when a category is created. You can view properties for any category, but you can only edit properties for user-defined categories.

- To view or edit properties:
- 1 Click a category in the Category pane.
- 2 Choose Categories, then Properties.

The Category Properties dialog opens.

- 3 Optional: Edit the information in any field that appears white. Read-only fields are gray.
- 4 When you have finished viewing or editing category properties, click OK.

The Category Properties dialog closes and any property changes are saved.

Deleting Categories

You can delete user-defined categories and all the distributions within them, but be careful not to delete a distribution you will need to use later.

- To delete a category:
- Select the category in the Category pane.
- 2 Choose Categories, then Delete.

A confirmation message is displayed.

3 When you click OK, the category and all its distributions are immediately deleted.

Original built-in distributions shipped with Crystal Ball are never deleted; only copies can be deleted.

Rearranging Category Order

Changing the order of folders in the Category pane can be helpful if you want to arrange categories in alphabetical order or put frequently used categories near the top.

- To move a category folder:
- 1 Select the folder to move in the Category pane.
- 2 Choose Categories, then Move Up or Categories, then Move Down.

The folder moves up or down one position each time you choose the menu command.

3 Continue moving folders up and down until they are arranged as preferred.

Sharing Categories over Networks

The ability to publish and subscribe to categories is powerful. You can create large numbers of specific distributions, add them to a category, and then publish it on the network for other users in the organization. Likewise, you can easily subscribe to categories published by others. You can also email categories to others and use categories sent to you.

- To publish a category to a shared folder on the computer or to a location on a network:
- 1 Select a category in the Distribution Gallery.
- 2 Choose Categories, then Publish.

The Browse For Folder dialog opens.

3 Browse to locate the folder where you want to place the category. You can click + to expand a folder and view its subfolders. If you need to create a new folder, select the folder that should include the new folder.

The new folder is displayed and you can rename it. For example, if you want to create a folder under My Documents, select My Documents. Then, when you click New Folder, the new folder is displayed under My Documents at the same level as My Pictures.

4 Click OK.

The dialog closes and a copy of the category is saved there.

Using Shared Categories

If other users have published categories in a shared folder on their computers or the network, you can access them for use with the Distribution Gallery. This is called subscribing to categories.

- To subscribe to a category, learn its name and location, then:
- Open the Distribution Gallery and choose Categories, then Subscribe.

The Subscribe To A Category dialog opens.

2 Click Add.

The Browse For Folder dialog opens.

- 3 Locate the target folder and click **OK** to add the new path to the Subscribe To A Category dialog box.
- 4 Click OK to load all categories in the listed paths and close the Subscribe To A Category dialog box.

All loaded categories are available for use as if they were on the local computer.

Note: Shared categories can be used like local categories in most ways. However, they cannot be modified unless they also exist in a folder on the local computer. If several users copy a published category locally and then modify it, they can publish their own versions and overwrite each others' changes. If you are publishing a category, you might want to make the shared folder read-only to avoid this problem.

- To edit or delete a path, or rearrange the path order:
- 1 Open the Subscribe To A Category dialog as described in step 2 in the previous steps.
- 2 Select the path to a target category.
- 3 Click an action button: Edit, Delete, Move Up, or Move Down.
- 4 When you are finished, click OK.

If you delete a path to a subscribed category, that category disappears from the Category pane of the Distribution Gallery. You can resubscribe to it at any time to use it again.

Sharing Categories Through Email

You can use email to share a user-defined category with a Crystal Ball user who is not on the network.

- To share categories using email:
- 1 Create and address an email message.
- 2 Use the file attachment command and browse to the category folder on the machine.
 - By default, this is C:\Documents and Settings\\the username\Application Data\Oracle\Crystal Ball\11.0\Store.
- 3 Select the user-defined category to send.
 - Crystal Ball category files have this format: CBCategory_name.cbc. For example, the Favorites category The Browse For Folder dialog opens. as CBCategory_Favorites.cbc.
- 4 Send the email.
- To use a category received in an email:
- 1 Open the email.
- 2 Save the attachment in the Crystal Balll category folder, by default C:\Documents and Settings\the username\Application Data\Oracle\Crystal Ball\11.0\Store.

The next time you start Crystal Ball, the new category is displayed in the Category pane of the Distribution Gallery.

Alternately, you can save the attachment in another folder, then follow the steps in "Using Shared Categories" on page 61 to subscribe to it.

Editing Shared Categories

- To edit a published user-defined category:
- 1 Make the changes in the Distribution Gallery.
- 2 Follow the steps in "Sharing Categories over Networks" on page 61 to publish the category to the same place.

3 When you are prompted to overwrite the existing file, click **OK**.

This publishes the modified category. Now, the next time a subscriber loads Crystal Ball, the modified category is displayed in the Distribution Gallery.

Unpublishing Shared Categories

- ➤ If you have published a user-defined category and want to remove it from publication, or "unpublish" it:
- 1 Use My Computer or Windows Explorer to browse to the folder where the category was published.
 - The category file is displayed as CBCategory_name.cbc. For example, the Favorites category is displayed as CBCategory_Favorites.cbc.
- 2 Select the category file and delete it.

The next time a subscriber loads Crystal Ball, the deleted category will not be displayed in the Distribution Gallery. Existing assumptions based on distributions from that category will still work but those distributions won't be available for creating new assumptions.

A Note About Modifying Shared Categories

If a shared category is moved, deleted, or otherwise modified, subscribers to that category will not notice the changes until they close Crystal Ball and open it again. It is wise to notify subscribers when a change of that type occurs so they can start a new Crystal Ball session and refresh their view of the Distribution Gallery.



Defining Other Model Elements

In This Chapter

Introduction	65
Defining Decision Variable Cells	65
Defining Forecasts	66
Working with Crystal Ball Data	70
Setting Cell Preferences	74
Saving and Restoring Models	75

Introduction

Chapter 3 describes how to start building a spreadsheet model by defining assumption cells. This chapter provides step-by-step instructions for completing models in Crystal Ball so you can run simulations against them. As you work through these instructions, you will learn to define decision variable cells and forecasts and how to cut, copy, and paste Crystal Ball data.

After you complete this chapter, read Chapter 5, "Running Simulations" for information on setting run preferences and running simulations.

Defining Decision Variable Cells

Decision variables are the variables that you can control, such as rent to charge or the amount of money to invest. Decision variables are not required for simulation models, but they can be useful when comparing and optimizing alternate scenarios. Several of the Crystal Ball tools discussed in Chapter 9, "Crystal Ball Tools", use and benefit from decision variables.

You also use decision variables with OptQuest, if available.

- To define one or more decision variable cells:
- Select a cell or range of cells.

Select value cells or blank cells only. You cannot define a decision on a formula or non-numeric cell.

2 Select Define , then Define Decision.



The Define Decision Variable dialog opens.

- 3 Click the More button, \(\bigsim \), to display all settings.
- 4 Complete the Define Decision Variable dialog:
 - Name is the name of the decision variable.
 - Bounds are the upper and lower limits for the decision variable range.
 - Type defines whether the variable type is:
 - o Continuous can assume any value between the lower and upper bounds
 - Discrete can assume values at specific intervals between the lower and upper bounds

With Discrete selected, Step specifies the interval between values for discrete variables. For example Step = 1 could be used to specify whole dollars, while Step = .5 could specify 50-cent increments.

- O Binary is 0 or 1 to represent a yes-no decision, where 0 = no and 1 = yes
- O Category can assume any discrete integer between the lower and upper bounds (inclusive), where the order of the values does not matter. This type, used for attributes or indexes, is mostly used when numeric values represent conditions or groups instead of numeric values. Example model Groundwater Cleanup.xls contains a decision variable named Remediation Method, expressed as integers 1, 2, and 3. These are not numeric values, but instead they represent three different remediation methods for groundwater cleanup and can be defined using the Category type.
- Custom can assume any value from a list of two or more specific values or a cell range reference. When values are entered directly, separate them by a valid list separator a comma, semicolon, or other value specified in the Windows regional and language settings. If a cell range reference is used, it must include more than one cell so there will be two or more values. Blanks and non-numerical values in the range are ignored.

You can use cell referencing to name a decision variable, define the lower and upper limits, set the step size, and define custom values ("Entering Cell References and Formulas" on page 44).

- 5 Click OK.
- 6 Repeat these steps for each decision variable in the model.

Note: There is no absolute limit to the number of data cells you can define per worksheet. In general, you should define less than 1000 assumptions, decision variables, and forecasts per worksheet.

Defining Forecasts

After you define the assumption cells and decision variables, you are ready to select and define forecast cells. Forecast cells usually contain formulas that refer to one or more assumption and

decision variable cells. The forecast cells combine cells in the model to produce the results you need.

- ➤ To define forecast cells:
- Select a formula cell or a range of formula cells.
- 2 Select Define, then Define Forecast,

The Define Forecast dialog opens.

If you have activated Crystal Ball's process capability features, additional fields are displayed. These are described in "Setting Specification Limits and Targets" on page 372.

- 3 Complete the Define Forecast dialog:
 - Name is the name of the forecast.
 - Units is the name of the units that are displayed at the bottom of the forecast chart, such as Millions.
- 4 To set additional forecast preferences, click the More button, , to expand the Define Forecast dialog.

 The expanded Define Forecast dialog has four tabs of additional options and fields, listed in "Setting Forecast Preferences" on page 67.
- 5 Click OK.
- 6 Repeat steps 1-4 for each forecast in the model.

Clicking Defaults in the expanded Define Forecast dialog restores the original default settings in place of any new settings you have made. You can also click Apply To to use the settings in other charts and worksheets.

Setting Forecast Preferences

The forecast preference settings are displayed when you click the More button, , in the Define Forecast dialog box or select Preferences, then Forecast in a forecast chart menu bar. The tabs in this dialog control several important aspects of Crystal Ball:

- "Forecast Window Tab" on page 68 window display and distribution fitting for the forecast
- "Precision Tab" on page 68 precision control settings
- "Filter Tab" on page 69 value filtering, which discards values inside or outside a range for the current forecast
- "Auto Extract Tab" on page 70 automatic data extraction to Microsoft Excel when a simulation stops.

Note: See "Confidence intervals" in the "Statistical Definitions" chapter of the online *Oracle Crystal Ball Statistical Guide* for more information about how absolute and relative precision relate to the confidence interval.

When you expand the Define Forecast dialog or open the Forecast Preferences dialog, the Forecast Window tab is displayed by default.

Forecast Window Tab

The Forecast Window tab of the Forecast Preferences dialog offers the following settings to manage window display and distribution fitting for the forecast:

- View settings Set the forecast window's display type ("Changing the Distribution View and Interpreting Statistics" on page 104).
- Window settings Set whether to display the forecast window automatically while the simulation runs or when the simulation stops. You can display one or more forecasts while the simulation is running. If you decide not to display the forecast, the simulation continues to run. You can override this option and close all the forecast windows using the Suppress Forecast Windows option in the Run Preferences dialog ("Setting Run Preferences" on page 78).
 - Window, then While Running Simulation Displays the forecast window automatically during simulations. This slows down the simulation.
 - Window, then When Simulation Stops Displays the forecast window automatically after the simulation stops. This is faster than displaying the window during simulations.
- Fit Distribution Fits a probability distribution to the forecast. After selecting this checkbox in this group, you can click Fit Options to select the distribution and goodness-of-fit tests you want.

Click OK to apply settings on the current tab to the active forecast. Or, you can click Apply To to apply settings on the active tab to the active worksheet, the active workbook, or all workbooks. At any time, you can click Defaults to restore the original default settings on the active tab in the dialog.

Precision Tab

The Precision tab of the Forecast Preferences dialog manages the precision control settings that determine when to stop a simulation based on confidence intervals for selected statistics.

The current simulation must be reset before precision control settings will take effect.

Select from the following settings:

Specify The Desired Precision For Forecast Statistics — Activates the precision control settings for the forecast. Crystal Ball uses these settings only if the simulation is set to stop when it reaches the specified precision from the Run Preferences dialog ("Setting Run Preferences" on page 78). The statistics available for precision control are the mean, the standard deviation, and an indicated percentile. Check any or all. If you select Percentile,

you can enter any percentile value greater than 0 and less than 100 to use as a precision control statistic.

- Must Be Within Plus Or Minus Selects which range to use for precision control, absolute units or relative percentage.
 - Units Determines the size of the confidence interval, in actual forecast units, used to test the precision of the forecast statistics.
 - Percent Determines the size of the confidence interval, in percentage terms, used to test the precision of the forecast statistics.

Click OK to apply settings on the current tab to the active forecast. Or, you can click Apply To to apply settings on the active tab to the active worksheet, the active workbook, or all workbooks. At any time, you can click Defaults to restore the original default settings on the active tab in the dialog.

Filter Tab

You can use the Filter tab of the Forecast Preferences dialog to discard values inside or outside a range for the current forecast or globally for all forecasts in a model. The values are not permanently deleted, only discarded for the purposes of the current analysis.

Select from the following settings:

- Set A Filter On The Forecast Values Activates the filter settings for the forecast.
- Include Values In The Range Discards values from the forecast if they fall above or below the two values in the range fields. Endpoints are included, not excluded.
- Exclude Values In The Range Discards values from the forecast if they fall between the two values in the range fields. The range is inclusive; Crystal Ball discards values inside the range as well as values equal to the range endpoints.
- When filtering values for this forecast, remove values for the same trial from other forecasts too For each trial that a value is not included or is excluded, removes the value for that trial from all other forecasts in the model, and assumptions too. For example, if the filter for the current forecast is set to include values from 4 through 10 and the value for the third trial is 12, the value for the third trial will be filtered from the current forecast and all other forecasts and assumptions in the model, regardless of the values in the other forecasts. If this setting is checked and you run a Forecasts report, "globally filtered" is displayed for the forecast in the Summary data following the filter description.

Note: This setting can be checked in the Define Forecast dialog for several forecasts with different filter settings. In that case, filtering for each selected forecast will be applied across all forecasts.

Click OK to apply settings on the current tab to the active forecast. Or, you can click Apply To to apply settings on the active tab to the active worksheet, the active workbook, or all workbooks. At any time, you can click Defaults to restore the original default settings on the active tab in the dialog.

Auto Extract Tab

You can use the Auto Extract tab of the Forecast Preferences dialog to specify which statistics to extract automatically to Microsoft Excel after the simulation stops.

The Auto Extract settings create tables of unformatted statistics primarily for use in other analyses. To extract formatted data, see "Extracting Data" on page 165.

You can choose from the following settings:

- Extract forecast statistics automatically to the spreadsheet when the simulation stops —
 Activates the auto-extract feature.
- [listbox] The list of statistics you can extract. Check the statistics you want, then use the up and down arrows to rearrange their order, if you want.
- Starting Cell The first cell on the forecast's worksheet where the statistics will be copied. Be sure no data entries are displayed to the right of this cell and below it because data could be overwritten without warning.
- Formatting Whether to include labels on the extracted statistics and use Microsoft Excel's AutoFormat for the cells.
- **Direction** Whether the extraction proceeds vertically (downward) or horizontally (to the right) on the worksheet.

Click OK to apply settings on the current tab to the active forecast. Or, you can click Apply To to apply settings on the active tab to the active worksheet, the active workbook, or all workbooks. At any time, you can click Defaults to restore the original default settings on the active tab in the dialog.

Working with Crystal Ball Data

You can use special commands to copy, paste, and clear Crystal Ball cell definitions from cells. These are different from similar Microsoft Excel commands and must be used to copy Crystal Ball cell definitions (data). Other Crystal Ball commands select data and highlight it for review.

The following sections describe these special commands:

- "Editing Crystal Ball Data" on page 70
- "Selecting and Reviewing Data" on page 72

Editing Crystal Ball Data

You can use the edit commands to copy, paste, or clear assumptions, decision variables, or forecasts from cells. You can set up an entire row or column of assumptions, decision variables, or forecasts with just a few steps.

If you select a range of cells with more than one kind of Crystal Ball data — for example, an assumption and a forecast — Crystal Ball prompts for which data type to copy or clear.

Important! To copy Crystal Ball cell definitions, use only the Crystal Ball Copy Data, Paste Data, or Clear Data commands. Using the Microsoft Excel copy and paste commands only copies the cell value and attributes, including cell color or pattern.

The following sections describe how to edit Crystal Ball data:

- "Copying Crystal Ball Data" on page 71
- "Pasting Crystal Ball Data" on page 71
- "Clearing Crystal Ball Data" on page 72
- "Clearing All Crystal Ball Data of a Single Type" on page 72

Copying Crystal Ball Data

- To copy Crystal Ball assumptions, decision variables, or forecasts from one area of the spreadsheet to another area in the same workbook or a different workbook:
- Select a cell or a range of cells that contains the Crystal Ball data to copy.
- 2 Select **Define**, then **Copy** [**Data**].

If you select a range of cells with more than one kind of Crystal Ball data — for example, an assumption and a forecast — Crystal Ball prompts for which data type to copy.

Check the type or types to copy and click **OK**.

Pasting Crystal Ball Data

- To paste Crystal Ball data:
- Select a cell or a range of cells that to paste into.

It should contain cells with values if you are pasting assumptions or decision variables (unless the range is all blank) and cells with formulas if you are pasting forecasts.

If you are pasting assumptions or decision variables into a completely blank range of cells, Crystal Ball pastes them along with the underlying cell value from each copied cell. Forecasts must be pasted into a cell with a formula.

Choose **Define**, then **Paste** [**Data**],



Crystal Ball pastes all selected data types (assumptions, decision variables, and forecasts) from the copied range into the range selected in step 1. Any existing Crystal Ball data in the range selected for pasting will be overwritten.

For best results, use the Paste Data command *immediately* after the Copy Data command. You are probably already familiar with this from using Microsoft Excel; if you copy in Microsoft Excel and then do anything else but select a range for pasting, the copied area is unselected and you must select it again before continuing.

Clearing Crystal Ball Data

- To clear Crystal Ball data:
- 1 Select a cell or a range of cells that contains the Crystal Ball data to clear.
- 2 Select **Define**, then **Clear [Data]**,



If you select a range of cells with more than one kind of Crystal Ball data — for example, an assumption and a forecast — Crystal Ball prompts for which data type to clear.

3 Check the type or types to clear and click **OK**.

Crystal Ball clears the selected Crystal Ball data type (assumptions, decision variables, and forecast definitions) from each cell.

Clearing All Crystal Ball Data of a Single Type

- To clear all of one type of Crystal Ball data from all cells in the active worksheet:
- 1 Select one of the following options:
 - Define, then Select All Assumptions
 - Define, then Select All Decisions
 - Define, then Select All Forecasts

(In Microsoft Excel 2007 or later, choose Define , then Select and then choose one of the Select commands.)

2 Select **Define**, then **Clear [Data]**,



Crystal Ball clears the Crystal Ball data from all the selected cells in the active worksheet.

Edit, then Clear, then All also clears all Crystal Balldata from selected cells.

Selecting and Reviewing Data

After you define assumptions, decision variables, and forecast cells and return to the spreadsheet, you might want to check the cell definitions to make sure you have defined them as you intended. You can select all the assumption, forecast, and decision variable cells and check their corresponding dialogs in sequence.

The following sections describe how to review and select Crystal Ball data:

- "Reviewing Assumption Cells" on page 73
- "Reviewing Decision Variable Cells" on page 73
- "Reviewing Forecast Cells" on page 73
- "Reviewing Selected Cells" on page 74

Note: Regarding the following sections concerning selection, in Microsoft Excel 2007 or later, choose Define, then Select and then choose one of the Select commands.

Reviewing Assumption Cells

To review all assumption cells:



1 Select **Define**, then **Select All Assumptions**,



2 Select **Define**, then **Define Assumption**,



The distribution dialog for the first assumption cell opens.

- 3 Optional: Change the assumption parameters.
- 4 Click OK.

If you have more than one assumption cell, each assumption is displayed in turn. Repeat steps 3 and 4 to check the assumptions for each cell.

Reviewing Decision Variable Cells

- To review all decision variable cells:
- 1 Select **Define**, then **Select All Decisions**,



Crystal Ball selects each cell defined as a decision variable.

2 Select **Define**, then **Define Decision**,



The Define Decision Variable dialog for the first decision variable cell opens.

- 3 Optional: Change the decision variable parameters.
- 4 Click OK.

If you have more than one decision variable cell, each decision variable opens in turn. Repeat steps 3 and 4 to check the decision variables for each cell.

Reviewing Forecast Cells

To review all forecast cells:



1 Select **Define**, then **Select All Forecasts**,

Crystal Ball selects each cell defined as a forecast.

2 Select Define, then Define Forecast.



The forecast dialog for the first forecast cell opens.

- 3 Optional: Change the forecast definition.
- 4 Click OK.

If you have selected more than one forecast cell, each forecast is displayed in turn. Repeat steps 3 and 4 to check the forecast for each cell.

Reviewing Selected Cells

- To review Crystal Ball data cells on any open workbook:
- 1 Select **Define**, then **Select**.

The Select dialog opens.

By default, the dialog opens in hierarchical Tree view. All the assumptions are listed first, then all the decision variables, and finally all the forecasts.

You can click the Assumption, Decision Variable, and Forecast buttons to show and hide assumptions, decision variables, and forecasts, respectively. If you prefer to view available cells in list format, click the List button,

- 2 Check the cells to select.
- 3 Click **OK** to highlight all the selected cells so you can change their preferences or perform other actions on them.

You can use the Select dialog to select cells on more than one worksheet but you need to activate each worksheet in turn to review and, if desired, apply a command to all the selected cells.

Setting Cell Preferences

You can change the appearance of Crystal Ball assumption, forecast, and decision variable cells so you can quickly identify them in spreadsheets. You can set Crystal Ball to change the appearance of these cells as you define them, or you can change the appearance of predefined cells.

- ➤ To set cell preferences:
- 1 Choose Define, then Cell Preferences.

(In Microsoft Excel 2007 or later, choose Define, then Cell Prefs.)

The Cell Preferences dialog opens.

- 2 Click the tab for the kind of cell to format: Assumptions, Decision Variables, or Forecasts.
- 3 Make appropriate settings for the chosen cell type:

- Color Changes the color of each Crystal Ball data cell of the type modified by the selected tab.
- Pattern Changes the pattern of each Crystal Ball data cell of the type modified by the selected tab.
- Add Comment To Cell Adds a Microsoft Excel comment that provides more information about the Crystal Ball data within each cell. (Crystal Ball only updates cell comments when you define or redefine an assumption, decision variable, or forecast.)
- Set Cell Value To... Changes assumption cell values to the selected value (Mean or Median) when no simulation is running.
- **Set To Range Midpoint** Changes decision variable cell values to the selected value (range Midpoint, Minimum, or Maximum) when no simulation is running.
- 4 Click Apply To and choose whether to apply settings from only the current tab or all tabs in the Cell Preferences dialog.
- 5 Choose whether to apply the settings to all cell preferences of the chosen type(s) on the current Microsoft Excel worksheet, all worksheets in the current workbook, or all open workbooks and any new workbooks to be created later.
 - (The default is All Open And New Workbooks.)
- 6 Click OK to close the Apply To dialog and apply the settings to the chosen cell types and worksheets.

Note: Like some other preferences, cell preferences must be applied to all cells of the selected type(s) on the chosen worksheet(s) or workbook(s). If necessary, you can click the Defaults button before choosing Apply To to clear current cell preference settings and restore the original defaults.

Saving and Restoring Models

The distributions specified for each assumption cell, the settings specified for each forecast cell, and the range information for each decision variable cell are saved with their spreadsheet through the Microsoft Excel save process. When you open the spreadsheet again, Crystal Ball retains the assumption, forecast, and decision variable cells.

Note: When you run simulations, as described in the next chapter, you can also save and restore simulation results in a separate file for future display and analysis ("Saving and Restoring Simulation Results" on page 88).

5

Running Simulations

In This Chapter

Introduction	77
About Crystal Ball Simulations	77
Setting Run Preferences	78
Freezing Crystal Ball Data Cells	82
Running Simulations	83
Managing Chart Windows	86
Saving and Restoring Simulation Results	88
Functions for Use in Microsoft Excel Models	90
Running User-Defined Macros	91

Introduction

Chapter 3 and Chapter 4 describe how to define a spreadsheet model in Crystal Ball. This chapter provides step-by-step instructions for setting run preferences and running a simulation in Crystal Ball.

After you complete this chapter, read Chapter 6 for information about analyzing simulation results.

About Crystal Ball Simulations

After you define assumption, forecast, and decision variable cells in the spreadsheet model, you are ready to run the simulation. Crystal Ball uses a technique called Monte Carlo simulation to forecast the entire range of results most likely to occur in the model scenario. For more information on Monte Carlo simulation, see "Monte Carlo Simulation and Crystal Ball" on page 22.

Crystal Ball provides statistics that describe the forecast results. For more information on forecast results and statistics, see Chapter 6, "Analyzing Forecast Charts," and the "Statistical Definitions" chapter of the Oracle Crystal Ball Statistical Guide.

To run simulations in Crystal Ball, follow these basic steps:

1. Define assumptions ("Entering Assumptions" on page 40), forecasts ("Defining Forecasts" on page 66), and decision variable cells if appropriate ("Defining Decision Variable Cells" on page 65).

- 2. If you want, customize the appearance of each cell as described in "Setting Cell Preferences" on page 74.
- 3. Set run preferences as described in the next section, "Setting Run Preferences" on page 78.
- 4. Optionally, exclude certain data cells from the simulation by "freezing" them, described in "Freezing Crystal Ball Data Cells" on page 82.
- 5. Run the simulation as described in "Running Simulations" on page 83.

The following sections explain each step.

Setting Run Preferences

Run preferences determine how Crystal Ball runs a simulation.

- To change run preferences:
- 1 Reset the previous simulation if necessary, 44.



- 2 Choose Run, and then Run Preferences to display the Run Preferences dialog
- 3 Click the tab with the preferences to change, described in the following sections:
 - "Setting Trials Preferences" on page 79 Specify when to stop a simulation, namely number of trials, calculation errors, and precision control.
 - "Setting Sampling Preferences" on page 79 Set the sampling seed value, method, and sample size.
 - "Setting Speed Preferences" on page 80 Determine whether a simulation runs in Normal, Demo, or Extreme (if available) and set additional speed control options.
 - "Setting Options Preferences" on page 81 Set a number of run preferences, including whether sensitivity data and assumption values are stored, whether assumption correlations are activated, whether user macros are run, whether the Crystal Ball Control Panel opens, and whether advanced accessibility settings should be switched on.
 - "Setting Statistics Preferences" on page 82 Determine how Crystal Ball displays percentiles and activates process capability features.
- 4 Change any preferences on any tab.
- 5 Click OK.
- 6 To reset settings on the active tab to the original defaults, click **Defaults**.

Setting Trials Preferences

The Trials tab of the Run Preferences dialog sets preferences that stop a simulation: number of trials, calculation errors, and precision control. For general instructions, see "Setting Run Preferences" on page 78.

The current simulation must be reset before precision control settings will take effect.

The Trials tab of the Run Preferences dialog has these settings:

- Number Of Trials To Run Defines the maximum number of trials that Crystal Ball runs before it stops the simulation. If you check either of the checkboxes on this dialog, Crystal Ball only uses the maximum number of trials if forecast results do not meet the other stop criteria first.
- Stop On Calculation Errors When checked, stops the simulation when a mathematical error (such as division by zero) occurs in any forecast cell. If a calculation error occurs, to help you find the error, Crystal Ball does not restore the cell values. If no calculation errors occur, the simulation continues until it reaches the Number Of Trials To Run or (if set) when the specified precision is reached.

Note: In Extreme Speed, the simulation stops at the end of a burst of trials when an error occurs, and not immediately upon detecting the error.

- Stop When Precision Control Limits Are Reached When checked stops the simulation when certain statistics reach a specified level of precision. You choose the statistics and define the precision that triggers this option in each Define Forecast dialog. For instructions, see "Precision Tab" on page 68. Any forecasts set to use precision control must all reach their specified precision within the confidence level to stop the simulation. If all the forecasts set to use precision control don't meet the specified precision, the simulation stops when it reaches the Number Of Trials To Run. By default, precision control is on.
- Confidence Level Sets the precision level (confidence level) that indicates when to stop a simulation.

Setting Sampling Preferences

The Sampling tab of the Run Preferences dialog sets the sampling seed value, sampling method, and sample size. For general instructions, see "Setting Run Preferences" on page 78.

The Sampling tab of the Run Preferences dialog has these settings:

- Use Same Sequence Of Random Numbers Sets the random number generator to generate the same set of random numbers for assumptions, so you can repeat simulation results. When you select this option, enter an integer seed value in the Initial Seed Value field.
- Initial Seed Value Determines the first number in the sequence of random numbers generated for the assumption cells (integer).

Note: To reproduce the sample results shown in this manual, check Use Same Sequence Of Random Numbers and use a seed value of 999.

- Sampling Method Indicates whether to use Monte Carlo or Latin Hypercube simulation sampling. Latin Hypercube sampling generates values more evenly and consistently across the distribution, but requires more memory.
- Sample Size For Latin Hypercube sampling, divides each distribution into the specified number of intervals (bins). A higher number increases the evenness of the sampling method, while reducing the randomness.

Note: If you are using Microsoft Excel 2007 or later with multi-threading, there is no guaranteed order of execution for user-defined functions included in Crystal Ball models. For this reason, they will not always return consistent results, even if a seed is set.

Setting Speed Preferences

The Speed tab of the Run Preferences dialog adjusts how fast a simulation runs. Extreme speed is only available in Crystal Ball Decision Optimizer. Extreme speed is the default simulation speed, if it is available. Otherwise, Crystal Ball only runs in Normal or Demo speed. If you choose Normal or Demo speed, the Options button is active and you can make additional settings.

Note: If the license includes Extreme speed, be sure to read Appendix C, "Using the Extreme Speed Feature" on page 293, for important information about model compatibility.

For general instructions, see "Setting Run Preferences" on page 78.

The Speed tab of the Run Preferences dialog has these settings:

- Run Mode settings Determine overall simulation speed.
 - Extreme Speed Available as a Crystal Ball option. This setting runs simulations up to 100 times faster than Normal mode but is not suitable for some models. For details, see Appendix C, "Using the Extreme Speed Feature" on page 293.
 - o Normal Speed The standard simulation option for general model processing.
 - Demo Speed Runs simulations in "slow-motion" to make it easier to watch values change in spreadsheet cells and charts.
- Options settings Set update rules for the active worksheet (Normal and Demo speeds only). See "Speed Tab Options Settings" on page 81 for details.
- Chart Windows settings Set the redraw rate for any charts open during a simulation.
 - Redraw Every _ Seconds Defines the redraw rate in terms of time. The default value is 0.5.
 - O Suppress Chart Windows (Fastest) Closes all charts during simulation. Selecting this option overrides the Show Window preferences set for any charts. This option produces the fastest simulations.

Speed Tab Options Settings

Speed options are provided for Normal and Demo speed.

The following options are provided for Normal speed:

- Update Every Trial Updates Crystal Ball data in Microsoft Excel after each simulation trial. Dynamic references are still updated internally if another setting is chosen.
- Update Every _ Seconds Defines the update rate in terms of time. The default value is 0.5
- Minimize Workbooks (Fastest) Minimizes the Microsoft Excel window. This option produces the fastest simulations.
- Bring Microsoft Excel 2007 To The Foreground While Running Normal Speed Simulations
 (Faster) This option is available when Crystal Ball is running on Microsoft Excel 2007 or
 later. When checked, Microsoft Excel runs in the foreground in Normal speed and
 performance improves.

The following option is provided for Demo speed:

• Maximum Number of Trials/Second — With optimal processing, what is the greatest number of trials to run each second. The default value is 10.

Setting Options Preferences

The Options tab of the Run Preferences dialog sets a variety of run preferences. For general instructions, see "Setting Run Preferences" on page 78.

The Options tab of the Run Preferences dialog has these settings:

- Store assumption values for sensitivity analysis Stores the randomly generated values used during the simulation for display while running. To see those values during a simulation, select this option and the appropriate Window preference for each assumption to display. Values can also be exported to a spreadsheet after the simulation using the Extract Data command. This setting also generates sensitivity data during simulations. The sensitivity chart can display this information to show how much each assumption influences a particular forecast. For more information on the sensitivity chart, see "Understanding and Using Sensitivity Charts" on page 142. The sensitivity chart is not available unless you select this option before you run a simulation.
- Enable correlations Activates any defined correlations between assumptions.
- Run user-defined macros Runs any user-defined macros as part of the simulation process. For details, see "Running User-Defined Macros" on page 91.
- Show control panel When checked, activates the Crystal Ball Control Panel. For more information, see "The Crystal Ball Control Panel" on page 85.
- Leave control panel open on reset When checked, continues to display the Control Panel after a simulation is reset.

• Enable accessibility options — When checked, activates a number of user interface enhancements of special interest to individuals with visual impairments and other disabilities. See "Activating and Using Accessibility Mode" on page 388.

Setting Statistics Preferences

The Statistics tab of the Run Preferences dialog determines how Crystal Ball displays percentiles. Settings on this tab also activate capability metrics to support Six Sigma and other quality programs. For general instructions, see "Setting Run Preferences" on page 78.

The Statistics tab of the Run Preferences dialog has these settings:

- Calculate Percentiles As settings Determine how Crystal Ball defines percentiles. Selecting either of these options also affects the percentiles used for the assumption alternate parameters.
 - Probability Below A Value Defines percentiles as the percent chance (probability)
 that the associated variable value is at or below a particular value. This setting is the
 default.
 - O Probability Above A Value Defines percentiles as the percent chance (probability) that the associated variable value is at or above a particular value.
- Format Percentiles As settings Determine how Crystal Ball displays percentiles in charts and reports, with a percent sign or the percentile preceded by "P."
- Calculate Capability Metrics Activates the process capability features in Crystal Ball; when checked Crystal Ball displays capability metrics indicating process quality, provided at least an upper or lower specification limit is entered in the Define Forecast dialog.
- Options button When Calculate Capability Metrics is checked, displays the Capability
 Options panel for specifying short-term or long-term capability metrics formulas, a Z-score
 shift value, and other capability metrics calculation settings. For more information, see
 "Setting Capability Calculation Options" on page 370.

Freezing Crystal Ball Data Cells

You can use the Freeze command to "freeze" or exclude certain Crystal Ball assumption, decision variable, and forecast cells from a simulation. Then, you can investigate the effect certain cells have on the model while holding others to their worksheet values.

The Freeze command is useful when you have multiple workbooks open and do not want to include all of their data cells in a simulation. You can freeze any unwanted cells instead of closing the workbooks that contain them.

- To freeze Crystal Ball data cells:
- 1 Choose **Define**, then **Freeze**.
- 2 Select one or more assumptions, decision variables, or forecasts listed in the Freeze dialog.

Use the Show buttons to hide and show cells of a particular type. Use the Select buttons to select all or none of the listed cells.

You can also click View List to change from Tree view to List view,



3 Click OK.

Running Simulations

After you have defined assumption and forecast cells — and, optionally, decision variable cells — in the worksheet model, you are ready to run a simulation. During the simulation, Crystal Ball creates a forecast chart for each forecast cell using frequency distributions to show the range of possible results. By default, these charts are displayed as the simulation runs. As you run the simulation, observe the shape of the distribution to get a feel for the problem you are trying to solve. If you are new to Crystal Ball, try running simulations in Demo speed to watch the forecast values change from trial to trial.

During a Crystal Ball simulation, you can stop, reset, and continue the simulation at any time and manipulate the forecast charts independently, displaying or hiding them as needed. You can use the Crystal Ball Control Panel to perform many of the procedures described in this section. For details, see "The Crystal Ball Control Panel" on page 85.

During the simulation, Crystal Ball saves the forecast values in a list that grows as the simulation progresses. In "Extracting Data" on page 165, you will learn how to export these forecast values to other programs, such as a statistical analysis program.

The following sections describe how to run and control simulations:

- "Starting Simulations" on page 83
- "Stopping Simulations" on page 84
- "Continuing Simulations" on page 84
- "Resetting and Rerunning Simulations" on page 84
- "Single-Stepping" on page 84
- "The Crystal Ball Control Panel" on page 85

Starting Simulations

To start a simulation, choose Run, then Start [Simulation],



You can then stop, continue, single step, or reset the simulation.

After you start the simulation, the Run, then Start [Simulation] command changes to Stop [Simulation]. If you then choose Stop [Simulation], the simulation halts and the command changes to Continue Simulation (in Microsoft Excel 2003 or earlier). Choosing Continue Simulation restarts the simulation where it left off, and changes the menu command back to Stop [Simulation].

Stopping Simulations

To stop a simulation, select **Run**, then **Stop** [**Simulation**] or click the **Stop** button on the Crystal Ball toolbar or Control Panel,

Continuing Simulations

To continue a simulation, select Run, then Continue [Simulation] or click the Continue button on the Crystal Ball toolbar or Control Panel,

(In Microsoft Excel 2007 or later, choose Run, then Start or click the Start button.)

Resetting and Rerunning Simulations

- To reset the simulation and start over again:
- Select **Run**, then **Reset [Simulation]** or click the **Reset** button on the Crystal Ball toolbar or Control Panel,

A dialog opens with a message asking you to confirm the request to reset the simulation.

- 2 Click OK.
 - Crystal Ball resets the number of trials to 0 and clears the list of values and statistics for each assumption and forecast. However, the assumption and forecast definitions remain.
- 3 Change any assumptions or forecasts as needed.
- 4 Select **Run**, then **Start** [**Simulation**] or click the **Start** button on the Crystal Ball toolbar, The simulation starts over again.

Single-Stepping

Before you run a simulation or after you have stopped it, you can use the Single Step command to watch the simulation process generate one set of values (a *trial*) at a time for the assumption cells and recalculate the spreadsheet. This feature is useful if you are trying to track down a calculation error or verify that the values being produced for the assumption cells are valid.

- To observe an individual trial:
- Select **Run**, then **Reset [Simulation]** or click the **Reset** button on the Crystal Ball toolbar or Control Panel.
- Select Run, then [Single] Step or click the Single-Step or Step button to run one trial of the simulation.
 Click the button again to run another.

The Crystal Ball Control Panel

You can use the Crystal Ball Control Panel to perform many simulation and analysis commands. By default, the Control Panel opens when you run a simulation.

To hide the Control Panel, uncheck **Show Control Panel** on the Options tab of the Run Preferences dialog (click the Run Preferences tool or choose Run, then Run Preferences).

By default, the Control Panel stays open following a reset. To close it, uncheck Leave Open On Reset on the Options tab of the Run Preferences dialog.

For a description of the Control Panel menu commands, see "The Crystal Ball Control Panel Menubar" on page 85.

Figure 20 The Crystal Ball Control Panel



When you click the More button, , simulation statistics are displayed below the controls. They show how fast the simulation ran and how many assumptions, decision variables, and forecasts were included in it.

The Crystal Ball Control Panel Menubar

The following table describes the Control Panel menus and their commands:

Table 4 Crystal Ball Control Panel Menus

Menu	Commands
Run	Start Simulation/Continue Simulation ("Resetting and Rerunning Simulations" on page 84)
	Reset Simulation ("Resetting and Rerunning Simulations" on page 84)
	Single Step ("Single-Stepping" on page 84)
	OptQuest
	Predictor
	More Tools ("Overview" on page 171)
	Save Results ("Saving Crystal Ball Simulation Results" on page 88)
	Restore Results ("Restoring Crystal Ball Simulation Results" on page 89)
	Run Preferences ("Setting Run Preferences" on page 78)
Analyze	Assumption Charts ("Understanding and Using Assumption Charts" on page 151)
	Forecast Charts ("Defining Forecasts" on page 66)
	Overlay Charts ("Understanding and Using Overlay Charts" on page 130)
	Trend Charts ("Understanding and Using Trend Charts" on page 136)
	Sensitivity Charts ("Understanding and Using Sensitivity Charts" on page 142)
	Cascade ("Multiple Windows" on page 86)
	Close All ("Multiple Windows" on page 86)
	Create Report ("Creating Reports" on page 161)
	Extract Data ("Extracting Data" on page 165)
Help	Control Panel Help

Managing Chart Windows

You can close the forecast and other chart windows at any time and the simulation will continue. Running a simulation with the windows closed decreases the time required to run the simulation.

The following sections describe how to work with single and multiple chart windows:

- "Single Windows" on page 86
- "Multiple Windows" on page 86

Single Windows

To close one window, click the Close icon in the corner of the chart window.

Multiple Windows

You can use several commands on the Analyze menu and Crystal Ball toolbar to:

- Close and cascade all open forecast windows and other chart windows.
- Open charts that have been closed.
- Create new charts.
- Delete overlay, trend, and sensitivity charts.

If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Analyze, then View Charts and then choose a command.

You can click a chart button (shown below) to open all windows for that chart type, or use one of the following commands:

- Analyze, then Open Selected Cells (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Open From Selection) Opens charts for all assumption and forecast cells in the selected range. [xxx confirm command names; may be the same now]
- Analyze, then Assumption Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Assumption Charts) — Opens a dialog so you can choose which assumption charts to open or close.
- Analyze, then Forecast Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Forecast Charts) — Opens a dialog so you can choose which forecast charts to open or close.
- Analyze, then Overlay Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Overlay Charts) — Opens a dialog so you can choose which overlay charts to open or close, create, or delete.
- Analyze, then Trend Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Trend Charts) — Opens a dialog so you can choose which trend charts to open or close, create, or delete.
- Analyze, then Sensitivity Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Sensitivity Charts) Opens a dialog so you can choose which sensitivity charts to open or close, create, or delete.
- Analyze, then Scatter Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Scatter Charts) Opens a dialog so you can choose which scatter charts to open or close, create, or delete.
- Analyze, then OptQuest Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then OptQuest Charts) — If OptQuest is available and you have just run an optimization, opens a dialog so you can choose which OptQuest charts to open or close, create, or delete.

- Analyze, then Predictor Charts (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Predictor Charts) If you have just run an time series forecast, opens a dialog so you can choose which Predictor charts to open or close, create, or delete.
- Analyze, then Close All (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Close All) — Closes all chart windows and clears the current simulation results and stored results from memory.
- Analyze, then Cascade (In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Cascade) — Neatly stacks all windows in front of Microsoft Excel.

Note: OptQuest and Predictor chart buttons are not available in Microsoft Excel 2003.

See "Setting Speed Preferences" on page 80 for ways to control whether and how often chart windows display and redraw.

Saving and Restoring Simulation Results

You can save all open forecast windows and other charts as well as simulation data after you run a simulation in Crystal Ball. You can only save results after a simulation stops. Although only results are saved and not an entire model, restored results files are displayed in Crystal Ball chart, report, and Extract Data dialogs so you can work with them. You can run new charts and reports against them and extract their data to worksheets. Only results from the current simulation are saved.

Because the saved files contain only results and are not complete models, you can load more than one results file at a time and you do not need to reset the current simulation before loading results.

The following sections describe how to save, restore, and work with saved simulation results:

- "Saving Crystal Ball Simulation Results" on page 88
- "Restoring Crystal Ball Simulation Results" on page 89
- "Using Restored Results" on page 89
- "Restored Results with Capability Metrics" on page 90

Saving Crystal Ball Simulation Results

- ➤ To save Crystal Ball results:
- 1 Choose Run, then Save Results. The Save Results dialog opens.
 - (In Microsoft Excel 2007 or later, choose Run, then Save or Restore, and then choose Save Results.)
- 2 Navigate to the correct folder to save the results file.
- 3 Name the results file. The default name is the name of the active workbook.

4 Click OK.

The saved results file has a .cbr extension. Crystal Ball saves all results data and charts that existed when the results were saved.

Note: Only results from the current simulation are saved. Previously restored results are not saved. Suppose current and restored forecasts are both used in an overlay chart. If the overlay chart is then saved as part of the current simulation results, when it is restored it will only contain forecasts from the saved current results set. Forecasts from the previously restored results will no longer be included in the chart.

Restoring Crystal Ball Simulation Results

- To restore Crystal Ball simulation results that you saved earlier:
- Select Run, then Restore Results. If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Run, then Save or Restore, then Restore Results.
- 2 When the Restore Results dialog opens, select the results file (.cbr file type) to restore and click Open.

Because you are restoring results and not simulation cell definitions or data, you do not need to reset the simulation before restoring results.

Results files can be restored at any time, regardless of whether the original workbooks are open or whether another simulation has run or not. You can open as many results files as you want — not just one as in Crystal Ball 2000 or earlier versions of Crystal Ball. However, you can only choose one at a time in the Restore Results dialog.

Note: If you try to load results from Crystal Ball 7.3 or later into an earlier version of Crystal Ball, an error message might be displayed, explaining that the Crystal Ball data contained in the selected result set cannot be loaded. One reason this message might be displayed is that a new chart type, scatter charts, was added in Crystal Ball 7.3. In this case, it could be helpful to load the entire model, not just saved results. If you load a model containing scatter charts into an earlier version of Crystal Ball, the model will run and the scatter charts will be ignored.

Using Restored Results

After you have restored one or more Crystal Ball results files, you can open and close restored charts, create new reports using them, and extract their data to spreadsheets. You can create overlay and trend charts with restored results and results from the current simulation to compare data. The results are displayed in dialogs following those for the current simulation.

To remove restored results from memory, choose Analyze, then Close All. In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Close All.

Restored Results with Capability Metrics

If you calculate capability metrics, store the results in a .cbr file, and then restore the results, the restored results use the preference settings on the machine where the results are restored. These might be different from the settings when the original simulation was run and stored. Crystal Ball refits the data when the results are restored, so results might differ somewhat from the original results.

Functions for Use in Microsoft Excel Models

You can use subroutines and functions in the Crystal Ball Developer Kit to automate certain Crystal Ball operations. The following Crystal Ball functions are available as spreadsheet functions for use in Microsoft Excel models:

- CB.GetAssumFN Retrieves information for a specific assumption cell.
- CB.GetAssumPercentFN Returns the value corresponding to a percentile for an assumption cell.
- CB.GetCertaintyFN Returns the certainty level of achieving a forecast value at or below a specific threshold.
- CB.GetForeDataFN Returns the value for the given trial for a specific forecast.
- CB.GetForePercentFN Returns the value corresponding to a percentile for a specific forecast.
- CB.GetForeStatFN Returns statistic for a specific forecast cell.
- CB.GetRunPrefsFN Returns a Run Preference setting.
- CB.IterationsFN Returns the number of trials run in a simulation.
- ➤ To use these functions with Crystal Ball loaded:
- Select an empty spreadsheet cell and select Insert, and then Function. In Microsoft Excel 2007 or later, select Formulas, and then Insert Function.
- 2 Select the Crystal Ball category in the category list box.
- 3 Find the function in the function list. Select it and click **OK**. Always select the function ending in FN, since the other runs significantly slower in most cases.
- 4 In the Function Arguments dialog, enter any required arguments and click OK.

The requested value is displayed in the cell with the function. Be sure to save the workbook with the new function while Crystal Ball is open.

Note: CB.GetForeStatFN only pulls N-1 trials during a recalculation for either a single step or a full simulation, but will reflect full trials at the end of the single step or full simulation since a final recalculation is always performed.

Running User-Defined Macros

You can run user-defined Microsoft Excel VBA macros automatically during a simulation by naming them with predefined names:

- CBBeforeSimulation Runs immediately after the Start Simulation command or Single-Step command
- CBBeforeTrial Runs before random numbers have been placed in assumption cells
- CBAfterRecalc Runs after Microsoft Excel has recalculated the model but before a trial value has been retrieved from the forecast cells
- CBAfterTrial Runs after the forecast trial values have been retrieved and entered into the forecast charts
- CBAfterSimulation Runs when the simulation is complete or stopped with a Stop command, by single-stepping (after each step), or for any other reason

The following figure (Figure 21) shows where each macro fits into the Crystal Ball simulation cycle.




Figure 21 When User-Defined Macros Run

Crystal Ball searches for macros with these names in open workbooks and runs them at the appropriate points during the simulation. The macros must be contained within workbooks and

worksheets (not VBA modules) and Run User-Defined Macros must checked on the Options tab of the Run Preferences dialog.

Crystal Ball checks for the presence of any CBBeforeTrial, CBAfterTrial, and CBAfterRecalc macros before the simulation runs by attempting to execute them. If this poses a problem, you may skip the logic in the macro by first testing for aTrial = 0.

In Extreme speed, only two user-defined macros are allowed: CBBeforeSimulation and CBAfterSimulation. For more information on running user-defined macros in Extreme speed, see "Running User-Defined Macros" on page 298.

The following sections provide details on user-defined macros:

- "User-Defined Macro Interfaces" on page 92
- "Priority Rules" on page 93
- "Global Macros" on page 94
- "Toolbar Macros" on page 94

User-Defined Macro Interfaces

You must use special VBA interfaces for the macros, described in the following sections:

- "CBBeforeSimulation" on page 92
- "CBAfterSimulation" on page 92
- "CBBeforeTrial" on page 93
- "CBAfterTrial" on page 93
- "CBAfterRecalc" on page 93

Note: If a macro returns anything other than zero (0), the simulation stops.

CBBeforeSimulation

```
Public Function CBBeforeSimulation() As Integer the code is here End Function
```

CBAfterSimulation

```
Public Function CBAfterSimulation( _ aIterations As Long, _ aTrials As Long, _ aMaxTrials As Long, _ aProgress As Double, _ aWasAborted As Boolean, _ aWasStopOnError As Boolean, _ aWasStoppedOnPrecision As Boolean, _ aHasNewBestSolution As Boolean) _ As Integer
```

```
'the code is here 
End Function
```

In this interface, parameters or arguments are defined as follows:

- alterations = The number of the simulation in a multi-simulation run, usually 1
- aTrials = Number of trials run
- aMaxTrials = Number of trials scheduled to run
- aProgress = Percent of trials completed
- aWasAborted = True if manually stopped
- aWasStopOnError = True if stopped on error
- aWasStoppedOnPrecision = True if stopped on precision control
- aHasNewBestSolution = New value of objective

CBBeforeTrial

```
Public Function CBBeforeTrial(aTrial As Long) As Integer 'the code is here
End Function
```

CBAfterTrial

```
Public Function CBAfterTrial (aTrial As Long) As Integer 'the code is here
End Function
```

CBAfterRecalc

```
Public Function CBAfterRecalc (aTrial As Long) As Integer 'the code is here
End Function
```

Priority Rules

The following rules govern the running order of macros:

If multiple workbooks are open, Crystal Ball searches the workbooks in the order of their
priority. The priority number is set by the CB.SetCBWorkbookPriority macro. If this macro
has not been called on the workbook, no workbook order is guaranteed.

CB.SetCBWorkbookPriority sets the running order for macros when multiple workbooks are open. For example, a workbook with order=10 will run before a workbook with order=20. CB.SetCBWorkbookPriority runs only at the workbook level and not at the worksheet level.

The definition of this macro is:

```
Public Function CB.SetCBWorkbookPriority (aPriority As Long)
```

A syntax example, used in a subroutine at the module level, is:

```
Application.Workbooks("Book1.xls").Activate
CB.SetCBWorkbookPriority(10)
```

This example assigns a priority of 10 to the workbook Book1.xls.

- Within each workbook, macros on worksheets run in left-to-right sheet order in the
 workbook provided there are no macros at the workbook level. If there are macros at the
 workbook level, macros in that workbook at the worksheet level do not run.
- If there are no Crystal Ball data cells in a workbook, no macros run against it.
- Macros in modules are not supported.

Global Macros

If you want to run certain macros with any or all of the Crystal Ball models, you need to make sure that the workbook containing those macros is open in Microsoft Excel. Then, you can use calls within each model to run those macros as required.

Previous versions of Crystal Ball handled global macros differently. Now, you can run specific macros only for certain models and are no longer required to turn the macros on and off depending on the model that is being run.

Toolbar Macros

If you want to use custom Microsoft Excel toolbar macros defined in a separate workbook while Crystal Ball is loaded, be sure to save that workbook to the folder C:\Documents and Settings\thename\Application Data\Microsoft\Excel\XLSTART.

6

Analyzing Forecast Charts

In This Chapter

Introduction	95
Guidelines for Analyzing Simulation Results	95
Understanding and Using Forecast Charts	97
Using Additional Forecast Features	112
Setting Chart Preferences	115
Managing Existing Charts	124

Introduction

This chapter describes how to use forecast charts to analyze simulation results. As you read through the chapter, you will learn to focus on a particular range of simulation results, analyze forecast charts, and interpret relevant descriptive statistics.

You will also learn how to customize forecast charts, work with chart windows, and share charts with other applications.

Chapter 7 introduces additional charts. Chapter 8 describes how to extract Crystal Ball simulation data for use in Microsoft Excel and other applications and how to prepare reports.

Guidelines for Analyzing Simulation Results

While a simulation runs, Crystal Ball creates a forecast chart for each forecast cell. Forecast charts condense much information into a small space. You can display that information in several different ways, both graphically and numerically.

When you analyze a simulation, you can focus on forecast charts. You can also display other kinds of charts, generate reports, and extract data for further processing using Microsoft Excel or other analysis tools.

The following steps can help guide the simulation analysis by focusing on details as well as general trends:

- 1. Look at the "big picture."
 - Consider each forecast chart from a high-level viewpoint. Look at the shape of the distribution:
 - Is it distributed normally or skewed negatively or positively?

- Is it "flat" (spread out on both sides of the mean) or "peaked" (with most values clustered closely around the mean)?
- Does it have a single mode (most likely value) or is it bimodal with several "peaks" or "humps"?
- Is it continuous or are there groups of values separated from the rest, maybe even extreme values that fall outside of the display range?

The statistical concepts in the "Statistical Definitions" chapter of the online *Oracle Crystal Ball Statistical Guide* can help with this part of the analysis.

2. Look at the certainty level, the probability of achieving values within a certain range.

You can enter a range, such as all values greater than \$0 dollars if you're analyzing profits, and view the certainty of falling within the range (\$0 to + Infinity, in this case). You can also enter a certainty, say 75%, and see what range of values would be required to meet that level. "Determining the Certainty Level" on page 98 explains these and other ways of analyzing certainty.

3. Focus on the display range.

You can change the display range to focus on different sections of the forecast chart. For example, you can set the display range to focus on just the upper or lower tail of the forecast. "Focusing On the Display Range" on page 102 discusses this analysis technique.

4. Look at different views of the forecast.

Use the View menu to switch among different ways of viewing the forecast distribution graphically (frequency, cumulative frequency, or reverse cumulative frequency) or numerically (statistics, percentiles, goodness of fit data, or capability metrics). You can also choose whether to show charts and statistics simultaneously or separately. For instructions, see "Changing the Distribution View and Interpreting Statistics" on page 104.

5. Customize the forecast chart.

Use the chart preferences to change graphic presentations from bars to areas or lines, or experiment with different colors, 2D versus 3D, more or fewer plotted intervals or data points, and other display variations. "Setting Chart Preferences" on page 115 describes the many ways you can format charts for analysis and presentation of results.

6. Create other kinds of charts.

You can display overlay charts ("Understanding and Using Overlay Charts" on page 130), trend charts ("Understanding and Using Trend Charts" on page 136), sensitivity charts ("Understanding and Using Sensitivity Charts" on page 142), assumption charts ("Understanding and Using Assumption Charts" on page 151), and scatter charts ("Understanding and Using Scatter Charts" on page 153). As with forecast charts, you can customize these charts as well. Selecting different views of the data can help you analyze it and present it to others.

7. Create reports.

Create reports with charts and data ("Creating Reports" on page 161). You can quickly produce pre-defined or custom reports for analysis and presentation.

8. Extract data for numeric analysis and presentation.

You can extract simulation results to Microsoft Excel for further analysis or for further export into other analytical tools. See "Extracting Data" on page 165 for details.

9. Use other Crystal Ball tools for different types of analyses.

Depending on the types of problems you are analyzing, the Crystal Ball Batch Fit, Bootstrap, Correlation Matrix, Decision Table, Scenario Analysis, Tornado Chart, and 2D Simulation tools offer addition analysis techniques. These tools are described in Chapter 9. You can use Predictor for time-series analysis and, if you have OptQuest, you can use it to optimize simulation results.

Understanding and Using Forecast Charts

Forecast charts use frequency distributions to show the number (frequency) of values occurring in a given interval. The highest value on the frequency scale of the chart is the frequency for the interval that contains the greatest number of forecast values — the mode of the graphed distribution. You can estimate the frequency for other points on the forecast chart using this frequency scale.


Note: You can now display forecast charts in Split View, which shows charts and statistics side by side. For more information about Split View and its features, see "Using Split View" on page 109.

The highest value on the Probability scale is the probability for the mode. You can estimate the probability for other points on the forecast chart using this probability scale.

In the example below, the mode (the x-axis value that occurred the most frequently) has a frequency of about 300, meaning that there are 300 values in the interval expressed by that column. The mode has a probability of about 0.06 (or 6%), meaning that there is a 6% chance of a value falling within this interval.

Figure 22 shows the elements of the forecast chart.

Figure 22 Forecast chart



Crystal Ball forecasts the entire range of results for a given situation. However, the forecast chart shows only the display range, which is a subset of the whole. By default, the display range includes all trials within 2.6 standard deviations of the mean (approximately 99% of the forecast values). Crystal Ball then rounds the display range to the next even number of units. For this reason, outlying trials might be excluded from the display range.

Note: To display all trials, change the chart axis preferences to display fixed endpoints between –Infinity and +Infinity ("Focusing On the Display Range" on page 102).

The number of trials run for a forecast are displayed at the top of the forecast chart, near the Probability scale. The number of trials actually shown (the number of trials run minus the number of outlying values) are displayed at the top of the chart near the Frequency scale (right vertical axis). The display range is the linear distance between the minimum and maximum trials on the chart (from about -\$15.00 to \$35.00 in Figure 22).

The following sections provide details about reading and analyzing forecast charts:

- "Determining the Certainty Level" on page 98
- "Focusing On the Display Range" on page 102
- "Formatting Chart Numbers" on page 103
- "Changing the Distribution View and Interpreting Statistics" on page 104
- "Setting Forecast Preferences" on page 111
- "Setting Forecast Chart Preferences" on page 112

Determining the Certainty Level

The forecast chart also shows the certainty range for the forecast, all trials between the certainty grabbers. In the next section, you will learn how to change the certainty range.

Note: When the certainty grabbers are at –Infinity and +Infinity, the certainty range includes every forecast value regardless of the size of the display range.

Crystal Ball compares the number of values in the certainty range with the number of values in the entire range to calculate the certainty level for the forecast. The example in Figure 22 shows a certainty level of 100% since the default certainty range includes all possible values. By default, Crystal Ball calculates the certainty level based on the entire range of forecast values.

On either side of the Certainty field, the certainty minimum and maximum are displayed.

The certainty level is one of Crystal Ball's key statistics because it shows the probability of achieving the values within a specific range. With Crystal Ball, you can determine the certainty level for specific value ranges either by moving the certainty grabbers on the forecast chart or typing the certainty minimum and maximum in the fields. You can also type a certainty level in the certainty field to get a certainty range centered around the median.

When you move the certainty grabbers, the certainty range changes and Crystal Ball recalculates the certainty level. When you type minimum and maximum values, Crystal Ball moves the certainty grabbers for you and recalculates the certainty level. When you type the certainty level in the certainty field, Crystal Ball moves the certainty grabbers to show you the value range for the certainty level you specified.


The following sections describe how to adjust and analyze certainty levels in Crystal Ball:

- "Moving Certainty Grabbers" on page 99
- "Changing the Certainty Minimum and Maximum Fields" on page 101
- "Entering Certainty Directly" on page 101
- "Anchoring a Grabber and Entering Certainty Directly" on page 101
- "Resetting the Certainty Range" on page 102

Moving Certainty Grabbers


- To determine the certainty level for a specific value range using certainty grabbers:
- 1 Make sure the forecast chart to use is the active (selected) window.
- 2 Move the certainty grabbers on the forecast chart, as shown in Figure 23.
 - Click the grabber and then drag it when the cursor changes to a hand.

Figure 23 Moving a Certainty Grabber



The certainty minimum field shows the value that corresponds to the position of the lowervalue (left) certainty grabber; the certainty maximum field shows the value that corresponds to the position of the higher-value (right) certainty grabber. The Certainty field shows the certainty level for the area between the certainty grabbers. Crystal Ball shades the columns outside the certainty grabbers a different color to show that those values have been excluded, as illustrated in Figure 24.

Figure 24 Certainty level: values \$0 to \$20 (in Millions)



The Net Profit forecast chart in Figure 24 is the same as the example preceding it, except that the certainty grabbers have been moved. The certainty minimum shows \$0.0 and the certainty maximum shows \$20.0. The key statistic is the certainty level of 65.6% in the Certainty field. By moving the certainty grabbers, you have changed the certainty range. Crystal Ball compares the number of values lying within the certainty range to the number of values in the entire range from negative \$14.7 million to positive \$34.4 million — to recalculate the certainty level. With a certainty level of 65.6%, you can be 65.6% confident of making a net profit between \$0 and \$20 million.

Changing the Certainty Minimum and Maximum Fields

- To determine the certainty level for a specific value range using the certainty minimum and maximum fields:
- 1 In the forecast chart, type a value in the certainty minimum field.
- 2 Type a value in the certainty maximum field.
- Press Enter.

The certainty grabbers move to correspond to the values you entered.

Entering Certainty Directly

Alternately, you can type the certainty level on the forecast chart and allow Crystal Ball to calculate the minimum and maximum for you. Make sure that the certainty grabbers are free (whenever you click an certainty grabber, it turns to a lighter color and is considered anchored). To free the lighter color certainty grabber, click it. Then when you type in a certainty level, you will center the certainty range on the median statistic.

Note: To free or anchor both certainty grabbers, press Ctrl+click or Shift+click.

- In the following procedure, you can enter a certainty level of 10% and watch Crystal Ball move the certainty grabbers:
- Press Ctrl+click (or Shift+click) until both certainty grabbers are free (a lighter color).
- 2 Type 10 in the Certainty field.
- Press Enter.

Crystal Ball moves the certainty grabbers to include 10% of the trials centered around the median statistic of the entire range.

Anchoring a Grabber and Entering Certainty Directly

You can also anchor a certainty grabber and type the certainty level. Then, Crystal Ball moves the free grabber to correspond to the value range for the level.

- To anchor a certainty grabber:
- Move the lowest-value (left) certainty grabber until the certainty minimum value is \$0.0.

The certainty grabber turns to a lighter color and is considered anchored.

- Type 75 in the Certainty field.
- Press Enter.

Crystal Ball moves the free certainty grabber toward the anchored certainty grabber to include 75% of the values as in Figure 25.

If both grabbers are free, the distribution is centered on the median.

Forecast: Net Profit (MM) Edit View Forecast Preferences Help 5,000 Trials Frequency View 4,993 Displayed Net Profit (MM) 0.07 320 0.05 240 200 Teg 0.04 160 PG 120 V 0.03 120 0.02 80 40

\$10.0 Millions

Certainty: 75.00

Figure 25 Certainty Level: 75% from Lowest-value (Left) Certainty Grabber

You also can cross over the certainty grabbers to determine the certainty level for the two tails (ends).

\$27.4

You can determine the certainty level for specific value ranges at any time, either during or after the simulation.

Resetting the Certainty Range

- To reset the original certainty range before beginning the next section:
- 1 In the forecast chart, either:

0.00

\$0.0

- Move the higher-value (right) certainty grabber until the certainty maximum field shows positive infinity, or
- Type +Infinity in the right (maximum) field (accessible with the Tab key).
- 2 Either:
 - Move the lowest-value (left) certainty grabber until the certainty minimum field shows negative infinity, or
 - Type –Infinity in the left (minimum) field, accessible with the Tab key.
 The forecast chart displays a certainty level of 100%, as displayed in the original forecast.

Focusing On the Display Range

With Crystal Ball, you can focus on a particular range of the forecast results by changing axis settings in the Chart Preferences dialog. For instructions, see "Customizing Chart Axes and Axis Labels" on page 123.

- To define the display range:
- 1 In a forecast chart, select Chart Preferences, then Axis, then Scale, then Type.

Choose a scale type:

- Auto Crystal Ball uses a default display range of 2.6 standard deviations from the mean, which includes about 99% of the forecast values. (See Standard Deviation, following.)
- Fixed Sets the display range end points manually so you can focus on particular value ranges. For example, you can focus on positive values only to look at the profit for a profit/loss forecast.
- **Standard Deviation** Sets the display range end points in terms of standard deviations; defines the number of standard deviations worth of value to display on each side of the mean and centers values around the mean. If you choose to set the display range in terms of standard deviations, you can change the display range to 1 standard deviation from the mean to look at approximately 68% of the forecast values.
- Percentile Sets the display range end points in terms of percentiles.
 - By default, the x-axis value numbers are automatically adjusted to round numbers to make the forecasts easier to read. The Axis chart preference settings include Round Display Range, which constrains the display range to round numbers. You can uncheck this setting to display actual, unrounded numbers.

Other chart customizations are available to help you interpret the results of the simulation by viewing the data in different ways. For details, see "Setting Chart Preferences" on page 115.

Also see "Setting Preferences with Shortcut Keys" on page 116 for ways to change the appearance of charts without using menu commands.

Showing Statistics for the Display Range

After you have changed the display range, you can display statistics for just that range.

- To show statistics for a display range:
- 1 Set the display range as described in "Focusing On the Display Range" on page 102.
- Note the values for the display range minimum and maximum.
- In the forecast chart menubar, choose Preferences, then Forecast, then Filter.
- 4 On the Filter tab of the Forecast Preferences dialog, set a filter on the forecast values and include values in the range between the display range minimum and maximum.
- 5 When the settings are complete, click OK.
- 6 Choose View, then Statistics in the forecast chart menubar to show statistics for the display range (or, in Split View, look at the statistics table).

Formatting Chart Numbers

By default, the number format displayed on the forecast chart comes from the underlying format of the forecast cell. You can select another cell format using the Chart Preferences dialog.

- To change the format of numbers in a forecast chart:
- 1 In the forecast window, choose Preferences, then Chart.
 - The Chart Preferences dialog opens.
- 2 Click the Axis tab.
 - Find the Format Number group.
- 3 Choose a format from the drop-down list. Available formats are similar to Microsoft Excel formats. For most formats, you can specify the number of decimal places and whether to use a thousands separator.
- 4 Click OK to apply the settings, or use Apply To as described in "Applying Settings to the Current Chart and Other Charts" on page 124.

Changing the Distribution View and Interpreting Statistics

The forecast settings related to distribution type determine the overall appearance of a forecast chart. You can also choose to display a table of statistics or percentiles instead of or in addition to a chart.

- To set the distribution type or display a data table:
- 1 Open the View menu in the forecast window.
- 2 Choose a distribution type or other view to display on the forecast chart:
 - Frequency Shows the number or frequency of values occurring in a given interval. This is the default distribution type.
 - Cumulative Frequency Shows the number or proportion (percentage) of values less than or equal to a given amount.
 - Reverse Cumulative Frequency Shows the number or proportion (percentage) of values greater than or equal to a given amount.
 - Statistics Shows a full set of descriptive statistics for a simulation in the forecast window.
 - Percentiles Shows percentile information in 10% increments, where a percentile is the percent chance, or probability, of a forecast value being less than or equal to the value that corresponds to the percentile (by default).
 - Goodness Of Fit If distribution fitting is selected in the Forecast or Preferences, then Forecast menus, shows goodness-of-fit statistics for the selected distributions and ranking methods.
 - Capability Metrics If process capability metrics are set for display, shows a table of
 process capability (quality) statistics for the simulation ("Viewing Capability Metrics"
 on page 372).
 - Split View Shows all selected views simultaneously ("Using Split View" on page 109).

See "View Examples" on page 105 for descriptions and illustrations of each view.

View Examples

The following sections provide descriptions and illustrations of each view:

- "Frequency" on page 105
- "Cumulative Frequency" on page 105
- "Reverse Cumulative Frequency" on page 106
- "Statistics" on page 107
- "Percentiles" on page 107
- "Goodness Of Fit" on page 108
- "Capability Metrics" on page 108"Split View" on page 109

Frequency

Frequency, the default forecast view, shows a simple count of values (the frequency) for each interval on the x-axis. Figure 26 shows a frequency chart of net profit values for a simulation where there is a 75% probability of net profit falling between \$0.00 and \$27.4 million. The chart has a median of \$8.4 million. This is also the 50th percentile. By default, there is a 50% probability that net profit will be at or below this value.





Figure 26 Forecast Chart - Frequency

Cumulative Frequency

Figure 27 shows the Net Profit forecast chart as a cumulative distribution. This chart shows the number or proportion (percentage) of values less than or equal to a given amount.

Figure 27 Forecast Chart—Cumulative Frequency




To create this chart, the frequencies are added cumulatively, starting from the lower end of the range, and then plotted as a cumulative frequency curve. To understand the cumulative distribution, look at a particular value, \$8.4 million (in the previous example). The chart shows that the probability of \$8.4 million is about 50%; approximately 50% of the values are less than \$8.4 million, while approximately 50% are greater. This would be correct for a median value. Note also that the chart shows that the probability for \$27.4 million is about .95 while the probability for \$0 is about .20. This is also correct, since the probability of Net Profit falling between those two values is .75 (.95 - .20 = .75) or Certainty = 75%.

Reverse Cumulative Frequency

Figure 28 shows the Net Profit forecast chart as a reverse cumulative distribution. This chart shows the number or proportion (percentage) of values greater than or equal to a given amount.

Figure 28 Forecast Chart—Reverse Cumulative Frequency




To create this chart, the frequencies are added cumulatively starting at the higher end of the range, and then plotted as a declining cumulative frequency curve. To understand the reverse cumulative distribution, look at a particular value, \$8.4 million (in the previous example). The chart shows that the probability of \$8.4 million is about 50%; approximately 50% of the values are less than \$8.4 million, while approximately 50% are greater. This would be correct for a median value. Note also that the chart shows that the probability for \$27.4 million is about .05

(of having a greater value) while the probability for \$0 is about .80. This is also correct, since the probability of Net Profit falling between those two values is .75 (.80 - .05 = .75) or Certainty = 75%. Notice in this chart that the reverse cumulative frequency values are complements of the cumulative frequency values: .20 + .80 = 1.00 and .95 + .05 = 1.0 (the probability values for \$0.0 and \$27.4 million, respectively).

Statistics

You can display a full set of descriptive statistics for a simulation in the forecast window by choosing View, then Statistics.

Figure 29 Forecast Window-Statistics



The example in Figure 29 shows statistics for the entire range of values (100% of the forecast values, including extreme values excluded from the default display range). Statistical terms listed in this table are discussed in the "Statistical Definitions" chapter of the Oracle Crystal Ball Statistical Guide and the Glossary in this User's Guide.


Note: If the Precision Control feature is checked in the Run Preferences dialog and the forecast has Precision Control options set, the Precision column opens in the Statistics view.

Percentiles

You can display percentile information in 10% increments in the forecast window by choosing View, then Percentiles. A percentile is the percent chance, or probability, of a forecast value being less than or equal to the value that corresponds to the percentile (the default). For example, Figure 30 displays the percentile view of the Net Profit forecast, where the 90th percentile corresponds to \$21.6 million, meaning that there is a 90% chance of a forecast value being equal to or less than \$21.6 million. Another interpretation is that 90% of the forecast values are equal to or less than \$19.3 million.

Notice that the Median in Statistics view is the same as the 50th percentile in Percentiles view — in this case, \$8.4 million.

Figure 30 Forecast-Percentiles View




If the Precision Control feature is checked in the Run Preferences dialog and the forecast has Precision Control options set, the Precision column is displayed in the Percentiles view.

Goodness Of Fit

If you have selected distribution fitting, described in the next section, you can choose the Goodness Of Fit view to display comparative fit statistics for each of the selected distribution types. The distributions are ordered according to the selected ranking method. Figure 31 shows statistics for the Anderson-Darling ranking method and each continuous distribution type. Notice that Beta is ranked highest for this forecast.


Figure 31 Forecast-Goodness Of Fit View



Capability Metrics

If the process capability features are activated on the Statistics tab of the Run Preferences dialog and if either an LSL, USL, or both are entered into the Define Forecast dialog, the Capability Metrics view is available for the forecast chart. For a definition of each statistic, see "Capability Metrics List" on page 378.

Figure 32 Capability Metrics View



Split View

Split View shows forecast charts and related statistics onscreen at the same time. For more information, see "Using Split View" on page 109.

Using Split View

Split View displays charts and statistics at the same time. If you are using Crystal Ball's process capability features, Split View is the default. Otherwise, you can activate Split View through the View menu or the Preferences, then Forecast menu in the forecast chart window.

- ➤ To activate Split View:
- In the forecast window, choose View to open the View menu.
- Choose Split View at the bottom of the menu.

The Frequency chart and statistics are both displayed, similar to the next figure.

Profit or Loss Edit View Forecast Preferences Help 5.000 Trials Solit View 4,994 Displayed **Profit or Loss** Forecast values Statistic 5.000 Mean \$2,473.72 0.05 \$2,486.58 240 Median Mode 0.04 200 \$1,850,66 Standard Deviat bability 160 requency Variance \$3,424,956,14 0.03 Skewness 0.0053 년 0.02 Kurtosis 2.40 0.7481 Coeff. of Variabi 80 Minimum (\$2,732.96) Maximum \$8.275.85 40 Mean Std. Error \$26.17 0.00 vertical pane \$6,000.00 \$3,000.00 splitter Monthly rental income Certainty: 100.00 Infinity -Infinity

Figure 33 Frequency Chart and Statistics in Split View

You can resize the window and use the vertical pane splitter to adjust the size of the chart and the statistics pane.

3 If you want, continue opening the View menu and choosing charts or data.

The following figure shows a Frequency chart, a Cumulative Frequency chart, plus Statistics and Percentiles tables.




Figure 34 Frequency and Cumulative Frequency Charts with Statistics and Percentiles in Split View

You can click in any of the Split View panes and use the chart shortcut keys to modify them without using the View or Preferences menus. For a list, see Table 5.

You can also resize the chart window and drag the horizontal and vertical pane splitters to resize each part of the Split View window.

To clear Split View or remove any of the views from the window, open the View menu and uncheck each view you want to close.

Setting Forecast Preferences

You can set a number of specific forecast preferences to customize how Crystal Ball calculates and displays forecast charts. These are in addition to the general chart preferences discussed in "Setting Chart Preferences" on page 115.

You can control several forecast features with preference settings, described in the section in parentheses. You can:

- Change the forecast chart view ("Forecast Window Tab" on page 68)
- Determine when the forecast window opens ("Forecast Window Tab" on page 68)
- Fit a distribution to the forecast ("Fitting a Distribution to a Forecast" on page 113)
- Set precision controls for forecast statistics ("Precision Tab" on page 68)
- Filter ranges of forecast values ("Filter Tab" on page 69)
- Automatically extract forecast data to a spreadsheet ("Auto Extract Tab" on page 70)

For an overview, see "Basic Instructions for Setting Forecast Preferences" on page 111.

Basic Instructions for Setting Forecast Preferences

Like the chart preferences, forecast preferences can be set differently for each forecast chart.

- To set forecast preferences:
- 1 Choose Preferences, then Forecast in a forecast chart's menubar.

The Forecast Preferences dialog opens.

Click a tab and set preferences as required.

These tabs are also displayed in the expanded Define Forecast dialog. For details see the following sections:

- "Forecast Window Tab" on page 68 Manages window display and distribution fitting for the forecast.
- "Precision Tab" on page 68 Manages precision control settings.
- "Filter Tab" on page 69 Discards values inside or outside a range for the current
- "Auto Extract Tab" on page 70 Specifies which statistics to extract automatically to Microsoft Excel when a simulation stops.

Also see the previous section, "Setting Forecast Preferences" on page 111.

See "Confidence intervals" in the "Statistical Definitions" chapter of the online *Oracle Crystal Ball Statistical Guide* for more information about how absolute and relative precision relate to the confidence interval.

- 3 To reset forecast preference defaults to the original settings shipped with Crystal Ball, click Defaults.
- 4 To copy preferences to other forecasts, click **Apply To**.

You can specify whether to apply all forecast preferences or just the current tab and whether to apply them to the current Microsoft Excel sheet, all sheets in the workbook, or all open and new workbooks.

5 When all settings are complete, click **OK** to apply them.

Setting Forecast Chart Preferences

The chart preferences can help you customize the appearance of forecast charts. To make these settings, choose Preferences, then Chart Preferences in the forecast chart menubar. For details, see "Setting Chart Preferences" on page 115.

The following preference settings help interpret charts:

- Chart type Shows forecast charts as columns, areas, or lines in two or three dimensions so you can view the data plot from different perspectives and grasp "the big picture" more easily
- Chart density Increases and decreases the number of bars or data points so you can spot trends more easily.
- Grid lines Make it easier to determine frequencies and probabilities.
- Marker lines Make it easier to locate means, medians, modes, percentiles, and other important values.
- Axis scaling and rounding Shows more or fewer axis values to read chart frequencies and probabilities more easily.

You can copy forecast charts and paste them into other applications. For more information, see "Copying and Pasting Charts to Other Applications" on page 125.

Using Additional Forecast Features

Previous topics in this chapter described how to analyze forecast charts by switching among different views, how to set forecast preferences, and how to set forecast chart preferences. The topics in this section include "Fitting a Distribution to a Forecast" on page 113 and "Defining Assumptions from Forecasts" on page 114.

You can also create a sensitivity chart from within a forecast chart to show which assumptions contributed the most to that forecast distribution. For details, see "Understanding and Using Sensitivity Charts" on page 142.

Fitting a Distribution to a Forecast

You can fit distributions to forecast charts, a feature that is similar to distribution fitting for assumptions (described in "Fitting Distributions to Data" on page 46).

You can fit distributions to forecasts two ways:

- You can choose Forecast, then Fit Probability Distribution in the forecast chart menubar to do a quick fit with the default or currently selected distributions and ranking method. You can also use this command to switch off distribution fitting that is set with either the Forecast menu or Preferences menu.
- You can choose Preferences, then Forecast, then Forecast Window in the forecast chart menubar to specify particular distributions and to choose a fit ranking method. Then, you can also change the fit options or use Apply To to set these preferences for other forecasts.
- To fit a probability distribution to a forecast chart using the Preferences, then Forecast command:
- Create a model and run a simulation.
- Select a forecast chart.
- 3 In the forecast chart menubar, choose Preferences, then Forecast.
- In the Forecast Window tab of the Forecast Preferences dialog, check Fit A Probability Distribution To The Forecast and click Fit Options.
 - The Fit Options panel opens.
- Specify which distributions to fit:
 - AutoSelect performs a basic analysis of the data to choose a distribution fitting option and ranking method. If the data includes only integers, fitting to all discrete distributions (with the exception of Yes-No) is completed using the Chi-square ranking statistic choice.
 - All Continuous fits the data to all of the built-in continuous distributions (these distributions are displayed as solid shapes on the Distribution Gallery).
 - All Discrete fits to all discrete distributions except Yes-No.
 - Choose displays another dialog from which you can select a subset of the distributions to include in the fitting.
- Specify how the distributions should be ranked. In ranking the distributions, you can use one of three standard goodness-of-fit tests:
 - Anderson-Darling This method closely resembles the Kolmogorov-Smirnov method, except that it weights the differences between the two distributions at their tails greater than at their mid-ranges. This weighting of the tails helps to correct the Kolmogorov-Smirnov method's tendency to over-emphasize discrepancies in the central region.
 - Chi-Square This test is the oldest and most common of the goodness-of-fit tests. It gauges the general accuracy of the fit. The test breaks down the distribution into areas

- of equal probability and compares the data points within each area to the number of expected data points.
- Kolmogorov-Smirnov The result of this test is essentially the largest vertical distance between the two cumulative distributions.
 - The first setting, **AutoSelect**, enables Crystal Ball to choose the ranking statistic. If all data values are integers, Chi-square is chosen.
- If you know location, shape, or other parameter values that might help create a more accurate fit with certain distributions, check Lock Parameters and enter appropriate values in the Lock Parameters dialog ("Locking Parameters When Fitting Distributions" on page 49).
- 8 By default, values for all appropriate ranking statistics are calculated but only values for the selected ranking statistic are displayed in Goodness Of Fit view. To show values for all three statistics, check Show All Goodness-of-fit Statistics at the bottom of the Distribution Options panel.
- 9 Click OK to perform the fit.

During a simulation, Crystal Ball disables distribution fitting on forecast charts and overlay charts after 1,000 trials and until the simulation stops to enhance performance. A final fit is performed at end of the simulation.

Defining Assumptions from Forecasts

It is sometimes convenient to use results from one simulation as inputs to another simulation. For example, the simulation results from a departmental revenue model can be used as input assumptions to a total company revenue model. It is not necessary for the two models to share the same simulation. You can use the Define Assumption from Forecast feature in Crystal Ball to convert a forecast distribution into an assumption in one of two ways. You can either fit a standard probability distribution to the forecast data, or use the forecast data directly as a custom distribution.

- To define an assumption from a forecast:
- 1 Run a Crystal Ball simulation and open a chart for the target forecast.
- 2 In the forecast chart menubar, select Forecast, and then Define Assumption from Forecast.
- 3 In the Define Assumption from Forecast dialog:
 - Enter a cell location for the new assumption. You can type it or click the cell selector to point to the cell.
 - Select a distribution type for the assumption. You can select the best fitting distribution or define a custom distribution.
 - o If you select **Best fitting distribution**, the current fitting defaults are used. You can click **Fit Options** to display the **Fit Options** dialog described in "Fitting a Distribution to a Forecast" on page 113. If you select **Show comparison chart**, you can view a chart of each fitted distribution and optionally override the best fit selection.

- If you select Custom distribution (with forecast data), when you click OK, the Define Assumption dialog opens for the custom distribution. It contains the data taken from the filtered range (if any) of the forecast. If you prefer, you can follow the instructions in "Using the Custom Distribution" on page 270 to modify the data.
- Optional: If you select Custom distribution (with forecast data), you can select Sample sequentially (instead of randomly) to use sequential sampling. For details, see "Sequential Sampling with Custom Distributions" on page 291.
- Optional: Indicate whether to display the new assumption's type and parameters in adjacent cells, either downward or to the right, and indicate whether to show names (labels) with the values.
- When settings are complete, click OK to perform distribution fitting (unless you selected Custom distribution) and open the Define Assumption dialog.
 - You can save the assumption with the indicated defaults, or modify the assumption as if you were defining it in the usual way. Most assumption definition features are available. You can enter different parameters and correlate the assumption with another. However, you cannot add the assumption to the Gallery until after it has been created.
- To finish defining the assumption, click **OK** in the **Define Assumption** dialog.

After you have defined the new assumption, you can select it and choose Define Assumption to change the distribution type or perform any other modifications, including adding it to the Gallery.

Setting Chart Preferences

You can set a number of chart preferences to customize the appearance of Crystal Ball charts. The following customizations can help you analyze and present the data, followed by references to instructions:

- Add or edit and format a title ("Adding and Formatting Chart Titles" on page 118)
- Change the chart type ("Setting the Chart Type" on page 118)
- Show more or fewer columns or data points ("Changing the Chart Density" on page 120)
- Show or hide grid lines ("Showing Grid Lines" on page 120)
- Show or hide the chart legend ("Showing the Chart Legend" on page 121)
- Set special chart effects such as transparency or 3D lines, areas, and columns ("Setting Special Chart Effects" on page 121)
- Set chart colors ("Setting Chart Colors" on page 122)
- Show mean, median, mode, standard deviation, percentile, or capability limit/target marker lines ("Showing the Mean and Other Marker Lines" on page 122)
- Hide and show vertical and horizontal axes, create and edit axis labels, and change an axis scale ("Customizing Chart Axes and Axis Labels" on page 123)

- Format chart numbers ("Formatting Chart Numbers" on page 103)
- Specify whether to use these preferences for more than the current chart ("Applying Settings to the Current Chart and Other Charts" on page 124)

Also see "Setting Preferences with Shortcut Keys" on page 116, for ways to change the appearance of charts without using menu commands. See "Basic Customization Instructions" on page 117 and "Specific Customization Instructions" on page 117 for other general customization tips.

Setting Preferences with Shortcut Keys

Table 5 lists key combinations that can be used to cycle through settings available in the Chart Preferences dialog. Most of these commands work on the primary distribution — the probability distribution for assumptions, and frequency for forecasts and overlay charts.

Note: You can also use Ctrl+(view number) to cycle through views in Split View, while Ctrl+ (chart number) cycles through multiple open charts. For more information, see "Using the Tab and Arrow Keys in the Crystal Ball User Interface" on page 390.

Table 5 Shortcut keys for Chart Preferences

Shortcut Key	Command Equivalent	Description	
Ctrl+d	View menu; Preference , then <i>chartname</i> Preferences , then View	Cycles through chart views — Frequency, Cumulative Frequency, Reverse Cumulative Frequency (for assumption and forecast charts)	
Ctrl+b; Ctrl+g	Preferences, then Chart Preferences, then General, then Density	Cycles through bin or group interval values to adjust the number of columns or data points	
Ctrl+l	Preferences, then Chart Preferences, then General, then Gridlines	Cycles through gridline settings: None, Horizontal, Vertical, Both	
Ctrl+t	Preferences, then Chart Preferences, then Chart Type, then Type	Cycles through chart types: Area, Line, Column; for sensitivity charts: Bar (direction), Bar (magnitude), Pie (in Contribution To Variance view)	
Ctrl+w	Preferences, then Chart Preferences, then General, then 3D Chart	Cycles between two-dimensional and three-dimensional chart display	
Ctrl+m	Preferences, then Chart Preferences, then Chart Type, then Marker Lines, then central tendencies	Cycles through central tendency marker lines: None, Mean, Median, Mode (except for sensitivity and trend charts)	
Ctrl+n	Preferences, then Chart Preferences, then General, then Legend	Toggles the legend display on and off	
Ctrl+p	Preferences, then Chart Preferences, then Chart Type, then Marker Lines, then Percentiles	Cycles through percentile marker lines: None, 10%, 20%, 90%	

Shortcut Key	Command Equivalent	Description
Ctrl+Spacebar	View menu; Preferences, then <i>chartname</i> Preferences	Cycles through window views: Chart, Statistics, Percentiles, Goodness Of Fit (if distribution fitting is selected — except for trend charts)

Basic Customization Instructions

These instructions apply most specifically to forecast charts. However, many of them apply to other charts as well. For that reason, they are as general as possible, although not all settings will apply to every type of chart.

- To customize a chart:
- 1 Create or display a chart and be sure it is the active chart window.
- Either double-click the chart or choose Preferences, then Chart Preferences from the chart's menubar.

The Chart Preferences dialog opens. Tab contents are as follows:

- General Chart title, overall appearance of the chart
- Chart Type Data sets (series) to display in the chart, chart type and color of plotted series, marker lines to display (optional)
- Axis Vertical and horizontal axes to display, axis labels, axis scales, and axis number formats
- Make appropriate settings.
- Optional step: To apply the settings to more than one chart, click Apply To. Then, specify whether to apply all chart preferences or just the current tab and whether to apply them to the current Microsoft Excel sheet, all sheets in the workbook, or all open and new workbooks and click OK. Otherwise, go to step 5.
- Click OK to apply the settings on all tabs to the active chart.

For a list of customizations you can perform on each tab, see "Specific Customization" Instructions" on page 117.

Specific Customization Instructions

This section explains how you can customize each part of the Crystal Ball charts to present the information most effectively. Topics for each Chart Preferences dialog tab are as follows.

General tab:

- "Adding and Formatting Chart Titles" on page 118
- "Changing the Chart Density" on page 120
- "Showing Grid Lines" on page 120
- "Showing the Chart Legend" on page 121

• "Setting Special Chart Effects" on page 121

Chart Type tab:

- "Setting the Chart Type" on page 118
- "Setting Chart Colors" on page 122
- "Showing the Mean and Other Marker Lines" on page 122

Axis tab:

• "Customizing Chart Axes and Axis Labels" on page 123

All tabs:

"Applying Settings to the Current Chart and Other Charts" on page 124

For basic chart customization instructions, see "Basic Customization Instructions" on page 117.

Adding and Formatting Chart Titles

- To add or change a chart title:
- 1 Display the General tab of the Chart Preferences dialog.

The Chart Title group is displayed at the top of the General page.

By default, Auto is checked and a default title is displayed.

- 2 To edit the title, uncheck Auto and type the new title in the text box.
- 3 When all the settings are complete, click OK.

The chart is displayed with the new settings.

Setting the Chart Type

Depending on the basic chart type (assumption, forecast, trend, overlay, or sensitivity), you can choose from among several chart display types, such as column, line, area, bar, or pie.

- To change the chart display type:
- Display the Chart Type tab of the Chart Preferences dialog. The Chart Type group is displayed in the middle of the Chart Type page.
 - If more than one series is displayed in the list box at the top of the tab, select one to work with. The settings on the tab apply to the selected series.
- To change the chart display type, open the Type drop-down list and select a display type. Depending on the basic chart and series types, you can choose from among these display types (not including scatter charts):

Table 6 Chart Types

Example	Туре	Description
Carrent Green Purific Represed (MM)	Column	Displays data as vertical columns that correspond to the group intervals (chart bins) of the data. The column chart is the default chart type for generated data in assumption, forecast, and overlay charts.
Consult Consult Printed Approved (MM) Salt two Printed Inference 1986 Salt two Printed Inference	Line	Shows data as an outline of the peaks and valleys.
Comment Great Payer of Approved (MM) Still the Present Inference	Area	Shows data as darkened peaks and valleys.
25.5% 25.5% 2.0%	Bar (directional)	Shows sensitivity data as horizontal bars to the right and left of the 0 line showing magnitude and direction of sensitivity
65.2% 25.5% 7.2% 2.0%	Bar (magnitude)	Shows sensitivity data as horizontal bars to the right of the 0 line showing magnitude of sensitivity but not direction
2.0% 7.2% -25.5%	Pie	Shows sensitivity data as a circle divided into proportional "slices" showing the magnitude of sensitivity

- 3 While you have a series selected, consider adjusting the chart color ("Setting Chart Colors" on page 122) and marker line settings ("Showing the Mean and Other Marker Lines" on page 122) too.
- When settings for the current series are complete, follow steps 2 through 3 to customize settings for any other series in the chart.


5 When all the settings are complete, click OK.


The chart is displayed with the new settings.

Changing the Chart Density

You can show more or less detail (such as data points and columns) in a chart by changing the number of bins used to group similar values. The level of detail or resolution is called the density of the chart. Higher densities more accurately reflect the actual distribution of data and lower densities highlight the overall shape and trend of data.

- To change the chart density:
- 1 Display the General tab of the Chart Preferences dialog.
- 2 Locate the Chart Bins group in the middle of the page.
- 3 Choose a density level from the Density drop-down list. For example, the following figure shows the lowest and highest density settings for a column chart:





Lowest density

Highest density

4 To show a space between each column, or "bin," check Show Column Gaps.

Even if you uncheck Show Column Gaps, gaps will always be displayed in a discrete distribution.

5 When all the settings are complete, click OK.

The chart is displayed with the new settings.

Showing Grid Lines

Grid lines are vertical or horizontal lines that can be displayed on charts to help you compare the charted data with the values along the horizontal and vertical axes of the chart.

For basic chart customization instructions, see "Basic Customization Instructions" on page 117.

- To hide or show grid lines:
- 1 Display the General tab of the Chart Preferences dialog.
- 2 Locate the Options group in the middle of the page.

- Choose a setting from the Gridlines drop-down list to show only horizontal grid lines (Horizontal), only vertical grid lines (Vertical), both horizontal and vertical grid lines (Both), or choose None to hide both horizontal and vertical grid lines.
- When all the settings are complete, click OK.

The chart is displayed with the new settings.

You can press Ctrl+l to toggle the horizontal grid lines on and off.

Showing the Chart Legend

The legend shows the name and chart color for each series in the chart.

For basic chart customization instructions, see "Basic Customization Instructions" on page 117.

- To hide or show a chart legend:
- Display the General tab of the Chart Preferences dialog.
- 2 Locate the Options group in the middle of the page.
- 3 Choose a setting from the Legend drop-down list to show the legend at the right side of the chart (Right), the left side of the chart (Left), or the bottom of the chart (Bottom). To hide the legend, choose None.
- 4 When all the settings are complete, click **OK**.

The chart is displayed with the new settings.

You can press Ctrl+n to toggle the legend on and off.

Setting Special Chart Effects

You can use special effects to help present the data most effectively. Transparency ensures that all chart series and values are visible and 3D effects add graphic depth that can be so useful when many series are charted (for example, bars become blocks as shown in the chart density figure in "Changing the Chart Density" on page 120).

For basic chart customization instructions, see "Basic Customization Instructions" on page 117.

- To set special chart effects:
- Display the General tab of the Chart Preferences dialog.
- Locate the Effects group at the bottom of the page.
- You can check any or all of the available effects to see how they enhance the chart. If you check Transparency, you can also choose a percent. 0% is completely opaque and 100% is completely transparent.
- When all the settings are complete, click OK.

The chart is displayed with the new settings.

You can press Ctrl+w to toggle 3D display on and off.

Setting Chart Colors

This preference sets the color of the current chart series. This is the color that is displayed for the series in the chart legend, if visible.

For basic chart customization instructions, see "Basic Customization Instructions" on page 117.

- To change the chart colors:
- Display the Chart Type tab of the Chart Preferences dialog. The Chart group is displayed in the middle of the Chart Type page.
 - If more than one series is displayed in the list box at the top of the tab, select one to work with. The settings on the page apply to the selected series.
- 2 Open the Color drop-down list and choose a color.
- While you have a series selected, consider adjusting the chart type ("Setting the Chart Type" on page 118) and marker line settings ("Showing the Mean and Other Marker Lines" on page 122) too.
- 4 When settings for the current series are complete, follow steps 2 through 3 to customize settings for any other series in the chart.
- 5 When all the settings are complete, click OK.

The chart is displayed with the new settings.

Showing the Mean and Other Marker Lines

You can display mean, mode, median, standard deviation, certainty, and other marker lines on assumption, forecast, and overlay charts. These lines help you locate various values in the charted distribution.

Note: If you have activated the process capability features and have entered an LSL, USL, or Target value, you can include marker lines for them in forecast charts ("Viewing LSL, USL, and Target Marker Lines" on page 375).

Base Case is the value in an assumption, decision variable, or forecast cell prior to running the simulation. For forecasts, Certainty Range shows lines at the certainty range endpoints. Marker lines are shown with labels, such as **Mean** = \$125.

You can press Ctrl+m to cycle through the median, mean, and base case or mode, depending on chart type. Press Ctrl+p to cycle through every 10th percentile.

- To display marker lines:
- Display the Chart Type tab of the Chart Preferences dialog. The Marker Lines group is displayed on that tab.

If more than one series is displayed in the list box at the top of the tab, select one to work with. The settings on the page apply to the selected series.

- 2 To display a marker line, check the box for that item. If you choose Standard Deviation, Percentile, or Value, another dialog opens:
 - For **Standard Deviation**, enter the standard deviation(s) where you want a marker to be displayed. If you enter more than one, separate them with commas. Then, choose whether you want the marker(s) to be displayed below the mean (technically indicating negative standard deviations), above the mean, or both above and below.
 - For Percentile, select the group of percentiles where you want markers to be displayed or choose Custom and create the own group of percentile points separated by commas.
 - For Value, enter the x-axis value where the line should be displayed in the Value field and click Add. If you want, enter an optional label and check Show Value On Marker Line to display the value on the chart. You can click New to add another value.
- While you have a series selected, consider adjusting the chart type ("Setting the Chart Type" on page 118) and color ("Setting Chart Colors" on page 122) too.
- 4 When settings for the current series are complete, follow steps 2 and 3 to customize settings for any other series in the chart.
- When all the settings are complete, click OK.
 - The chart is displayed with the new settings.

If the marker lines fall outside the maximum or minimum value displayed on a chart, they will not be displayed on the chart. For example, this can happen with standard deviations of plus or minus 2 or 3 for uniform distributions.

Customizing Chart Axes and Axis Labels

You can customize the label, scale, and format of the main axis in Crystal Ball charts.

- To customize chart axes:
- Display the Axis tab of the Chart Preferences dialog.
- Consider the label to use for the axis.
 - By default, Auto is checked in the Axis Label group. A label is automatically assigned depending on the chart type. For example, the value (x) axis of a forecast chart uses the Units entry in the Define Forecast dialog as the label. To enter a custom axis label, uncheck Auto and type a more descriptive label in the edit field.
- Next, adjust the Scale settings. By default, Auto is displayed in the Type list and the chart is displayed with the most appropriate end points for that chart type. To use another scale, choose it from the Type drop-down list and enter the minimum (Min) and maximum (Max) values to use for the scale.
 - Most chart/axis combinations offer Fixed as an alternative. The value axis for assumption, forecast, and overlay charts also offers Standard Deviation and Percentile.
- The Format Number settings control the format of the axis label numbers. Choose appropriate settings for the chart.

- For Format settings, Cell Format uses the format of the underlying cell. Most choices are similar to those used in Microsoft Excel: General, Number, Currency, Scientific, Percentage, or Date.
- The Decimal settings control the number of decimal points.
- When checked, Thousand Separator inserts a thousands-separator symbol where appropriate (except when Scientific formatting is set). The thousands separator that is displayed is the one defined in Windows International or Regional Options settings.

The Format Number settings also control the format of assumption parameters in the Define Assumption dialog and assumption charts.

5 When all the settings are complete, click OK.

The chart is displayed with the new settings.

Applying Settings to the Current Chart and Other Charts

If you would like to apply the current settings to other charts in the model, you can choose what to apply and where to apply them. (These instructions apply anywhere the Apply To button is displayed.)

- To specify how chart settings should be applied:
- 1 Click the Apply To button.

The Apply To dialog opens.

- 2 Indicate the tab or tabs of settings to apply:
 - This Tab applies only settings on the current tab.
 - All Tabs applies all the current settings on the entire dialog.
- 3 Indicate where the settings should be applied.
 - This Sheet applies the settings to only the current sheet of the current workbook.
 - This Workbook applies the settings to all sheets of the current workbook.
 - All Open And New Workbooks applies the settings to all workbooks that are currently open and all workbooks to be created.

All Workbooks effectively changes the global Chart Preferences defaults to the settings on the current tab or all tabs, depending on the setting in the previous dialog group.

Apply To settings are both flexible and powerful. They can be used as focused or broad defaults.

Managing Existing Charts

Previous sections of this chapter have described how to create and customize new charts. The following sections describe how to open, copy, paste, print, close, and delete existing charts:

• "Opening Charts" on page 125

- "Copying and Pasting Charts to Other Applications" on page 125
- "Printing Charts" on page 126
- "Closing Charts" on page 127
- "Deleting Charts" on page 127

Opening Charts

After you create an assumption or forecast chart, it is saved with the workbook that contains it. Other charts are saved with the active workbook model. You can display charts again, with current data, anytime you rerun the model with associated workbooks open.

- To open a chart:
- 1 Open the model containing the chart and run a simulation or restore saved results ("Restoring Crystal Ball Simulation Results" on page 89).
- 2 Open the Analyze menu and select the type of chart to display: Assumption Charts, Forecast Charts. Overlay Charts, Trend Charts, or Sensitivity Charts.
- When the dialog for that chart opens, check the box in front of each chart to display.
- Click Open.

Crystal Ball and Microsoft Excel icons in the Windows task bar to activate charts that have disappeared behind the spreadsheet.

In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Forecast Charts.

To open a number of charts at one time, select the Crystal Ball data cells and choose Analyze, then Open Selected Cells. All charts for the selected cells open and are displayed in front of any other open charts. In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Open From Selection.

Copying and Pasting Charts to Other Applications

You can copy and paste assumption, forecast, overlay, trend, and sensitivity charts to other applications such as Microsoft Word, Powerpoint, and Microsoft Excel.

For instructions, see the following sections:

- "Copying Charts" on page 125
- "Pasting Charts from the Clipboard" on page 126

Copying Charts

- To copy charts for use in other applications:
- Select the chart to copy.

Open its View menu and choose the view to be copied.

If you choose a data view such as Percentiles, Statistics, or Goodness Of Fit, the data will be pasted into many applications as alphanumeric data, ready to edit, add, and so on. This is true for Microsoft Excel and Word, but not Powerpoint. Data is pasted into Powerpoint as a graphic.

Graphic views, such as Frequency, are pasted as bitmap images.

3 In the chart's menubar, choose **Edit**, then **Copy**.

The chart is copied to the Clipboard, ready to paste into another application.

Pasting Charts from the Clipboard

- To paste a chart into another application using its own Paste commands:
- 1 Copy the Crystal Ball chart as described in the previous section.
- 2 Open a document (spreadsheet, slide, and so on) in the application to receive the chart.
- 3 Within that application, choose Edit, then Paste or Edit, then Paste Special.

(To paste in Microsoft Excel 2007 or later, press Ctrl+v or click the Home tab and choose Clipboard, then Paste.)

As described previously, if you copied a data view such as Percentiles, Statistics, or Goodness of Fit, the data is pasted into many applications as editable numbers or text.

Graphic views, such as Frequency, are pasted as bitmap images.


Printing Charts

To print a chart, display it and choose Edit, then Print in the chart's menubar.

Before printing, can choose Edit, then Page Setup to format the chart on the page. Then, choose Edit, then Print Preview to view the chart as it will print on the selected paper size.

For example, Figure 35 shows the Print Preview dialog for a forecast chart from Toxic Waste Site.xls in Landscape orientation on Letter paper.

Figure 35 Print Preview Dialog for a Forecast Chart



Closing Charts

When you close a chart, you remove it from memory but do not delete it permanently.

- To close a chart:
- Open the Analyze menu and select the type of chart to close.

If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Analyze, then View Charts and then choose a command.

- When the dialog for that chart opens, check the box in front of each chart to close.
- Click Close.

The selected chart or charts are closed without prompting.

You can use Analyze, then Close All to close all chart windows from the current simulation and restored results. In Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Close All.

Deleting Charts

You do not need to open a chart to delete it, as long as the model or saved results file containing it is open.

- To delete a chart (except assumption and forecast charts):
- 1 Open the model containing the chart.
- Open the Analyze menu and select the type of chart to delete.

If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Analyze, then View Charts and then choose a chart type.

- 3 When the dialog for that chart opens, check the box in front of each chart to delete.
- 4 Click Delete.

The selected chart or charts are deleted without prompting. Assumption and forecast charts cannot be deleted in this way.

7

Analyzing Other Charts

In This Chapter

Overview	129
Understanding and Using Overlay Charts	130
Understanding and Using Trend Charts	136
Understanding and Using Sensitivity Charts	142
Understanding and Using Assumption Charts	151
Understanding and Using Scatter Charts	153

Overview

The previous chapter gave guidelines for analyzing and customizing simulation results with a focus on forecast charts.

This chapter extends the information on analyzing simulation results provided in Chapter 6. You will learn how to use the following additional charts to interpret and present data:

- Overlay charts to view forecasts superimposed over each other ("Understanding and Using Overlay Charts" on page 130)
- Trend charts to view the certainty ranges of all the forecasts on a single chart ("Understanding and Using Trend Charts" on page 136)
- Sensitivity charts to view the impact the assumptions have on forecast results ("Understanding and Using Sensitivity Charts" on page 142)
- Assumption charts to confirm random number sampling ("Understanding and Using Assumption Charts" on page 151)
- Scatter charts to compare paired data ("Understanding and Using Scatter Charts" on page 153)

If you have OptQuest, you can also display OptQuest charts with optimization results.

For information on customizing charts, managing chart windows, and printing charts, see Chapter 6.

Selecting Assumptions, Forecasts, and other Data Types

When defining Crystal Ball charts and performing other procedures, you must sometimes select assumptions, forecasts, and other types of Crystal Ball data or objects. The following instructions apply to several types of selection situations.

- To select a Crystal Ball data cell or other object:
- Perform an operation that displays a Choose... dialog. By default, these dialogs open in a hierarchical Tree view. If you prefer, click the List button to change the view from a tree to a list.
- Check the boxes in front of the assumptions, forecasts, decision variables, or other objects to include.
- When the selection is complete, click **OK**.

Understanding and Using Overlay Charts

After completing a simulation with multiple related forecasts, you can create an overlay chart to view the relative characteristics of those forecasts on one chart. You can compare differences or similarities that otherwise might not be apparent. There is no limit to the number of forecasts you can view at one time on an overlay chart.

For example, if a model has several forecasts based on slightly different interest rate assumptions, an overlay chart can be used to show how these slight differences manifest themselves in the variability of the forecasts. You can customize the overlay chart to accentuate these differences or similarities.

The following overlay chart (Figure 36) shows the relative reliabilities of three manufacturing materials.




Figure 36 An Overlay Chart with 3D Formatting

Note: The default overlay chart is two-dimensional, not rotated, and displays the legend below the chart. This and other illustrations may differ from the default view.

After the simulation stops, you can also use the overlay chart to fit probability distributions to the forecasts. This process is similar to the distribution fitting feature described in "Fitting a Distribution to a Forecast" on page 113, except that the fit is applied to forecast data, not historical data.

The following sections describe how to create and use overlay charts:

- "Creating Overlay Charts" on page 131
- "Customizing Overlay Charts" on page 134
- "Using Distribution Fitting with Overlay Charts" on page 134

Creating Overlay Charts

The example model used in this section, Reliability.xls, shows the reliability of a design component when it is manufactured from different materials and is subjected to varying stresses.

Each forecast chart displays the design component's reliability as a ratio of two distributions, strength and stress. A reliable component has values greater than 1, because the component's strength exceeds the stress subjected to it. An unreliable component has values less than 1, because its strength is less than the stress. For this model, the overlay chart can be used to compare component reliabilities.

- To create an overlay chart:
- Run a simulation in Crystal Ball.

To produce a meaningful overlay chart, the simulation should have more than one forecast. Reliability.xls is used in these instructions.

Choose Analyze, then Overlay Charts.



(If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Overlay Charts.)

The Overlay Charts dialog opens. If you have not yet created any overlay charts for open workbooks or restored results files, the dialog is blank.


To create a new overlay chart, click New.

The Choose Forecasts dialog opens.

By default, this dialog opens in a hierarchical Tree view. If you prefer, click the List button to change the view from a tree to a list.

- Check the boxes in front of the forecasts to include in the overlay chart.
- Click OK to create a new overlay chart with the selected forecasts as shown in Figure 37.

Figure 37 Overlay Chart for Selected Forecasts




The overlay chart is displayed with the frequency distributions for the selected forecasts superimposed over each other. By default, they are displayed as columns. You can change them to lines or areas in two or three dimensions and can change the number of data points or intervals (bins) with the Chart Preferences settings ("Setting Chart Preferences" on page 115). A default name is assigned; you can also change it with the Chart Preferences settings.

- Follow the steps in "Customizing Overlay Charts" on page 134 and "Setting Chart Preferences" on page 115 to change a variety of chart features and highlight those of greatest interest.
 - You can also use "shortcut keys," or keyboard equivalents for commands to quickly change the chart preferences. For a list of these, see Table 5 on page 116.
- In this example, press Ctrl+d until the reverse cumulative chart view is displayed. Then, press Ctrl+t until the outline chart type is displayed.

As shown in Figure 38, the chart view changes to show all three distributions completely. This chart in outline view most clearly suggests that Material 3 has superior reliability and is dominant since a greater proportion of its distribution is to the right of 1.00 and its values for all probability levels are higher than the others.

Figure 38 Overlay Chart with Three Distributions



Now, experiment with chart shortcut keys.

Press Ctrl+d again until the frequency distribution is displayed. Press Ctrl+t to display the column chart. Try pressing Ctrl+b to change the number of frequency bins (columns, in this view). Then, press Ctrl+w to make the chart three-dimensional.

Overlay Chart 1 Edit View Overlay Preferences Help ✓ Enable Rotation 5.000 Trials Frequency View Overlay Chart 1 400 0.08 350 0.07 300 0.06 250 0.05 Material 1 Reliability 200 Material 2 Reliability 0.04 Material 3 Reliability 150 0.03 100 0.02 50 0.01 0.00 1.10 1.20 1.30

Figure 39 Overlay Chart, 3D View

If you want, you can drag either side of the chart to give it a taller, narrower look, as shown in Figure 39, or a stretched look as in Figure 40.

In 3D view, the Enable Rotation checkbox is displayed at the top of the chart, accessible with the Tab key. When it is checked, you can click inside the chart and drag to rotate it. This can enhance the data display for both analysis and presentation. Figure 40 shows a rotated overlay chart, stretched to emphasize x-axis differences.

Overlay Chart 1 ▼ Enable Rotation 5,000 Trials Overlay Chart 1 0.06 240 200 0.03 Material 2 Reliability 160 0.02 Material 3 Reliability 120 80 40 1.05 1.10 1.15 1.20 1.25

Figure 40 A Rotated and Stretched Overlay Chart

Note: Rotation settings are for the current session only and are not saved with the chart.

For more customization suggestions, see the next section.

Customizing Overlay Charts

You can customize overlay charts in a variety of ways:

- Use the View menu in the overlay chart window to switch among several graphic and numeric views.
- Use the Overlay menu to add additional forecasts to the chart or remove them all and toggle between the default view and Goodness Of Fit view.
- Use Preferences, then Overlay to choose a view, determine when the overlay chart window should be displayed, and specify whether to fit distributions to all forecasts (described in "Using Distribution Fitting with Overlay Charts" on page 134).
- Use Preferences, then Chart Preferences to further customize the chart's appearance as described in "Setting Chart Preferences" on page 115.

Customizing overlay charts helps you compare forecasts by viewing their differences in several ways. For example, the area and column chart types might obscure parts of some distributions behind other distributions, but the outline and line chart types show virtually all of each distribution. For best display of many types of data, you can choose 3D view and then rotate the chart as shown in Figure 40.

Using Distribution Fitting with Overlay Charts

Similar to the distribution fitting with historical data described in "Fitting Distributions to Data" on page 46, you can fit standard probability distributions to forecasts in an overlay chart, in addition to comparing forecasts to each other.

You can fit distributions to forecasts in overlay charts two ways:

- Choose Overlay, then Fit Probability Distributions in the overlay chart menubar to do a quick fit with the default or currently selected distributions and ranking method. You can also use this command to switch off distribution fitting that is set with either the Overlay menu or Preferences menu.
- Choose Preferences, then Overlay, then Overlay Window in the overlay chart menubar to specify particular distributions and to choose one of three fit ranking methods. Then, you can also change the fit options or use Apply To to set these preferences for other overlay charts.
- To fit a probability distribution to all forecasts in an overlay chart using the Preferences, then Overlay command:
- 1 Follow the steps for forecast charts given in "Fitting a Distribution to a Forecast" on page 113. Wherever the instructions say Forecast, as in Preferences, then Forecast, substitute Overlay.
- Click **OK** to perform the fit.

Crystal Ball fits the distributions, and then displays a probability distribution for each forecast as shown in Figure 41. As the legend shows, the forecast in the middle fits best to a gamma distribution, while the other two are beta fits. The Series tab of the Chart Preferences dialog was used to change the colors of the lines of best fit for greater contrast in the figure.




Figure 41 Overlay Chart with Forecasts and Lines of Best Fit

Note: This overlay chart is shown in 3D, rotated view with legend to the right.

You can copy overlay charts and paste them into other applications. For more information, see "Copying and Pasting Charts to Other Applications" on page 125.

Understanding and Using Trend Charts

After completing a simulation with multiple related forecasts, you can create a trend chart to view the certainty ranges of all the forecasts on a single chart. A trend chart summarizes and displays information from multiple forecasts, making it easy to discover and analyze trends that might exist between related forecasts.




Figure 42 Upward Trending Sales Figures, By Quarter

Trend charts display certainty ranges for multiple forecasts in a series of colored bands. Each band represents the certainty ranges into which the actual values of the forecasts fall. For example, the band which represents the 90% certainty range shows the range of values into which a forecast has a 90% chance of falling. By default, the bands are centered around the median of each forecast.

The trend chart in Figure 42 displays certainty ranges on a quarterly basis over a three year period. Because the model contains quarterly forecast formulas dependent on the previous quarter's results, the bands widen in the future. This occurs because the standard deviation of the forecasts increases or widens for each quarter. Trend charts like this one demonstrate the compounding of uncertainty that occurs as predictions are made farther and farther into the future.

You can customize trend charts to display the probability that given forecasts will fall in a particular part of a value range. For example, if a model contains forecasts related through time, you can use trend charts to view the certainty ranges for each forecast side by side. You can also compare at a glance the certainty ranges for an early time period and a later time period.

Trend charts are only meaningful when you have multiple forecasts that are related to each other.

The following sections describe how to create and work with trend charts:

- "Creating Trend Charts" on page 137
- "Customizing Trend Charts" on page 137

•

Creating Trend Charts

The following are basic steps for creating trend charts. The steps are illustrating using an example spreadsheet, Sales Projection.xls, shipped with Crystal Ball.

- To create a trend chart:
- Run a simulation in Crystal Ball.

To produce a meaningful trend chart, the simulation should have more than one forecast.

When the simulation stops, choose Analyze, then Trend Charts.



(If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Trend Charts.)

The Trend Charts dialog opens. If no trend charts have been created for the open workbooks or restored results files, the dialog is blank.

3 To create a new trend chart, click New.

The Choose Forecasts dialog opens.

By default, this dialog opens in a hierarchical Tree view. If you prefer, click the List button to change to List view.

Select two or more forecasts to include in the trend chart.

To select a forecast, check the box in front of its name. For this example, select all forecasts in the Sales Projection worksheet.

To select all forecasts in this example, you can click the Select All button or check the box

in front of Sales Projection.xls or the Model tab name. lacktree lacktree



Click OK.

The trend chart is displayed as shown in Figure 42.

As with overlay charts, you can change the scale and proportions of the chart by dragging its edges. For other customization possibilities, see the next section.

Customizing Trend Charts

You can customize trend charts in several ways. Most of the following customizations are available through trend chart menus or the Trend Preferences dialog:

- "Changing Trend Chart Views" on page 138
- "Specifying When Trend Charts Are Displayed" on page 139
- "Adding, Removing, and Ordering Forecasts" on page 139

The following sections describe how to make settings using the Trend Chart Preferences dialog:

- "Changing the Overall Appearance of Trend Charts" on page 140
- "Setting Certainty Band Type and Colors" on page 140
- "Setting Certainty Bands" on page 141
- "Changing Value Axis Preferences" on page 141

For some of the options, you can use shortcut keys to bypass the Trend Preferences dialog. These shortcut keys are listed in Table 5 on page 116 and Appendix F, "Accessibility."

Changing Trend Chart Views

Use the trend chart View menu to change the placement of the certainty bands within the trend chart. The default setting centers the bands around the median of each forecast. You can change the location of the bands so that they are anchored at either the high end or the low end of the projected forecast ranges.

Smaller bands are always displayed on top of larger bands. This obscures the larger bands. Don't confuse the actual width of a band with the portion that is visible. You can display the true size of a band using the certainty band preferences, described in "Setting Certainty Bands" on page 141. Optionally, you can use these preferences to display the bands one at a time on the trend chart.

- To change the placement of the certainty bands:
- 1 In the trend chart, open the View menu.
- Choose one of the views listed in Table 7, following.

Table 7 Trend Chart Views

View	Effect	
Centered On Median		The default; displays forecasts centered around the median of each forecast value.
Cumulative		Displays the forecasts anchored at the lower end of the forecast range; shows the certainty that the forecast values will be at or below a given value (cumulative probability).

View	Effect	
Reverse Cumulative		Displays the forecasts anchored at the higher end of the forecast range; shows the certainty that the forecast values will be at or above a given value (reverse cumulative probability).

Alternately, you can choose Preferences, then Trend and choose a view in the View list of the Trend Preferences dialog.

You can use a Crystal Ball shortcut key to bypass the View menu. Each time you press Ctrl+d, the next trend chart view opens.

Specifying When Trend Charts Are Displayed

You can set preferences that determine whether the trend chart opens automatically, and if so, whether it is displayed while the simulation is running or after it stops.

- To set these display preferences:
- 1 Choose Preferences, then Trend.

The Trend Preferences dialog opens.

2 To change the trend chart view, use the drop-down View list.

The views are discussed in the previous section, "Changing Trend Chart Views" on page 138.

- 3 Use settings in the Windows group to determine whether the chart opens automatically.
 - If Show Automatically is checked, you can choose whether to display the chart while the simulation is running or after it stops.
- 4 You can click Defaults at any time to restore original default settings for the Trend Preferences dialog.
- When all settings are complete, click OK to activate them and close the dialog.

Adding, Removing, and Ordering Forecasts

- To add and remove forecasts from a trend chart:
- 1 In the trend chart menubar, choose Trend, then Choose Forecasts.
 - The Choose Forecasts dialog opens.
- 2 Check and uncheck forecasts to add and remove them from the chart.
 - To clear all forecast selections, choose Trend, then Remove All in step 1.
- 3 Click OK to accept the settings and modify the chart.

After you select forecasts and create a trend chart, you can change the order of forecast display.

- To change the order of forecasts:
- 1 In the trend chart menubar, choose Preferences, then Chart, then Chart Type.
 - The Chart Type dialog opens.
 - All charted forecasts are displayed in the Series list at the top of the dialog in their display order.
- Select a forecast in the list and use the up and down arrow keys to move the forecast up or down the list.
- 3 Optional step: You can choose Defaults at any time to restore all settings to their original default values.
- Optional step: To apply the settings to more than one chart, click Apply To. Then, specify how they should be applied (see "Applying Settings to the Current Chart and Other Charts" on page 124 for details) and click OK. Otherwise, go to step 5.
- 5 When all forecasts are positioned in order, click OK to accept the current Chart Preferences settings and modify the chart.

Changing the Overall Appearance of Trend Charts

When you first choose Preferences, then Chart in the trend chart menubar, the General tab of the Chart Preferences dialog opens.

Except for the disabled Chart Bins preferences, the General tab settings are the same as those for forecast and other charts.

You can set the following features, described in the sections in parentheses:

- Chart title ("Adding and Formatting Chart Titles" on page 118)
- Gridlines ("Showing Grid Lines" on page 120)
- Legend ("Showing the Chart Legend" on page 121)
- Chart effects ("Setting Special Chart Effects" on page 121)

For descriptions of each setting while you are viewing the dialog, click Help.

Setting Certainty Band Type and Colors

By default, the trend chart certainty bands are displayed as solid-colored areas. You can change the bands from solid areas to uncolored areas bordered by lines. You can also change the color of each certainty band.

- To change the chart type or color settings:
- 1 Choose Preferences, then Chart, then Chart Type.
 - The Chart Type tab of the Chart Preferences dialog opens.
- 2 To change all certainty bands from areas to lines, choose Line in the Chart Type drop-down list.

- To change the color of a certainty band:
 - a. Select the certainty band to change from the list below the Certainty Bands button.
 - b. Select a color from the Band Color list.
- To select a different set of certainty levels or define the own, click the Certainty Bands button and follow the steps in the next section, "Setting Certainty Bands" on page 141
- 5 Optional: You can choose Defaults at any time to restore all settings to their original default values.
- 6 Optional: To apply the settings to more than one chart, click Apply To. Then, specify how they should be applied (see "Applying Settings to the Current Chart and Other Charts" on page 124 for details) and click OK. Otherwise, go to step 7.
- 7 When settings are complete, click OK to accept current settings on all tabs of the Chart Preferences dialog.

You can use the Chart Series list at the top of the Chart Type tab to change the order of forecasts in the forecast axis. For instructions, see "Adding, Removing, and Ordering Forecasts" on page 139.

Setting Certainty Bands

You can easily choose a set of certainty bands (levels) to display in the trend chart or create the own custom set.

- To change or define a set of certainty bands:
- 1 In the trend chart menubar, choose Preferences, then Chart.
- 2 Click the Chart Type tab.
- 3 On the Chart Type tab, click the Certainty Bands button.
- 4 The Percentiles dialog opens.
- Choose a set of certainty bands to display on the trend chart.
- To create the own set, choose Custom and enter a series of certainty bands, separated by commas.
- Click OK.

The trend chart opens with the selected certainty bands. A legend is displayed on the right side of the trend chart indicating which levels are represented by which bands.

If the legend does not include all bands, drag the top or the bottom of the trend chart to increase its height until all bands are displayed.

Changing Value Axis Preferences

You can change several trend chart axis preferences, as described below. You can add a name for the value axis, set number formats, and set value rounding. You can also change Scale settings from Auto to Fixed and specify a range minimum and maximum value. You can use these scale settings to display the probability that given forecasts will fall in a particular part of a value range.

- To change value axis settings:
- 1 In the trend chart menubar, choose Preferences, then Chart, then Axis.
 - The Axis tab of the Chart Preferences dialog opens.
- 2 By default, no name displays for the value axis. To add one, type it in the Axis Label field.
- 3 By default, Scale is set to Auto and displays all selected bands completely. To limit the display to a subset of values, set Scale to Fixed and enter a minimum and maximum value.
 - By changing the minimum or maximum endpoint values, you can zoom in or out on selected ranges of the trend chart.
- 4 The Format settings are similar to those for forecast and other charts. For more information, see "Customizing Chart Axes and Axis Labels" on page 123.
 - The number format for the axis values is taken from the first forecast that is displayed on the trend chart.
- 5 Optional: You can choose Defaults at any time to restore all settings to their original default values.
- Optional: To apply the settings to more than one chart, click Apply To. Then, specify how they should be applied (see "Applying Settings to the Current Chart and Other Charts" on page 124 for details) and click OK. Otherwise, go to step 7.
- 7 When settings are complete, click OK.
 - The trend chart opens with the value axis changed.
 - You can copy trend charts and paste them into other applications. For more information, see "Copying and Pasting Charts to Other Applications" on page 125.

Understanding and Using Sensitivity Charts

As you become more proficient at building spreadsheet models, you will want to know how much a given assumption affects the result. In other words, you want to determine the sensitivity of the forecast to each assumption. The overall sensitivity of a forecast to an assumption is a combination of two factors:

- The model sensitivity of the forecast to the assumption
- The assumption's uncertainty

To determine the model sensitivity, you would need to algebraically analyze the various relationships between the forecast cell and the assumption cells. Calculating these complex relationships could be a difficult and time-consuming task without a program like Crystal Ball.

The following sections describe how to create and use sensitivity charts:

- "About Sensitivity Charts" on page 143
- "Creating Sensitivity Charts" on page 143
- "How Crystal Ball Calculates Sensitivity" on page 145
- "Limitations of Sensitivity Charts" on page 146

"Customizing Sensitivity Charts" on page 147

About Sensitivity Charts

Sensitivity charts show the influence of each assumption cell on a particular forecast cell. During a simulation, Crystal Ball ranks the assumptions according to their importance to each forecast cell. The sensitivity chart displays these rankings as a bar chart, indicating which assumptions are the most important or least important in the model (Figure 43). You can output print the sensitivity chart on the report or copy it to the clipboard.

Sensitivity: Risk Assessment Edit View Sensitivity Preferences Help 5 000 Trials Contribution to Variance View Sensitivity: Risk Assessment 0.0% 21.0% 42.0% 63.0% Volume of Water per Day CPF 27 4% Body Weight Concentration of Contaminan...

Figure 43 Assumptions and Their Effects on Toxicity Risk

Sensitivity charts provide these key benefits:

- You can find out which assumptions are influencing the forecasts the most, reducing the amount of time needed to refine estimates.
- You can find out which assumptions are influencing the forecasts the least, so that they can be ignored or discarded altogether.
- As a result, you can construct more realistic spreadsheet models and greatly increase the accuracy of the results because you know how the assumptions affect the model.

Creating Sensitivity Charts

The example model used in this section, Toxic Waste Site.xls, illustrates the risk assessment of a toxic waste site. The sensitivity chart created in the steps below displays, in descending order, the assumptions in this model. The assumption with the highest level of sensitivity can be considered as the most influential assumption in the model.

- To create a sensitivity chart:
- Close any spreadsheets that are currently open.
- Open the spreadsheet to analyze.

Toxic Waste Site.xls is used in these instructions. You can choose Help, then Crystal Balll, then Examples Guide to choose the model from a list.

(In Microsoft Excel 2007 or later, choose Resources, then Examples Guide in the Help group.)

3 Choose Run, then Run Preferences, then Options.



4 Make sure Store Assumption Values For Sensitivity Analysis is checked and click OK.

Sensitivity analysis requires that all of the random numbers from the simulation be kept for comparison with forecast values.

- 5 Run a simulation until it stops.
- 6 Choose Analyze, then Sensitivity Charts.

Charts, then Sensitivity Charts.)



The Sensitivity Charts dialog opens. If you have not yet created any sensitivity charts for the active spreadsheet or restored results file, the dialog is blank.

7 Click the New button.

The Choose Forecast dialog opens.

By default, this dialog opens in a hierarchical Tree view. If you prefer, click the List box to change the view from a tree to a list. For more information on the Choose Forecast dialog, click the Help button.


8 Check the box in front of the forecast name to include it in the sensitivity chart.

Toxic Waste.xls has only one forecast, Risk Assessment.

9 Click **OK** to create a new sensitivity chart as shown in Figure 44.

The illustrated chart has a transparency effect applied using the chart preferences to make sensitivity values easier to read. For instructions, see "Setting Special Chart Effects" on page 121.

Figure 44 Sensitivity Chart for the Selected Forecast



The assumptions are listed beside the bar chart, starting with the assumption with the highest sensitivity. If necessary, use the scroll bar to view the entire bar chart. You can drag the edges of the chart to resize it — make it narrower, wider, taller, or shorter. This often changes the tick labels along the top of the chart.

Note: If you try to create a sensitivity chart but Store Assumption Values For Sensitivity Analysis was not selected in the Run Preferences dialog, select it, and then reset the simulation and run the simulation again.

In this example, Figure 44, there are four assumptions listed in the sensitivity chart. The first assumption, Volume Of Water Per Day, accounts for approximately 65% of the variance in forecast values and can be considered the most important assumption in the model. A researcher running this model would want to investigate this assumption further in the hopes of reducing its uncertainty and, therefore, its effect on the target forecast. The last assumption, Concentration Of Contaminant In Water, contributes the least to forecast variance (about 2%). In fact, this assumption has such a small effect, it could be ignored or altogether eliminated by clearing it from the spreadsheet.

Sensitivity charts like this one illustrate that one or two assumptions typically have a dominant effect on the uncertainty of a forecast.

How Crystal Ball Calculates Sensitivity

Crystal Ball calculates sensitivity by computing rank correlation coefficients between every assumption and every forecast while the simulation is running. Correlation coefficients provide a meaningful measure of the degree to which assumptions and forecasts change together. If an assumption and a forecast have a high correlation coefficient, it means that the assumption has

a significant impact on the forecast (both through its uncertainty and its model sensitivity). Positive coefficients indicate that an increase in the assumption is associated with an increase in the forecast. Negative coefficients imply the opposite situation. The larger the absolute value of the correlation coefficient, the stronger the relationship.

To help interpret the rank correlations, Crystal Ball provides a default chart view called the Contribution To Variance view. This view makes it easier to answer questions such as "What percentage of the variance or uncertainty in the target forecast is due to assumption X?"

It is important to note that the Contribution To Variance method is only an approximation and is not precisely a variance decomposition. Crystal Ball calculates Contribution To Variance by squaring the rank correlation coefficients and normalizing them to 100%.

Both the alternate Rank Correlation View and the Contribution To Variance view display the direction of each assumption's relationship to the target forecast. Assumptions with a positive relationship have bars on the right side of the zero line. Assumptions with a negative relationship have bars on the left side of the zero line. To show just the absolute magnitude of the relationship, you can change the Chart Type preference setting described in Table 8 to Bar (Magnitude).


Note: To ensure appropriate accuracy in Contribution To Variance view, consider running at least 10,000 trials.

Limitations of Sensitivity Charts


Sensitivity charts have several limitations:

- Correlated assumptions. Sensitivity calculations might be inaccurate for correlated assumptions. For example, if an important assumption were highly correlated with an unimportant one, the unimportant assumption would likely have a high sensitivity with respect to the target forecast. Assumptions that are correlated are flagged as such on the sensitivity chart. In some circumstances, turning off correlations in the Run Preferences dialog might help you to gain more accurate sensitivity information.
- Non-monotonic relationships. Sensitivity calculations might be inaccurate for assumptions whose relationships with the target forecast are not monotonic. A monotonic relationship means that an increase in the assumption tends to be accompanied by a strict increase in the forecast; or an increase in the assumption tends to be accompanied by a strict decrease in the forecast.

For example, the relationship y = Log(x) is monotonic:



While y = Sin(x) is not:



The Tornado Chart tool can help you discover if any of the assumptions have nonmonotonic relationships with the target forecast. For more information, see "Tornado Chart Tool" on page 187.

Discrete distributions. Sensitivity calculations might be inaccurate for assumptions or forecasts that have a small set of discrete values. Because the sensitivity calculation relies on rank correlation, a slight loss of information occurs when the assumption or forecast values are replaced by ranks. This loss of information is generally offset by the advantage of being able to measure sensitivity between dissimilar types of distributions. However, when a large percentage of assumption or forecast values are similar or identical, this loss of information grows and can significantly distort the calculation of correlations.

Be aware of this problem, for example, for:

- assumptions, when using distributions such as Binomial with a small Trials parameter (for example, < 10).
- forecasts, when formulas in the spreadsheet result in identical values (for example, ifthen logic, INT functions, and so on).

Customizing Sensitivity Charts

You can customize sensitivity charts by changing the assumptions, forecasts, and sensitivity preferences. As you become more familiar with the sensitivity chart, practice choosing the options that help you get the answers you seek and that are appropriate for the data.

The following sections describe how you can customize sensitivity charts:

- "Adding and Removing Assumptions" on page 147
- "Changing the Target Forecast" on page 148
- "Setting Sensitivity Preferences" on page 148
- "Setting Sensitivity Chart Preferences" on page 149

Adding and Removing Assumptions

By default, the sensitivity chart includes all of the assumptions from the simulation. To remove or change the assumptions included, follow the instructions in this section.

The total number of assumptions included in the chart affects the calculation of the Contribution to Variance percentages.

- To remove or change assumptions included in the sensitivity chart:
- In the Sensitivity Chart window, choose Sensitivity, then Choose Assumptions.

The Choose Assumptions dialog opens.

- 2 Check the assumptions to add to the sensitivity chart and uncheck those to remove from the chart.
- 3 Click OK to activate the selected assumptions.

Changing the Target Forecast

- To change the target forecast in the sensitivity analysis:
- 1 In the Sensitivity Chart window, choose Sensitivity, then Choose Target Forecast.

The Choose Forecasts dialog opens.

- 2 Check a new target forecast.
- 3 Click OK to activate the selected forecast.

Setting Sensitivity Preferences

You can set a number of preferences that determine:

- The sensitivity view that is displayed
- Whether the sensitivity chart opens automatically and whether it is displayed while the simulation is running or after it stops
- How many assumptions are shown in the chart, starting with the most sensitive
- Whether sensitivities are limited to a certain sensitivity value or higher
- To set sensitivity preferences:
- 1 Choose Preferences, then Sensitivity.

The Sensitivity Preferences dialog opens.

- 2 To change how the sensitivities are presented, use the View drop-down list:
 - Contribution To Variance shows sensitivities as values that range from 0% to 100% and indicate relative importance by showing the percentage of the forecast variance contributed by each assumption.
 - Rank Correlation shows sensitivities as rank correlations that range from -1 to +1 and indicate both magnitude and direction of the correlation of each assumption with the forecast.
 - Sensitivity Data shows a table of contributions to variance (%) and rank correlations for each assumption.
- 3 Use settings in the Windows group to determine whether the chart opens automatically.
 - If Show Automatically is checked, you can choose whether to display the chart while the simulation is running or after it stops.
- 4 To limit sensitivities by rank or value, click the Criteria tab to display it.

If you have a model with a lot of assumptions, you can check either or both boxes to limit the number of assumptions shown in the chart to a fixed number or to assumptions above a certain sensitivity value. If you check both, the more restrictive of the two criteria is used.

- You can click Defaults at any time to restore original default settings for the Sensitivity Preferences dialog.
- When all settings are complete, click OK to activate them and close the dialog.

You can copy sensitivity charts and paste them into other applications. For more information, see "Copying and Pasting Charts to Other Applications" on page 125.

Setting Sensitivity Chart Preferences

Settings in the Sensitivity Preferences dialog determine how sensitivity is viewed and set the criteria for showing assumptions. Additional settings in the Chart Preferences dialog determine the appearance of the chart.

- To set sensitivity chart preferences:
- 1 In the sensitivity chart window, choose Preferences, then Chart.

The chart preferences dialog opens.

- On the General tab, you can set the following features, described in the sections in parentheses:
 - Chart title ("Adding and Formatting Chart Titles" on page 118)
 - Gridlines ("Showing Grid Lines" on page 120)
 - Legend ("Showing the Chart Legend" on page 121)
 - Chart effects ("Setting Special Chart Effects" on page 121)

Except for the disabled Chart Bins preferences, the General tab settings are the same as those for forecast charts.

For descriptions of each setting while you are viewing the dialog, click Help.

- 3 Click the Chart Type tab for more settings.
- You can use the Chart Type drop-down list to choose one of these chart types:

Table 8 Sensitivity Chart Types


Chart Type	Description	Example
Bar (directional)	The default; horizontal bars to the right and left of the 0 line showing magnitude and direction of sensitivity	25.5% -7.2% 2.0%

Chart Type	Description	Example
Bar (magnitude)	Horizontal bars to the right of the 0 line showing magnitude of sensitivity but not direction	65.2% 25.5% 7.2% 2.0%
Pie	A circle divided into proportional wedges showing the magnitude of sensitivity (only available for Contribution To Variance view)	2.0% 7.2% -25.5%

- For bar charts, choose whether to use a different color for each assumption (the default), or whether to use the same color for all assumptions.
 - If you uncheck Show Multiple Colors, you can choose a specific color to use for all assumptions.
- Choose whether to show value labels on the chart (the default), or uncheck Show Values On Chart to show only graphics but no values.
- Optional step: You can choose Defaults at any time to restore all settings to their original default values.
- Optional step: To apply the settings to more than one chart, click Apply To. Then, specify how they should be applied (see "Applying Settings to the Current Chart and Other Charts" on page 124 for details) and click OK. Otherwise, go to step 9.
- Click OK to apply the settings on all tabs to the active chart.

You can apply different combinations of settings for special effects. For example, Figure 45 shows a sensitivity pie chart with 3D and Transparency chart effects. Assumptions have similar values and ranks to the directional bar chart shown in Figure 44.

Figure 45 Transparent, Three-Dimensional Sensitivity Pie Chart




Understanding and Using Assumption Charts

Assumption charts are created automatically when you run a simulation. They cannot be deleted, only opened or closed.

During or after a simulation runs, you can view assumption charts to show trial values (that is, random numbers) for the current simulation layered over the assumption's ideal probability distribution (Figure 46).

Figure 46 Assumption Chart



Assumption charts are useful for checking and comparing Run Preferences settings. For example, you can look at charts for the same assumption before and after increasing the number of trials and switching between Monte Carlo and Latin Hypercube sampling. If certain settings yield a noticeably better match between the random numbers and the ideal distribution, you might want to use those settings even if the simulation runs longer.

- To open an assumption chart:
- 1 Run a simulation.
- 2 Choose Analyze, then Assumption Charts.

(If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Assumption Charts.)

3 When the dialog opens, choose the assumptions to view and click OK.

For customization instructions, see "Customizing Assumption Charts" on page 152.

Customizing Assumption Charts

Because assumption charts look so similar to forecast charts, many of their menu commands and settings are the same.

The following sections describe how you can customize assumption charts:

- "Setting Assumption Chart Views" on page 152
- "Setting Assumption Preferences" on page 152
- "Setting Assumption Chart Preferences" on page 152

Setting Assumption Chart Views

You can use the View menu to choose five views: Probability, Cumulative Probability, Reverse Cumulative Probability, Statistics, and Percentiles. For a description of these views and how to select them, see "Changing the Distribution View and Interpreting Statistics" on page 104.

Setting Assumption Preferences

Assumption preferences, set with Preferences, then Assumptions, are similar to those for forecasts, described in "Setting Forecast Preferences" on page 111. By default, assumption charts are not shown when you run a simulation. You can change the Show Automatically setting to automatically show assumption charts while a simulation is running or when it stops.

While the Forecast Preferences dialog has a button for fitting distributions to forecasts, distribution fitting isn't available in the Assumption Preferences dialog. Instead, there is a Run Preferences button so you can easily change the Store Assumption Values... setting on the Options tab of the Run Preferences dialog. For more information on the Assumption Preferences dialog, click the Help button in the dialog.

Setting Assumption Chart Preferences

Assumption chart preferences are virtually identical to forecast chart preferences. To review or change them, choose Preferences, then Chart and follow the instructions in "Setting Chart Preferences" on page 115.

Understanding and Using Scatter Charts

Scatter charts show correlations, dependencies, and other relationships between pairs of forecasts and assumptions plotted against each other.

In its basic form, a scatter chart contains one or more plots of a target variable mapped against a set of secondary variables. Each plot is displayed as a cloud of points or symbols aligned in a grid within the scatter chart window. Figure 47 shows a set of all model assumptions plotted against a target forecast. In this case, the Material 3 Reliability forecast is the target.




Figure 47 Scatter chart, Scatter view, with Optional Lines and Correlations

In Figure 47, the line shows where the pairwise points would be displayed if they were sorted in ascending order. The closer the points conform to the line, the closer the relationship among the plotted variables. Lines sloped from lower values to higher (the lower left to the upper right) show positive relationships. If the relationship is negative, the line slopes from higher to lower values (the upper left to the lower right).

Figure 47 shows optional correlations displayed for each plot. Coil Diameter has the highest correlation with Material 3 Reliability, while Material 1 Strength has the lowest correlation.

In another form of scatter chart, the Matrix view, each selected variable is plotted against every other selected variable to show the relationships among them. Figure 48 shows intercorrelations among three forecasts in Matrix view. Material 2 Reliability and Material 3 Reliability have the highest intercorrelation while Material 1 Reliability and Material 2 Reliability have the lowest.

Figure 48 Scatter Chart, Matrix View, with Optional Lines and Correlations

The axis labels are indicated by the text in the diagonal cells. The text is the x-axis label for all plots in the same column as the text. It is the y-axis label for all plots in the same row. For example, in Figure 48, the y-axis label of the highlighted plot is Material 2 Reliability and the x-axis label is Material 3 Reliability.

You can plot scatter charts directly through the Analyze menu, or you can create a sensitivity chart and choose Sensitivity, then Open Scatter Chart to create a chart showing an exploded view of the effect each assumption has on the target forecast. The result is similar in form to Figure 47.

The following sections describe how to create and work with scatter charts:

- "Creating Scatter Charts" on page 154
- "Customizing Scatter Charts" on page 156

Creating Scatter Charts

- ➤ To create a scatter chart:
- 1 Run a simulation in Crystal Ball.
- When the simulation stops, choose **Analyze**, then **Scatter Charts**.



(If you are using Crystal Ball in Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Scatter Charts.)

The Scatter Charts dialog opens. If you have not yet created any scatter charts for the open workbooks or restored results files, the dialog is blank.

3 Click the New button.

The Choose Data dialog opens.

By default, this dialog opens in a hierarchical Tree view. If you prefer, click the List box to change the view from a tree to a list. For more information on the Choose Data dialog, click the Help button.

Check the box in front of an assumption or forecast name to include it in the scatter chart.


You can include up to 25 variables in a scatter chart. A warning message is displayed if you select more. If you try to create a scatter chart including an assumption but Store Assumption Values For Sensitivity Analysis was not checked in the Run Preferences dialog, check it, and then reset the simulation and run the simulation again.

To create a scatter chart in Scatter view, set a single assumption or forecast as the target. You do not need to set a target to display the chart in Matrix view.

To set a target, check the box in front of the target assumption or forecast, click its name, and then click Set As Target. In this case, Material 3 Reliability is set as the target and all assumptions are selected as secondary variables.

Click OK to create the new scatter chart shown in Figure 49.

Only a portion of the chart is shown here. For a view of the complete chart, see Figure 47.



Scatter Chart for the Selected Target, Scatter View Figure 49

Here is some information about features shown in Figure 49:

- You can choose Preferences, then Chart Preferences to change the chart title.
- To change the number of trials displayed in the plots, choose Preferences, then Scatter, then Criteria.
- The Y-axis labels indicate the scatter chart target. Each X-axis label indicates the secondary variable plotted against the target.
- The Ordered Fit line shows where the pairwise points would be displayed if they were sorted in ascending order. You can choose Preferences, then Chart Preferences, then Chart Type

to change it to a Linear Regression line, which uses a least-squares technique to show the linear relationship of the points.

- Auto is the default color for all symbols. With color set to Auto, plots are colored according to the combination of variables included in them:
 - Assumption vs. assumption = green
 - Forecast vs. forecast = dark blue
 - Assumption vs. forecast = dark teal (blue-green)
- You can use the Plot Sizer to increase or decrease the size of all plots and the amount of
 detail shown within them. To focus on a single plot, drag the Plot Sizer pointer toward the
 right to enlarge the plot, and then use the scroll bars to center it.
- In Scatter view, plots move to fill available window space when they are resized. In Matrix view, plots keep the same NxN configuration. You can scroll to view any plots that are not currently displayed onscreen.
- Frozen forecasts and assumptions are not included in scatter charts.

Customizing Scatter Charts

You can customize scatter charts by changing the plotted variables, scatter preferences, and chart preferences. As you become more familiar with the scatter chart, practice choosing the options that help you get the answers you seek and that are appropriate for the data.

To customize scatter charts, use the menus in the chart window or click portions of the chart:

- Double-click within a plot to open the Chart Preferences dialog.
- Double-click an axis to open the Axis dialog.
- Double-click outside a plot or axis to open the Scatter Preferences dialog.

The following sections describe ways to customize scatter charts:

- "Adding and Removing Assumptions and Forecasts" on page 156
- "Setting Scatter Preferences" on page 157
- "Setting Scatter Chart Preferences" on page 158
- "Scatter Charts and Filtered Data" on page 159

Adding and Removing Assumptions and Forecasts

When you create a new scatter chart, all assumptions and forecasts in the model are available to use as targets, secondary variables, or matrix variables. Some variables may be closely related to the target or other elements in the matrix and some variables may be completely unrelated.

You can include up to 25 variables, but scatter charts are easier to interpret with smaller numbers of variables. It can be helpful to edit the chart, removing some secondary variables or matrix elements that show little relationship to each other or a target and adding others that might relate more closely.

In complex models with lots of assumptions and forecasts, you might find it helpful to begin by creating a sensitivity chart and then creating a scatter chart from the data included within it. For example, you could open a forecast chart and choose Forecast, then Open Sensitivity Chart to view a sensitivity chart. Then, within the sensitivity chart, you could choose Sensitivity, then Open Scatter Chart to create a scatter chart using that forecast as the target.

- However you create a scatter chart, you can follow these steps to remove or change the variables (forecasts and assumptions) included in it:
- 1 In the Scatter Chart window, choose Scatter, then Choose Data.
 - The Choose Data dialog opens.
- 2 Check the assumptions or forecasts to add to the scatter chart and uncheck those to remove from the chart.
- If you want, set a different target.
 - Click the variable name and then click Set As Target.
- 4 Click OK to display the edited chart.

If you set a new target or remove a target, the view changes when the chart opens again. If you set a new target, the scatter chart is displayed in Scatter view with selected variables plotted against the new target. If you remove an existing target and do not set a new one, the chart is displayed in Matrix view.

Setting Scatter Preferences

You can set a number of preferences that determine:

- The scatter view that is displayed
- Whether the scatter chart opens automatically and whether it is displayed while the simulation is running or after it stops
- The plot size
- The number of trials that are plotted relative to the total number of trials in each simulation
- To set scatter preferences:
- 1 Choose Preferences, then Scatter.

The Scatter Preferences dialog opens.

- To change how the scatter chart is presented, use the View drop-down list:
 - Scatter View (1xN) shows secondary variables plotted against a target.
 - Matrix View (NxN) shows all selected variables plotted against each other.
- Use settings in the Window group to determine whether the chart opens automatically.

If Show Automatically is checked, you can choose whether to display the chart while the simulation is running or after it stops.

To change the size and amount of detail shown in each plot, use the Plot Sizer.

Slide it left to decrease or right to increase plot size. To set the new size as a global default or apply it to other scatter charts, click Apply To and choose appropriate settings. Click Help for additional information.

- To determine the number of trials that are plotted relative to the total number of trials in each simulation, click the Criteria tab to display it.
 - By default, only a portion of trials are plotted for each simulation, based on the total number of trials. To display all trials or a specific number, choose Show The First _____ Trials and enter the number of trials to display. Or, you can choose Show The First _____% Of The Trials to show a certain percent of the trials, from 1% to 100%.
- You can click Defaults at any time to restore original default settings for the Scatter Preferences dialog or click Apply To to set new defaults.
- 7 When all settings are complete, click OK to activate them and close the dialog.

You can copy scatter charts and paste them into other applications. For more information, see "Copying and Pasting Charts to Other Applications" on page 125.

Setting Scatter Chart Preferences

Settings in the Scatter Preferences dialog determine which scatter chart view is displayed and set the criteria for plotting all or a portion of trials actually run. Additional settings in the Chart Preferences dialog determine the appearance of the chart.

- To set scatter chart preferences:
- In the scatter chart window, choose Preferences, then Chart.

The Chart Preferences dialog opens.

- 2 On the General tab, you can set the following features, described in the sections in parentheses:
 - Chart title ("Adding and Formatting Chart Titles" on page 118)
 - Gridlines ("Showing Grid Lines" on page 120)
 - Legend ("Showing the Chart Legend" on page 121)
 - Chart effects ("Setting Special Chart Effects" on page 121)

Except for the disabled Chart Bins and 3D Chart preferences, the General tab settings are the same as those for forecast charts.

For descriptions of each setting while you are viewing the dialog, click Help.

- 3 Click the Chart Type tab for more settings.
- 4 You can use the Chart Type drop-down list to choose whether to:
 - Draw points, and if selected, to choose a symbol, color, and size for them.
 - Draw a line, and if selected, to choose a line type, color and size. The Ordered Fit line
 type shows where pairwise points would be displayed if they were sorted in ascending
 order. The Linear Regression line type uses a least-squares technique to show the linear
 relationship of the points.

- Display correlation coefficients for each plot. These are computed using the Spearman rank correlation method.
- Display points that have been filtered out (see the next section, "Scatter Charts and Filtered Data" on page 159).
- You can use the Axis tab to choose a number format for the chart axes and to indicate whether to round axis values, as described in "Customizing Chart Axes and Axis Labels" on page 123.
- Optional step: You can choose Defaults at any time to restore all settings to their original default values.
- 7 Optional step: To apply the settings to more than one chart, click Apply To. Then, specify how they should be applied (see "Applying Settings to the Current Chart and Other Charts" on page 124 for details) and click OK. Otherwise, go to step 7.
- Click OK to apply the settings on all tabs to the active chart.

Scatter Charts and Filtered Data


You can use the Filter tab of the Forecast Preferences dialog to include or exclude specific ranges of data from forecast charts ("Filter Tab" on page 69). If you include a filtered forecast in a scatter chart, you can choose whether or not to show filtered points in the chart.

- To change this setting:
- 1 Open the scatter chart and choose Preferences, then Chart.
- Click the Chart Type tab.
- Check or uncheck Show Points That Have Been Filtered Out to show or hide the filtered points, respectively.
- Click OK.

By default, filtered points are displayed in scatter charts as very light (grayed) points or symbols.

Figure 50 shows the same data as Figure 47 except Material 3 Reliability has been filtered to only include data between 1.08 and 1.23. Excluded data shows as very light triangles while included data is plotted normally, in this case as transparent blue triangles, size 4.

Figure 50 Scatter Chart Showing Filtered Points



8

Creating Reports and Extracting Data

In This Chapter

Introduction	161
Creating Reports	161
Extracting Data	165

Introduction

This chapter describes how to create reports with charts and data and how to extract data for use in other applications.

See the following sections for instructions:

- "Creating Reports" on page 161
- "Extracting Data" on page 165


Creating Reports

You can generate pre-defined reports for a simulation or you can create a custom report with any or all of the following items:

- Report summary
- Forecasts
- Assumptions
- Decision variables
- Overlay charts
- Trend charts
- Sensitivity charts
- Scatter charts

Figure 51 shows part of a forecast report for the Vision Research model.

Figure 51 Sample Forecast Report



The following sections describe how to create and work with reports:

- "Basic Steps for Creating Reports" on page 162
- "Setting Report Options" on page 163
- "Defining Custom Reports" on page 164
- "Report Processing Notes" on page 165

Basic Steps for Creating Reports

To create a report:



1 Choose Analyze, then Create Report.

(In Microsoft Excel 2007 or later, click the upper half of the Create Report icon. If you click the lower half, you can print a pre-defined report with current options. To change options settings, choose Report Preferences before you choose a report.)

The Create Report Preferences dialog opens.

- 2 Click an icon to choose a report:
 - Assumptions report summary plus assumption parameters, charts, and correlations.
 - Decision Variables decision variable bounds, variable types, and step size (if discrete).

- Forecasts report summary plus forecast summaries, charts, statistics, percentiles, and capability metrics if generated.
- Full, the default all sections and details except assumption statistics and percentiles.
- Index only forecast, assumption, and decision variable summaries.
- Custom displays the Custom Report dialog for report definition.
- OptQuest if you have OptQuest and have active optimization data, displays OptQuest results.
- Predictor if you have run Predictor and have active time series forecast data, displays Predictor results.
- If you choose Custom, click the Custom button and complete the Custom Report dialog ("Defining Custom Reports" on page 164). Otherwise, continue with step 4.
- 4 Click the Options tab to set a location and format for the report ("Setting Report Options" on page 163).

(In Microsoft Excel 2007 or later, if you clicked the lower half of the Create Report icon, choose Report Preferences to set a location and format for the report before you choose a report.)

5 When all settings are complete, click **OK**.

Crystal Ball creates the report as a Microsoft Excel worksheet. You can modify, print, or save the report in the same way as any other worksheet. For example, you can choose the File, then Print option for the spreadsheet model as you would for a normal spreadsheet.

Note: If ### is displayed in the report instead of a numeric value, try making the column wider to show the entire number.

Setting Report Options

Report options specify the report location and format.

- To set report options:
- 1 In the Create Report dialog, click the Options tab.
 - (In Microsoft Excel 2007 or later, if you clicked the lower half of the Create Report icon, choose Report Preferences to set a location and format for the report before you choose a report.)
- In the Location group, choose whether to create the report in a new Microsoft Excel workbook or the current workbook.
 - If you choose Current Workbook, a new sheet is created after the current sheet. You can enter a descriptive name for the new sheet in the Sheet Name field.
- In the Formatting group, indicate whether to include the cell location (workbook, worksheet, and cell address) in report headers and whether to include cell comments.
 - By default, these settings are checked.

- If you choose to include cell comments, only non-Crystal Ball comments are included; Crystal Ball cell comments are redundant and are filtered out.
- In the Chart Format group, choose Image to create a Crystal Ball chart or choose Microsoft Excel to create a Microsoft Excel chart.
 - If you choose Image, you can format charts using the Crystal Ball Chart Preference settings. Image is the default chart format.
- When all settings are complete, click **OK**.

Defining Custom Reports

- To define a custom report:
- Follow the steps in "Basic Steps for Creating Reports" on page 162 to display the Create Report dialog.
- Click the Custom button.
- In the Custom Report dialog, check one or more items in the Report Sections group to include in the
 - Report Summary The report title, date and time, Run Preferences settings, and run statistics.
 - Forecasts Forecast information, including the name, charts, percentiles, statistics, and more.
 - Assumptions Assumption information, including parameters, charts, percentiles, statistics, and correlations.
 - Decision Variables Decision variable information, including type (continuous or discrete) with step if discrete, plus lower and upper bounds.
 - Charts (Overlay, Trend, Sensitivity, Scatter, OptQuest) Includes the selected type of chart(s) in the report. You can scale the size of the charts by entering a percentage in the field.

Note: Microsoft Excel charts cannot be created for assumptions defined as custom distributions.

If you have activated the process capability features and have generated capability metrics, you can include them in the custom report ("Including Capability Metrics in Reports" on page 377).

- As each item is highlighted in the Report Sections group, check appropriate settings in the Details group:
 - Report Summary: Report Title, Date/Time, Run Preferences (Run Preferences settings for the report), Run Statistics
 - Forecasts: Summary, Chart And Size, Statistics, Percentiles, Parameters
 - Assumptions: Chart And Size, Statistics, Percentiles, Correlations
 - Decision Variables: Type, Step Size, Bounds

- Overlay Charts: Chart And Size
- Trend Charts: Chart And Size
- Sensitivity Charts: Chart And Size
- Scatter Charts: Chart And Size
- OptQuest Charts: Chart And Size
- Predictor Charts: Chart And Size

When no details are selected for a custom report section, only a single row is output with the Crystal Ball item name and the cell reference.

- For each item checked in Report Sections, choose whether to display all of that type of item, only selected items, or all open items. If you select Choose, a dialog opens so you can check the box in front of each item to display.
- 6 When these settings are complete, click OK to close the dialog and return to the Create Report dialog.
- 7 In the Create Report dialog, click the Options tab to display it ("Setting Report Options" on page 163).
- 8 When all the report options are set, click OK to close the dialog and create the report.

Report Processing Notes

The following are special notes concerning Crystal Ball reports:

- No section is created if there are no Crystal Ball items for that section.
- Statistics follow the chart by default.
- The Choose, then All options always include restored results if they exist.
- If scroll bars are present in a chart, they are displayed in the report.
- If an assumption has been truncated, you might want to add marker lines to show where the distribution has been truncated. To do this, display the Chart Type tab in the Chart Preferences dialog and set an appropriate Value marker.
- Sensitivity Charts and Scatter Charts in Scatter view include only the portion of those charts available onscreen at the time the report is requested.

Extracting Data

You can extract assumption and forecast information generated during a Crystal Ball simulation. Crystal Ball places the extracted data in the specified worksheet location. You can only extract data after you run a simulation or restore saved results.

- To extract data:
- 1 Choose Analyze, then Extract Data.



The Extract Data dialog opens.

- To select the type of data to extract, choose the appropriate setting in the Select Data To Extract list:
 - Statistics Descriptive statistics summarizing the assumption and forecast values.
 - Percentiles The probability of achieving values below a particular threshold in the selected increments. You can reverse the meaning of the percentiles by changing the setting in the Run Preferences, then Options panel ("Setting Statistics Preferences" on page 82).

Note: If you select Percentiles, a dialog opens so you can choose which percentiles to use. Choose Custom and enter a set of custom percentiles if the set you need is not already available in the dialog.

Chart Bins — For each group interval, or bin, the interval range as well as the probability and frequency of occurrences within the interval for the forecast. This setting is independent of the Chart Preferences density setting that controls how many bins or data points are displayed in a graphic chart.

Note: If you select Chart Bins, the Chart Bins dialog opens. You can enter the number of bins to use and can select whether to use the chart range as displayed or the entire chart range, including extreme values excluded from display.

- Sensitivity Data Sensitivity data (such as the rank correlation coefficient) for all pairs of assumptions and forecasts indicating the strength of the relationship. If you plan to extract sensitivity data, be sure to check Store Assumption Values For Sensitivity Analysis on the Options tab of the Run Preferences dialog ("Setting Options Preferences" on page 81). Note: Data is extracted for all assumptions regardless of assumptions selected for extraction.
- Trial Values The generated assumption and forecast values for each simulation trial.
- Capability Metrics Process capability metric values, if available. If you have activated the process capability features and have generated capability metrics, you can extract them ("Extracting Capability Metrics" on page 375).

Data types are extracted in the order they are displayed in the Select Data To Extract list. You can use the up and down arrows to rearrange the data types.

- In the Forecasts group, select forecasts for data extraction (All, Choose, or None):
 - To include the selected data and restored results for all forecasts in the current simulation, choose All.
 - To include the selected data for only selected forecasts or to include forecasts in selected restored results files, choose Choose.
 - Choose None if you don't want to extract any forecast data.
 - If you choose Choose, the Choose Forecasts dialog opens.

Only forecasts for which data was generated or restored are included in the list. Any restored results are displayed after current data and have file names with the .cbr extension.

By default, the Choose Forecasts dialog is in a hierarchical Tree view. You can click the List button at the top to change it to List view.

- 4 In the Assumptions group, select assumptions for data extraction (All, Choose, or None):
 - To include the selected data for all assumptions and restored results in the current simulation, choose All.
 - To include the selected data for only selected assumptions or to include assumptions in selected restored results files, choose Choose.

The Choose Assumptions dialog opens so you can choose from list of available assumptions. It works similarly to the Choose Forecasts dialog, described previously. For more information on the Choose Assumptions dialog, click the Help button.

- Choose None if you don't want to extract any assumption data.
 - You can only extract data for calculated assumption values using the Extract Data command. To view statistics for theoretical data, display an assumption chart and then choose View, then Statistics.
- If you have active OptQuest or Predictor data, make appropriate settings to extract the target data. For details, see the Crystal Ball Decision Optimizer OptQuest User's Guide or the Crystal Ball Predictor User's Guide.
- 6 Click the Options tab to specify a location or formatting for the extracted data.
- 7 In the Location area of the Options tab:
 - To extract data to a new workbook, choose New Workbook.
 - To extract data to a new worksheet in the active workbook, choose Current Workbook, then New Sheet.
 - To extract data to the current sheet, choose Current Workbook, then Current Sheet.
- 8 Specify the name of the sheet and the first cell of the range where the extracted data will be stored.
- 9 Check settings in the Formatting group to indicate how to format the extracted data:
 - When checked, Include Labels adds row and column headers to the data table. Otherwise, just the numeric values are extracted.
 - When checked, Include Cell Locations adds the workbook, worksheet, and cell address with the object name in the column header. Otherwise, only the object name is displayed.

	Book1	Cell Location labels
	Sheet1!A2	Cell Location labels
Statistics	A2	ĺ
Trials	1000	
Mean	0.02	

- When checked, AutoFormat applies the following formats to extracted data:
 - o Bold font for column headers
 - Border beside row labels
 - o Border beneath column headers
 - o Border before first assumption
 - Numeric formatting to values

- AutoFit width to columns
- 10 You can click the Defaults button at any time to restore the original settings to both tabs of the Extract Data dialog.
- 11 When both the Data and Options tab settings are complete, click OK.

Crystal Ball extracts the simulation data to the specified worksheet location. The extracted data is arranged as columns of forecasts and assumptions and rows of data. You can sort, modify, print, or save the data in the same way as any other spreadsheet.

For examples of extracted data, see "Data Extraction Examples" on page 168.

Data Extraction Examples

The previous section, "Extracting Data" on page 165, describes how to insert simulation data into a worksheet for further analysis. The following figures show samples of different types of data extracted with all Formatting settings checked (forecasts only).

Figure 52 A sample of extracted data, Statistics format

	A	В	С
1	Statistics	Ending Sales Year 3 - Q1	Ending Sales Year 3 - Q2
2	Trials	1000	1000
3	Mean	\$17,081,883	\$17,950,914
4	Median	\$17,042,665	\$17,930,927
5	Mode		
6	Standard Deviation	\$1,117,986	\$1,297,577
7	Variance	\$1,249,891,924,272	\$1,683,706,187,431
8	Skewness	0.19	0.15
9	Kurtosis	3.22	3.19
10	Coeff. of Variability	0.07	0.07
11	Minimum	\$13,695,983	\$14,060,365
12	Maximum	\$21,289,239	\$22,981,379
13	Range Width	\$7,593,255	\$8,921,013
14	Mean Std. Error	\$35,354	\$41,033

Figure 53 A sample of extracted data, Percentiles format

	Α	В	С
17	Percentiles	Ending Sales Year 3 - Q1	Ending Sales Year 3 - Q2
18	0%	\$13,695,983	\$14,060,365
19	10%	\$15,622,926	\$16,312,876
20	20%	\$16,129,311	\$16,837,542
21	30%	\$16,492,819	\$17,288,938
22	40%	\$16,826,501	\$17,603,671
23	50%	\$17,042,665	\$17,930,927
24	60%	\$17,307,813	\$18,215,054
25	70%	\$17,596,651	\$18,592,965
26	80%	\$17,986,610	\$19,018,804
27	90%	\$18,526,765	\$19,659,121
28	100%	\$21,289,239	\$22,981,379

Figure 54 A sample of extracted data, Chart Bins format

	A	В	С
31		Ending Sales Year 3 - Q1	
32	Chart Bins	Minimum	Maximum
33	1	\$13,951,523	\$14,076,737
34	2	\$14,076,737	\$14,201,952
35	3	\$14,201,952	\$14,327,166
36	4	\$14,327,166	\$14,452,381
37	5	\$14,452,381	\$14,577,595
38	6	\$14,577,595	\$14,702,809
39	7	\$14,702,809	\$14,828,024
40	8	\$14,828,024	\$14,953,238
41	9	\$14,953,238	\$15,078,453
42	10	\$15,078,453	\$15,203,667
43	11	\$15,203,667	\$15,328,881
44	12	\$15,328,881	\$15,454,096

Figure 55 A sample of extracted data, Sensitivity Data format

	A	В	C
85	Sensitivity Data		
86	Assumptions	Ending Sales Year 3 - Q1	Ending Sales Year 3 - Q2
87	Growth Year 1 - Q1	0.25	0.22
88	Growth Year 1 - Q2	0.31	0.26
89	Growth Year 1 - Q3	0.23	0.23
90	Growth Year 1 - Q4	0.26	0.24
91	Growth Year 2 - Q1	0.31	0.25
92	Growth Year 2 - Q2	0.22	0.21
93	Growth Year 2 - Q3	0.28	0.28
94	Growth Year 2 - Q4	0.45	0.41
95	Growth Year 3 - Q1	0.46	0.45
96	Growth Year 3 - Q2	0.05	0.43
97	Growth Year 3 - Q3	0.00	-0.03
98	Growth Year 3 - Q4	-0.02	-0.01
99	Coil Diameter, in.		
100	Material 1 Strength		

Figure 56 A sample of extracted data, Trial Values format

	A	В	С
109	Trial values	Ending Sales Year 3 - Q1	Ending Sales Year 3 - Q2
110	1	\$18,849,027	\$19,620,035
111	2	\$16,454,224	\$16,645,784
112	3	\$16,048,233	\$16,565,879
113	4	\$14,838,034	\$14,473,412
114	5	\$14,556,109	\$14,399,614
115	6	\$16,234,351	\$16,397,570
116	7	\$16,924,035	\$18,552,808
117	8	\$16,344,792	\$16,678,518

Crystal Ball Tools

In This Chapter

Introduction	171
Overview	171
Batch Fit Tool	173
Correlation Matrix Tool	181
Tornado Chart Tool	187
Bootstrap Tool	195
Decision Table Tool	203
Scenario Analysis Tool	209
2D Simulation Tool	214
Data Analysis Tool	223

Introduction

This chapter describes the Crystal Ball tools. For each tool, there is a general description, a tutorial, and references to descriptions of all panels, fields, and options.

In addition, see "Numerical Differences and the Compare Run Modes Tool" on page 300 for information about the Compare Run Modes tool, used to determine how much faster the model could run with Extreme Speed, if you have Crystal Ball Decision Optimizer.

Overview

Crystal Ball tools are features in the Run, then More Tools menu that extend the functionality of Crystal Ball. They cover two aspects of Crystal Ball modeling plus integration with other Oracle products:

- "Setup Tools" on page 172
- "Analysis Tools" on page 172
- "Integration Tools" on page 172

This chapter describes each tool and provides a step-by-step example for using it.

Note: Some of the tools can take a long time to run. Should you need to cancel a tool before it stops by itself, click in the Microsoft Excel window and hit Esc. A panel opens and you can choose to End or Continue.

Tip: If you are coding with Visual Basic for Applications (VBA), ensure that you are not at a breakpoint in VBA code when starting any of the Crystal Ball Tools.

Setup Tools

The following are Crystal Ball setup tools, described in the referenced sections:

- "Batch Fit Tool" on page 173 Automatically fits selected probability distributions to multiple data series.
- "Correlation Matrix Tool" on page 181 Rapidly defines and automates correlations of assumptions.
- "Tornado Chart Tool" on page 187 Individually analyzes the impact of each model variable on a target outcome.

Analysis Tools

The following are Crystal Ball analysis tools, described in the referenced sections:

- "Bootstrap Tool" on page 195 Addresses the reliability and accuracy of forecast statistics.
- "Decision Table Tool" on page 203 Evaluates the effects of alternate decisions in a simulation model.
- "Scenario Analysis Tool" on page 209 Displays what inputs created particular outputs.
- "2D Simulation Tool" on page 214 Independently addresses uncertainty and variability using a two-dimensional simulation.
- "Data Analysis Tool" on page 223 Imports one or more series of raw data and performs a variety of analyses on them.

Integration Tools

The integration tools, available with Crystal Ball EPM, enable you to run Crystal Ball simulations against Oracle Hyperion Strategic Finance, Fusion Edition and other Oracle Enterprise Performance Management products such as Oracle Essbase working through Oracle Hyperion Smart View for Office, Fusion Edition. For information about the Crystal Ball Decision Optimizer Strategic Finance and Enterprise Performance Management integration tools, see the Oracle Crystal Ball Enterprise Performance Management, Fusion Edition Integration Guide.

Tools and Run Preferences

For consistency from one run to the other, it can be helpful to choose Latin Hypercube sampling on the Sampling tab of the Run Preferences dialog. You can also choose to use the same sequence of random numbers with a specific random seed value. The default is 999.

Note: If you do not set a random seed value before running one of the Crystal Ball tools, Crystal Ball will set a value for you and will notify you of this when you run a tool.

Batch Fit Tool

The Batch Fit tool fits probability distributions to multiple data series. You can select any or all of the probability distributions (binomial, normal, triangular, uniform, and so on) to fit to any number of series limited only by the size of the spreadsheet.

Batch Fit is intended to help you create assumptions when you have historical data for several variables. It selects which distribution best fits each series of historical data, and gives you the distribution and its associated parameters for you to use in the model. This tool also gives you a table of goodness-of-fit statistics for the best fitting distribution and provides a matrix of correlations calculated between multiple data series so you can easily see which series are related and to what degree.

To use the Batch Fit tool, the data series must be contiguous (in adjacent rows or columns) in either rows or columns.

You can select any combination of probability distributions to fit to all the data series.

Using the Batch Fit Tool

To use the Batch Fit tool, you start it and then use the wizard panels to define inputs and outputs, set options, and run the tool.

Topics in this section include:

- "Starting the Batch Fit Tool" on page 174
- "Using the Batch Fit Welcome Panel" on page 174
- "Setting Batch Fit Input Data Options" on page 174
- "Setting Batch Fit Fitting Options" on page 175
- "Setting Batch Fit Output Options" on page 176

Also see "Batch Fit Example" on page 177.

Starting the Batch Fit Tool

- To start the Batch Fit tool:
- 1 In Microsoft Excel with Crystal Ball loaded, open or build the workbook to analyze.
- Select Run, then More Tools, then Batch Fit. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

If this is the first time you have opened the Batch Fit tool, the Welcome panel opens.

Using the Batch Fit Welcome Panel

The Welcome panel opens the first time you use the Batch Fit tool. It describes the tool and its use. The controls for this panel are:

- Next Opens the Input Data panel for specifying the location of the data series.
- Run Runs the Batch Fit tool.
- To continue with the Batch Fit tool, click Next.

The Input Data panel opens.

Setting Batch Fit Input Data Options

The Input Data panel of the Batch Fit wizard indicates the location of the data to fit to the distributions selected in the next panel. You can also specify other input-related options.

When you open this panel, the Batch Fit data selector chooses possible data for fitting. This information is displayed in the Location Of Data Series field and the illustration. You can choose different data if necessary. The fields and options in this panel are:

- Location Of Data Series Used to enter or interactively select the cells that contain data to fit. If the data has headers or labels at the beginning of the rows or columns of data, include them in the selection and select the appropriate Headers setting(s). The data must be in adjacent rows or columns.
- Orientation Indicates whether the data is in rows or columns. Data In Rows indicates
 that data is in horizontal rows. Data In Columns indicates that data is in vertical columns.
- Headers Indicates whether the data has headers and/or labels and whether they are located in the top (first) row or the left (first) column (varies with orientation). Checked items are used in the output. Top Row Has Headers/Labels, when checked, includes text in the top (first) row in the selection. Left Column Has Labels/Headers, when checked, includes text in the left (first) column in the selection.
- Back Returns to the Welcome panel.
- Next Advances to the Fitting Options panel.
- Run Runs the tool, automatically fitting distributions to the data and creating output of the assumptions and statistics.

➤ When all Input Data settings are complete, click Next to open the Fitting Options panel.

Setting Batch Fit Fitting Options

The Fitting Options panel of the Batch Fit wizard indicates which distributions to fit to each set of data. The Fitting Options panel has these settings:

- Distributions To Fit Indicates the distributions to use for fitting:
 - AutoSelect Crystal Ball chooses the best type of distribution for fitting
 - All Continuous Fits the data to all those distributions in which every value in the distribution's range is possible (these distributions are displayed as solid shapes on the Distribution Gallery)
 - All Discrete Fits the data to all discrete (non-continuous) distributions in the Distribution Gallery (except Yes-No)
 - Choose Displays another dialog from which you can select a subset of the distributions to include in the fitting.
- Rank By Goodness-of-Fit Statistic Indicates which ranking method to use for determining the best fit:
 - AutoSelect Crystal Ball chooses the best available goodness-of-fit statistic to use for ranking.
 - Anderson-Darling Closely resembles the Kolmogorov-Smirnov method, except that
 it weights the differences between the two distributions at their tails greater than at their
 mid-ranges; use this method when you need a better fit at the extreme tails of the
 distribution.
 - Ohi-Square Oldest and most common of the goodness-of-fit tests; gauges the general accuracy of the fit by breaking down the distribution into areas of equal probability and comparing the data points within each area to the number of expected data points.
 - o Kolmogorov-Smirnov Finds the largest vertical distance between the two cumulative distributions.
- Lock Parameters When you check this box or click the Ellipsis (...) button, opens the
 Lock Parameters dialog, where you can choose parameters to lock during fitting and specify
 their values.

Note: If you know location, shape, or other parameter values that might help create a more accurate fit with certain distributions, check Lock Parameters and enter appropriate values in the Lock Parameters dialog. For details, see "Locking Parameters When Fitting Distributions" on page 49.

- Back Returns to the Input Data panel.
- Next Opens the Output Options panel.
- Run Runs the tool, automatically fitting distributions to the data and creating output of the assumptions and statistics.

When all Fitting Options settings are complete, click Next to open the Output Options panel.

Setting Batch Fit Output Options

The Output Options panel of the Batch Fit wizard sets output options that control the tool. Available settings and buttons are:


- Fit Results (Assumptions), Locations Specifies locations of results:
 - New Workbook Places results in a new workbook.
 - Current Workbook Places results in the current workbook. You can choose New Sheet, which places results in a new sheet of the current workbook, or Current Sheet, which places results in an existing sheet of the current workbook; you can specify the first cell of the output range (results are displayed to the right and downward from that cell).
- Fit Results (Assumptions), Format Output When checked, uses special cell formatting for the data in the output.
- Correlations Specifies whether correlations are generated and defined:
 - Show Correlation Matrix Between Data Series When checked, correlates the data series with each other and displays the results in a matrix.
 - Define Correlations Between Assumptions If The Absolute Value Of The Correlation Is Above ___ — When checked, defines correlations between assumptions if the absolute value of the correlation is above the indicated level. The default is 0.0, which intercorrelates all assumptions.
 - Keep Correlations Linked To Matrix When checked, changes correlations between assumptions if a new matrix is generated.
- Reports Indicates which reports to generate and their sheet names:
 - Create Goodness-Of-Fit Report When checked, creates a goodness-of-fit report on a separate worksheet with the specified name.
 - Show All Goodness-of-Fit Statistics When checked, shows all goodness-of-fit statistics, not just the chosen type.
 - Create Assumptions Report When checked creates an assumptions report for all
 assumptions created by Batch Fit on a separate worksheet with the specified name.
 - Full Statistics When checked, the assumptions report includes values for all statistics and percentiles (deciles) for each assumption.
- Back Returns to the Fitting Options panel.
- Run Runs the tool, automatically fitting distributions to the data and creating output of the assumptions and statistics.

➤ When all Output Options settings are complete, click Run to run the Batch Fit tool.

Batch Fit Example

In the Crystal Ball Examples folder, there is a Magazine Sales.xls spreadsheet you can use to experiment with the Batch Fit tool. This spreadsheet model (Figure 57) shows the estimated gross profit resulting from newsstand sales of four of the company's most popular magazines.

Figure 57 Magazine Sales Workbook



In this model, cells C5 through F5 are formulas that refer to the first row of data on the Sales Data worksheet. However, the model would be more accurate if you replaced these formulas with assumptions based on historical data. To do this, you can use the Batch Fit tool to generate an assumption for each data column of the Sales Data worksheet. Then, you can use Crystal Ball commands to copy and paste those assumptions from the Batch Fit data to the first data row of the Magazine Sales model. Because Crystal Ball cannot define assumptions in formula cells, you need to convert the formulas in C5 to F5 to values before pasting the assumption cells.

Example topics include:

- "Generating Batch Fit Results" on page 177
- "Using Batch Fit Results in a Model" on page 179
- "Interpreting Batch Fit Results" on page 181

Generating Batch Fit Results

- To work through the Batch Fit example:
- 1 In Microsoft Excel with Crystal Ball loaded, open the example workbook Magazine Sales.xls.
- Start the tool; select Run, then More Tools, then Batch Fit. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)
 - If this is the first time you have opened the Batch Fit tool, the Welcome panel opens.
- 3 If the Welcome panel opens, click Next to display the Input Data panel.

On the Input Data panel, the Batch Fit data selector attempts to select relevant data and indicate its orientation. In this case, you need to select data on another worksheet.


- 4 On the Input Data panel, click the Select Cells icon beside the Location Of Data Series field.
- 5 Select the Sales Data worksheet.
- 6 Select cells A1 through D361.
- 7 Click **OK** to return to the Batch Fit panel.
- 8 Select **Data in Columns**, the default.
- 9 Check the Top Row Has Headers setting.
- 10 Click Next.

The Fitting Options panel opens.

- 11 In the Fitting Options panel, choose All Continuous in the Distributions To Fit group.
- 12 Choose Anderson-Darling in the Rank By Goodness-of-Fit Statistic group.
- 13 Click Next.

The Output Options panel opens, similar to Figure 58 but with different settings.

Figure 58 Output Options panel, Batch Fit tool



- 14 For Fit Results, choose Current Workbook and then choose Current Worksheet.
- 15 Click the Select Cells icon beside the Upper Left Corner Of Output Range box.
- 16 Select cell B13 on the Model worksheet.
- 17 Click OK to return to the Batch Fit panel.
- 18 Make sure the Format Output box is checked.
- 19 Include a correlation analysis of the data:
 - a. Check "Show correlation matrix between data series."

b. Check "Define correlations between assumptions..." For now, leave 0.0 in the box to display all correlations.

The threshold must be between 0 and 1 (inclusive). The tool displays the correlation matrix at the end of the results section.

- c. For now, do not check "Keep correlations linked to matrix."
- 20 Check the Create Goodness-Of-Fit Report but for now, do not check Create Assumptions Report.

Use the default Sheet Name setting.

the settings should look like those in Figure 58.

21 Click Run.

The tool fits each selected distribution to each data series. The Batch Fit results are displayed on the Model worksheet below the existing table, similar to Figure 59.

Magazine Sales.xls Magazine Checkout Sales 3 Magazine A Magazine B Magazine C Magazine D 4 5 6 7 8 9 Sales Volume 362 646 508 562 Retail Price \$7.95 \$4.95 \$3.95 \$5.95 Cost of Goods \$2.20 \$3.80 \$1.95 \$2.40 \$1 292 Gross Profit \$1.546 \$1.502 \$1,805 11 12 13 14 15 16 \$6,144 **Total Gross Profit: Data Series:** Sales Volume A Sales Volume B Sales Volume C Sales Volume D Distribution: Best Fit: Max Extreme Triangular Beta Gamma Anderson-Darling: 0.0824 0.2129 0.4141 0.0893 17 18 P-Value: 0.989 ---0.997 19 20 21 22 23 Correlations: Sales Volume A Sales Volume B Sales Volume C Sales Volume D Sales Volume A 1.0000 0.0166 -0.0128 -0.0124 0.0166 1.0000 -0.0059 0.0607 Sales Volume B Sales Volume C -0.0128 -0 0059 1.0000 -0 0392 Sales Volume D -0.01240.0607 -0.03921.0000 ◆ ▶ N Description / Goodness-of-fit \ Model / Sales Data /

Figure 59 Batch Fit Results Below the Model

Using Batch Fit Results in a Model

Now that assumptions have been generated from historic sales data, the next steps are to convert Sales Volume formulas to values and replace them with the new assumptions. Then, you can run a simulation with more accurate results.

- To use Batch Fit results in a model:
- 1 Select the formulas in cells C5 through F5.
- 2 Use Microsoft Excel's Copy and Paste Special commands to copy the formula cells and paste them back into the same locations. When the Paste Special dialog opens, choose Values as the Paste setting. Leave the other default settings and click OK.

- Copy assumption data into the worksheet.
 - a. Select cells C14 through F14.
 - b. Select Define, then Copy Data.

Notice this is Crystal Ball's Copy Data command, not Microsoft Excel's Copy command.

This function copies Crystal Ball data only, not the cell value.

- c. Select cells C5 through F5.
- d. Select Define, then Paste Data.

The assumptions copy to the first row of the table, leaving the current values.

- e. Select cells C14 through F14.
- f. Select Define, then Clear Data.

This deletes the original assumptions generated by Batch Fit to eliminate redundant assumptions and speed up the simulation, as shown in Figure 60.

Figure 60 Crystal Ball Assumptions in Cells C5 through F5

Α	В	С	D	Е	F
	Magazine Checkout Sales				
		Magazine A	Magazine B	Magazine C	Magazine D
	Sales Volume	562	362	646	508
	Retail Price	\$4.95	\$7.95	\$3.95	\$5.95
	Cost of Goods	\$2.20	\$3.80	\$1.95	\$2.40
	Gross Profit	\$1,546	\$1,502	\$1,292	
	Total Gross Profit:	\$6,144			
	Data Series:	Sales Volume A	Sales Volume B	Sales Volume C	Sales Volume D
	Distribution:	498.96	354.37	639.2	
	Best Fit:	Max Extreme	Triangular	Beta	Gamma
	Anderson-Darling:	0.0824	0.2129	0.4141	0.0893
	P-Value:	0.989			0.997
	Correlations:	Calaa Waliinaa A	Sales Volume B	Calaa Waliimaa C	Calaa Valuma D
	Sales Volume A	1.0000	0.0166	-0.0128	Carco Coranio B
	Sales Volume B	0.0166	1.0000		
	Sales Volume C	-0.0128	-0.0059	1.0000	
	Sales Volume D	-0.0124	0.0607	-0.0392	
	Gales Volume D	-0.0124	0.0001	-0.0332	1.0000
	Description / Goodness-of-fit / Model / Sales Dat	a /	< 1111		

- In the Run, then Run Preferences, then Sampling dialog:
 - Set Random Number Generation to use the Same Sequence Of Random Numbers and a seed value of 999.
 - Select Monte Carlo simulation.

When using this tool, use these options to make the resulting simulations comparable.

- In the Run, then Run Preferences, then Trials dialog, set the Number Of Trials To Run to 500. 5
- Click OK.
- Select Run, then Start Simulation.

The simulation runs until it stops.

Interpreting Batch Fit Results

In this example, you had historical data with no trend or seasonality for all the variables.

If the historical data have a time-series element, trend, or seasonality, you should use Predictor instead of the distribution fitting functionality.

When the Batch Fit tool runs, it fits each column of data to each selected distribution. For each fit of a distribution to a set of data, the tool calculates the indicated goodness-of-fit test statistic. The distribution with the best fit is placed in the spreadsheet to create an assumption cell that you can copy to the appropriate location in the model.

In this example, the four columns of data resulted in four assumptions that you copied and pasted to the appropriate locations. The forecast is already defined for you, and running a simulation produces a forecast chart of the total gross profit from the Magazine Sales workbook. In the Total Gross Profit forecast chart, if you replace –Infinity with \$5,500, you find that the certainty or probability of making this amount of profit is almost 75% (Figure 61).




Figure 61 Magazine Sales profit from newsstand sales

Correlation Matrix Tool

The following sections describe the Correlation Matrix tool and its use:

- "About Correlations" on page 182
- "About the Correlation Matrix Tool" on page 182
- "Using the Correlation Matrix Tool" on page 182
- "Correlation Matrix Example" on page 184

About Correlations

When the values of two variables depend on each other in any way, you should correlate them to increase the accuracy of the simulation's forecast results.

There are two types of correlations:

- Positive correlation Indicates that two assumptions increase or decrease together. For example, the price of gasoline and shipping costs increase and decrease together.
- Negative correlation Indicates that an increase in one assumption results in a decrease in the other assumption. For example, the more items you buy from a particular vendor, the lower the unit cost.

The correlation coefficient range is -1 to 1, where 0 indicates no correlation. The closer the coefficient is to ± 1 , the stronger the relationship between the assumptions. You should never use a coefficient of ± 1 ; represent relationships this closely correlated with formulas in the spreadsheet model.

About the Correlation Matrix Tool

In Crystal Ball, you usually enter correlations one at a time using the Correlation dialog. Instead of manually entering the correlations this way, you can use the Correlation Matrix tool to define a matrix of correlations between assumptions in one simple step. This saves time and effort when building the spreadsheet model, especially for models with many correlated assumptions.

The correlation matrix is either an upper or lower triangular matrix with 1s along the diagonal. When entering coefficients, think of the matrix as a multiplication table. If you follow one assumption along its horizontal row and the second along its vertical column, the value in the cell where they meet is their correlation coefficient. The matrix contains only the correlation coefficients you enter.

Figure 62 Correlation Matrix

	Assumption 1	Assumption 2	Assumption 3
Assumption 1	1.000		
Assumption 2		1.000	
Assumption 3			1.000

If you enter inconsistent correlations, Crystal Ball tries to adjust the correlations so they don't conflict. For more information on inconsistent correlations, see "Specifying Correlations Between Assumptions" on page 49.

Using the Correlation Matrix Tool

To use the Correlation Matrix tool, start it and then use the wizard panels to choose assumption, set options, and run the tool.

This section has the following topics:

- "Starting the Correlation Matrix Tool" on page 183
- "Selecting Assumptions to Correlate" on page 183
- "Setting Correlation Matrix Location and Orientation Options" on page 183
- "Running the Correlation Matrix Tool" on page 184

Also see "Correlation Matrix Example" on page 184.

Starting the Correlation Matrix Tool

To start the Correlation Matrix tool, select Run, then More Tools, then Correlation Matrix. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

The Select Assumptions panel opens.

Selecting Assumptions to Correlate

The Select Assumptions panel of the Correlation Matrix tool indicates which assumptions are to be correlated. The settings in this panel include:

Available Assumptions — Lists all the assumptions defined in the active workbook. To move an assumption into the Selected Assumptions list, either double-click an assumption or select an assumption and click, then>. You can also make an extended selection using the Shift or Ctrl keys.

By default, no assumptions are in this list.

- Top Moves one or more selected assumptions in the Selected Assumptions list to the top of the list.
- Bottom Moves one or more selected assumptions in the Selected Assumptions list to the bottom of the list.
- Up Moves one or more selected assumptions in the Selected Assumptions list up one position.
- Down Moves one or more selected assumptions in the Selected Assumptions list down one position.
- Next Opens the Specify Options panel for defining correlation matrix options.

Setting Correlation Matrix Location and Orientation Options

The Specify Options panel of the Correlation Matrix tool sets two kinds of options, described in these sections:

- "Location of Matrix Options" on page 184
- "Orientation Options" on page 184

This panel also has the following buttons:

- Back Returns you to the Select Assumptions panel.
- Start Creates the correlation matrix according to the options you specify.

Location of Matrix Options

The Location Of Matrix options indicate whether to create the correlation matrix on the active worksheet or on a new worksheet. Available options include:

- Create A Temporary Correlation Matrix On A New Worksheet Creates the matrix on a
 new worksheet. Use this option if you do not want to permanently alter the spreadsheet.
 After filling in the matrix, click Load The Matrix to enter the coefficients into the model.
 This is the default matrix location.
- Create And Link To A Correlation Matrix On The Active Worksheet Creates the matrix
 on the active worksheet and links it with cell references, so if you change the cells, the
 correlations update automatically. Select this option to embed the correlation matrix in the
 model. If you select this option, you must indicate where to put the matrix in the model
 using the Specify Upper Left Corner option.
- Specify Upper Left Corner Enables you to select the first (upper-left) cell for the matrix by either entering the cell reference in the field or clicking on the cell in the worksheet. This option is only available if you create the matrix on the active worksheet.
- Format Matrix Adds borders and headings, and changes column widths to make the
 correlation matrix more readable. This option is only available if you create the matrix on
 the active worksheet. This option is selected by default.

Orientation Options

The Orientation options indicate whether to make the correlation matrix an upper triangular or a lower triangular matrix.

- Upper Triangular Matrix Creates the matrix with cells to fill in above the diagonal. This is the default.
- Lower Triangular Matrix Creates the matrix with cells to fill in below the diagonal.

Running the Correlation Matrix Tool

To run the Correlation Matrix tool when all settings are complete, click **Start**.

Correlation Matrix Example

In the Crystal Ball Examples folder, there is a Portfolio Allocation.xls workbook you can use to experiment with the Correlation Matrix tool. This spreadsheet calculates the total expected return for an investment model. In this example, you will run a simulation without correlations and then add the correlations and rerun the simulation for comparison.

- To start the Correlation Matrix tool:
- In Microsoft Excel with Crystal Ball loaded, open the workbook Portfolio Allocation.xls.
- Set the following options in the Run, then Run Preferences dialog:
 - On the Trials tab, set Maximum Number Of Trials to 500.
 - On the Sampling tab, set Random Number Generation to Use Same Sequence Of Random Numbers and set an Initial Seed Value of 999.
 - Also on the Sampling tab, set Sampling Method to Monte Carlo.
- 3 Run a simulation by selecting **Run**, then **Start Simulation**.

Choose View, then Statistics in the forecast chart. The forecast statistics for the simulation are similar to those shown in Figure 63, following.




Figure 63 **Uncorrelated Simulation Statistics**

- Choose Run, then Reset to reset the simulation.
- Select Run, then More Tools, then Correlation Matrix. (In Microsoft Excel 2007 or later, select More **Tools** in the **Tools** group.)

The Select Assumptions panel opens.

- Include all the assumptions in the correlation matrix by moving all the assumptions from the Available Assumptions list to the Selected Assumptions list using one of these alternatives:
 - Double-click each assumption to move.
 - Select each assumption to move and click the to move it.
 - Make an extended selection using the Shift or Ctrl keys and then click
- Click Next.

The Specify Options panel opens.


- Select the following settings:
 - Create A Temporary Correlation Matrix On A New Worksheet

Upper Triangular Matrix

9 Click Start.

The tool creates a temporary matrix in a new workbook.

10 Enter the following correlation coefficients into the matrix.



Leaving a cell blank is not the same as entering a zero. Values that are not specified in the matrix will be filled in with estimates of appropriate values when the simulation runs.

11 Click Load The Matrix.

The tool loads the correlation coefficients from the matrix into the Crystal Ball model. Loading a full correlation matrix for all assumptions can take several minutes or even longer.

If a Matrix Successfully Loaded message is not displayed, press Tab or Return to exit the current cell and then click Load The Matrix again.

12 Reset the simulation.

13 Rerun the simulation.

The forecast statistics for the correlated simulation are shown in Figure 64.




Figure 64 Correlated Simulation Statistics

The standard deviation is now much higher than the original simulation due to the correlations. The original model without the correlations ignored this risk factor and its effects.

Tornado Chart Tool

The Tornado Chart tool measures the impact of each model variable one at a time on a target forecast. The tool displays the results in two ways, described in these sections:

- "Tornado Chart" on page 187
- "Spider Chart" on page 189

This method differs from the correlation-based sensitivity method built into Crystal Ball because this tool tests each assumption, decision variable, precedent, or cell independently. While analyzing one variable, the tool freezes the other variables at their base values. This measures the effect of each variable on the forecast cell while removing the effects of the other variables. This method is also known as "one-at-a-time perturbation" or "parametric analysis."

The Tornado Chart tool is useful for:

- Measuring the sensitivity of variables that you have defined in Crystal Ball.
- Quickly pre-screening the variables in the model to determine which ones are good candidates to define as assumptions or decision variables. You can do this by testing the precedent variables of any formula cell. See for more information on precedents.


The following sections describe how to create and use tornado charts and spider charts:

- "Using the Tornado Chart Tool" on page 189
- "Tornado Chart Example" on page 192
- "Limitations of the Tornado Chart Tool" on page 195

Tornado Chart

The Tornado Chart tool tests the range of each variable at percentiles you specify and then calculates the value of the forecast at each point. The tornado chart (Figure 65) illustrates the swing between the maximum and minimum forecast values for each variable. The variable that causes the largest swing is displayed at the top and the variable that causes the smallest swing is displayed at the bottom. The upper variables have the most effect on the forecast, and the lower variables have the least effect on the forecast.

Figure 65 Tornado Chart




The bars next to each variable represent the forecast value range across the variable tested, as discussed previously. Next to the bars are the values of the variables that produced the greatest swing in the forecast values. The bar colors indicate the direction of the relationship between the variables and the forecast.

For variables that have a positive effect on the forecast, the upside of the variable (shown in blue) is to the right of the base case (the initial value in the cell before running the simulation) and the downside of the variable (shown in red) is to the left side of the base case. For variables that have a reverse relationship with the forecast, the bars are reversed.

When a variable's relationship with the forecast is not strictly increasing or decreasing, it is called non-monotonic. In other words, if the minimum or maximum values of the forecast range do not occur at the extreme endpoints of the testing range for the variable, the variable has a nonmonotonic relationship with the forecast (Figure 66).

Figure 66 A non-monotonic variable



If one or more variables are non-monotonic, all the variable bars are the same color all the way across.

Spider Chart

The spider chart (Figure 67) illustrates the differences between the minimum and maximum forecast values by graphing a curve through all the variable values tested. Curves with steep slopes, positive or negative, indicate that those variables have a large effect on the forecast, while curves that are almost horizontal have little or no effect on the forecast. The slopes of the lines also indicate whether a positive change in the variable has a positive or negative effect on the forecast.

Target forecast Material 1 Reliability 1.20 ← Material 1 Strength 1:15 Coil Diameter, in. Forecast axis values Assumptions Number of Coils 1.10 and other variables Shearing Modulus o Elasticity 1.05 Wire Diameter, in 1.00 Spring Deflection, in 50.0% 70.0% 90.0% 10.0% Percentiles of the variables

Figure 67 Spider chart

A maximum of 250 variables can be displayed in these charts.

Using the Tornado Chart Tool

To use the Tornado Chart tool, start it, specify a target to analyze, specify input variables, set options, and run it as described in the following topics:

- "Starting the Tornado Chart Tool" on page 189
- "Specifying a Tornado Chart Analysis Target" on page 190
- "Specifying Tornado Chart Input Variables" on page 190
- "Specifying Tornado Chart Options" on page 191
- "Running the Tornado Chart Tool" on page 192

Also see "Tornado Chart Example" on page 192.

Starting the Tornado Chart Tool

To start the Tornado Chart tool, select Run, More Tools, and then select Tornado Chart. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

The Specify Target panel opens.

Specifying a Tornado Chart Analysis Target

The Specify Target panel in the Tornado Chart tool indicates which forecast to analyze. The Specify Target panel contains the following settings:

- Forecast List Lists all the forecast cells in all open spreadsheets. When you select a forecast
 from the list, its cell information is automatically displayed in the Enter Target Cell field.
 The first forecast is selected by default.
- Enter Target Cell Describes the cell location of the selected forecast or formula. If you select a forecast from the list, the cell information is automatically displayed in this field. You can use this field to select a formula cell instead of a forecast.
- Next Opens the Specify Input Variables panel.

When settings are complete, click Next to open the Specify Input Variables panel.

Specifying Tornado Chart Input Variables

The Specify Input Variables panel of the Tornado Chart tool specifies the assumptions, decision variables, and precedents to include in the tornado and spider charts. You can include any value cell in the tornado chart calculations. However, the cells are usually:

- Assumptions Cells defined as assumptions in Crystal Ball.
- Decision variables Cells defined as decision variables in Crystal Ball.
- Precedents All cells within the active spreadsheet that are referenced as part of the formula
 or a sub-formula of the target cell. Precedents for this tool are handled differently than
 standard Microsoft Excel precedents in that they cannot trace beyond the active spreadsheet.
 Therefore, you can only use precedents on the active spreadsheet as input variables.

The Specify Input Variables panel includes these settings:

- Input Variable List Lists all the variables selected for the tornado and spider charts.
- Add Assumptions Adds all assumptions from all open worksheets to the input variable list.
- Add Decision Variables Adds all decision variables from all open worksheets to the input variable list.
- Add Precedents Adds all precedents of the target cell from all open worksheets to the input variable list.
- Add Range Lets you select a range of cells from the open worksheet to add to the input variable list. If you click this button, an Input panel opens, asking you to enter a cell range or to select a range of cells from the spreadsheet. You must click OK to accept the selected range.
- Remove This button removes the selected variable from the input variable list.
- Remove All Removes all of the items from the input variable list.
- Back Returns to the Specify Target panel.
- Next Opens the Specify Options panel.

When settings are complete, click Next to open the Specify Options panel.

Specifying Tornado Chart Options

The Specify Options panel of the Tornado Chart tool sets options that control the tool. The groups of options in this panel are described in the following sections:

- "Tornado Method Options" on page 191
- "Tornado Input Options" on page 191
- "Tornado Output Options" on page 191

Other controls are:

- Back Returns to the Specify Input Variables panel.
- Start Starts the tool, generating the tornado and spider charts.

Tornado Method Options

The Tornado Method options in this panel are:

- Percentiles Of The Variables Indicates that the tool should test the variables using percentiles of the assumption distributions or percentiles of the decision variable ranges. This is the default.
- Percentage Deviations From The Base Case Indicates that the tool should test the variables using small changes that are specified percentages away from the base case. This is the only option available if you selected variables other than assumptions or decision variables. The tool treats discrete decision variables as continuous for this second method.

Tornado Input Options

The Tornado Input options include:

- Testing Range Defines the range in which the tool samples the variables. The choices are either: the percentile range (if the tornado method is Percentiles Of The Variables) or the percent from the base case (if the tornado method is Percentage Deviations From The Base Case). The default is 10% to 90% percentiles or -10% to 10% deviation.
- Testing Points Defines how many values to test in the testing range. The testing points are evenly distributed across the testing range. Testing more points than just end points detects non-monotonic variable relationships. The default is five testing points.
- For Base Case Specifies whether to define the base case as either the existing cell values or the median values of the variables. If the tornado is based on a percentage deviation, only the cell values option is available. The default is the median cell values.

Tornado Output Options

The Tornado Output options are:

- Tornado Chart Generates a tornado chart showing the sensitivity of the variables using range bars.
- Spider Chart Generates a spider chart showing the sensitivity of the variables using sloping curves.
- Show Top Variables Indicates how many variables to include in the tornado charts if there are a lot of variables. The charts can only clearly show about 20 variables.

Running the Tornado Chart Tool

➤ When all settings are complete, click Start to run the Tornado Chart tool and generate selected charts.

Tornado Chart Example

In the Crystal Ball Examples folder, there is a Reliability.xls workbook you can use to experiment with the Tornado Chart tool. This spreadsheet model predicts the reliability of a spring using three different construction materials. This section includes the following example topics:

- "Generating Tornado and Spider Charts" on page 192
- "Interpreting Tornado and Spider Chart Results" on page 194

Generating Tornado and Spider Charts

- To run the Tornado Chart tool:
- 1 In Microsoft Excel with Crystal Ball loaded, open the workbook Reliability.xls.
 - If you have any other spreadsheets open, close them first, because the tool gathers all Crystal Ball definitions from all open workbooks.
- Select Run, then More Tools, then Tornado Chart. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)
 - The Specify Target panel opens. All of the forecasts from Reliability.xls are displayed in the list.
- 3 Select the "Material 1 Reliability" forecast.
- 4 Click Next.

The Specify Input Variables panel opens.

You can include any value cell in the tornado chart calculations, using cell names if available. However, the cells are usually defined as one of the following types:

 Assumptions — Cells defined as assumptions in Crystal Ball. For more information about assumptions, see "About Assumptions and Probability Distributions" on page 40.

- Decision variables Cells defined as decision variables in Crystal Ball. For more information about defining decision variables, see "Defining Decision Variable Cells" on page 65.
- Precedents All cells within the active worksheet that are referenced as part of the formula or a sub-formula of the target cell. Precedents for this tool are handled differently than standard Microsoft Excel precedents in that they cannot trace beyond the active worksheet. Therefore, you can only use precedents on the active worksheet as input variables.

Click Add Assumptions.

Available assumptions are displayed in the panel.


- Remove Material 2 Strength and Material 3 Strength.
 - Select an assumption to remove.
 - b. Click Remove.
 - c. Repeat steps 6a and 6b for the second assumption to remove.

The last two assumptions have no impact on the target forecast. If you leave them in the list, they will be displayed in the charts even though they are unrelated to the target forecast.

Click Next.

The Specify Options panel opens, similar to Figure 68.

Figure 68 Specify Options panel, Tornado Chart tool



- Set the following options as shown in Figure 68:
 - Testing Range = 10% to 90%.
 - Testing Points = 5.
 - For Base Case = Use Existing Cell Values.
 - Tornado Method = Percentiles Of The Variables.

- Tornado Output = Tornado Chart and Spider Chart.
- Show $_$ Top Variables = 20.

9 Click Start.


The tool creates the tornado and spider charts in their own workbooks with data tables as shown in Figure 69 and Figure 70.

For more information about the Tornado Chart panels, click Help in a panel.

Interpreting Tornado and Spider Chart Results

In this example, six assumptions are listed in the tornado chart (Figure 69). The first assumption, Material 1 Strength, has the highest sensitivity ranking and is the most important. A researcher running this model would investigate this assumption further in the hopes of reducing its uncertainty and, therefore, its effect on the target forecast, Material 1 Reliability.


The last two assumptions, Wire Diameter and Spring Deflection, are the least influential assumptions. Since their effects on the Material 1 Reliability are very small, you might ignore their uncertainty or eliminate them from the spreadsheet.



Tornado Chart with Data Figure 69

The spider chart shows similar information (Figure 70). Material 1, at the top of the tornado chart, has the steepest positive slope in the spider chart. Notice that more reliability values are given, one for each of five levels within the testing range.

Figure 70 Spider chart with Data



Limitations of the Tornado Chart Tool

While tornado and spider charts are useful, there are some limitations:

- Since the tool tests each variable independently of the others, the tool does not consider correlations defined between the variables.
- The results shown in the tornado and spider charts depend significantly on the particular base case used for the variables. To confirm the accuracy of the results, run the tool multiple times with different base cases.

This characteristic makes the one-at-a-time perturbation method less robust than the correlation-based method built into Crystal Ball's sensitivity chart. Hence, the sensitivity chart is preferable, since it computes sensitivity by sampling the variables all together while a simulation is running.

Bootstrap Tool

Bootstrap is a simple technique that estimates the reliability or accuracy of forecast statistics or other sample data. Classical methods rely on mathematical formulas to describe the accuracy of sample statistics. These methods assume that the distribution of a sample statistic approaches a normal distribution, making the calculation of the statistic's standard error or confidence interval relatively easy. However, when a statistic's sampling distribution is not normally distributed or easily found, these classical methods are difficult to use or are invalid.

In contrast, bootstrapping analyzes sample statistics empirically by repeatedly sampling the data and creating distributions of the different statistics from each sampling. The term bootstrap

comes from the saying, "to pull oneself up by one's own bootstraps," since this method uses the distribution of statistics itself to analyze the statistics' accuracy.

There are two bootstrap methods available with this tool:


- One-simulation method Simulates the model data once (creating the original sample) and then repeatedly resamples those simulation trials (the original sample values). Resampling creates a new sample from the original sample with replacement. That is, it returns the selected value to the sample before selecting another value, letting the selector possibly reselect the same value. It then creates a distribution of the statistics calculated from each resample. This method assumes only that the original simulation data accurately portrays the true forecast distribution, which is likely if the sample is large enough. This method isn't as accurate as the multiple-simulation method, but it takes significantly less time to run.
- Multiple-simulation method Repeatedly simulates the model, and then creates a distribution of the statistics from each simulation. This method is more accurate than the one-simulation method, but it might take a prohibitive amount of time.

Note: When you use the multiple-simulation method, the tool temporarily turns off the Use Same Sequence Of Random Numbers option. In statistics literature, the one-simulation method is also called the non-parametric bootstrap, and the multi-simulation method is also called the parametric bootstrap.

Figure 71 Bootstrap simulation methods

One-simulation method Simulate forecast (create original sample) Create resample from original sample (with replacement) Compute statistics for resample Have you reached the number of bootstrap amples ves Form distribution for all resample statistics

Multiple-simulation method



Since the bootstrap technique does not assume that the sampling distribution is normally distributed, you can use it to estimate the sampling distribution of any statistic, even an unconventional one such as the minimum or maximum value of a forecast. You can also easily

estimate complex statistics, such as the correlation coefficient of two data sets, or combinations of statistics, such as the ratio of a mean to a variance.

To estimate the accuracy of Latin Hypercube statistics, you must use the multiple-simulation method.

The following sections describe how to use the Bootstrap tool:

- "Using the Bootstrap Tool" on page 197
- "Bootstrap Example" on page 199

Using the Bootstrap Tool

To use the Bootstrap tool, you start it, define a target forecast or formula cell, choose a bootstrap method, set options, and then run the tool. The following sections describe these steps:

- "Starting the Bootstrap Tool" on page 197
- "Using the Bootstrap Welcome Panel" on page 197
- "Specifying Forecasts to Analyze with the Bootstrap Tool" on page 198
- "Specifying a Bootstrap Tool Method" on page 198
- "Setting Bootstrap Options" on page 199
- "Running the Bootstrap Tool" on page 199

Also see "Bootstrap Example" on page 199.

Starting the Bootstrap Tool

To start the Bootstrap tool, select **Run**, then **More Tools**, and then **Bootstrap**. (In Microsoft Excel 2007 or later, select **More Tools** in the **Tools** group.)

The first time you start Bootstrap, the Welcome panel opens; otherwise, the Target Forecast panel opens.

Using the Bootstrap Welcome Panel

The Welcome panel opens the first time you use the Bootstrap tool. It describes the tool and its use. The controls for this panel are:

- Next Opens the Target Forecast panel for choosing a target forecast.
- Run Runs the Bootstrap tool and generates results.

Click Next to continue to the Target Forecast panel.

Specifying Forecasts to Analyze with the Bootstrap Tool

The Target Forecast panel of the Bootstrap tool indicates which forecast, formula cell, or cell range to analyze. The controls for this panel are:

- Forecast list Lists all the forecast cells in all open spreadsheets. When you select a forecast
 from the list, its cell information is automatically displayed in the data range field. The first
 forecast is selected by default.
- Data range Describes the cell location of the selected forecast or formula. If you select a forecast from the forecast list, the cell information is automatically displayed in this field. You can use this field to select a formula cell instead of a forecast. If you choose a data range, the data must be a single contiguous (connected) block.
- Back Opens the Welcome panel.
- Next Opens the Method panel for defining which simulation method to use.
- Run Runs the Bootstrap tool and generates results.
- Select a target forecast and then click Next to select a simulation method.

Specifying a Bootstrap Tool Method

The Method panel of the Bootstrap tool specifies a bootstrap method and analysis type. It has the following controls:

- Bootstrap Method Selects whether to use the one-simulation or multiple-simulation bootstrap method. For more information on these two methods, see "Bootstrap Tool" on page 195. The default is the one-simulation method.
- Analyze Distributions Of Selects whether to analyze distributions of statistics, percentiles, or capability metrics. If you select Percentiles, you must complete the Percentiles options. The default is Statistics.
- Percentiles If Percentiles is selected in Analyze Distributions Of, selects which target percentiles to analyze. You can select either: deciles (the 10, 20, 30, 40, 50, 60, 70, 80, and 90 percentiles); 2, 5, 50, 90, and 97.5 percentiles; or a custom list of percentiles you enter in the field (between 0 and 100, exclusive) separated by commas.
- Back Returns to the Target Forecast panel for specifying a target forecast.
- Next Opens the Options panel for defining sample and display options.
- Run Runs the Bootstrap tool and generates results.

➤ When Method settings are complete, click Next to open the Options panel.

Setting Bootstrap Options

The Options panel of the Bootstrap tool sets sample and display options for Bootstrap. This panel has the following controls:

- Sample Control Sets the number of bootstrap samples and the number of trials per sample. The default number of bootstrap samples is 200, and the default number of trials is the number set in the Crystal Ball Run Preferences dialog.
- While Running Specifies which forecasts to show while you are running the tool. You can view all the defined forecasts, only the target forecast, or none of the forecasts.
- Back Returns to the Method panel for specifying the bootstrap method to use.
- Run Runs the Bootstrap tool and generates results.
- When Options settings are complete, click Run to run the Bootstrap tool.

Running the Bootstrap Tool

To run the Bootstrap tool, click **Run** in the Options panel.

Results are generated as described in "Interpreting Bootstrap Results" on page 201.

Bootstrap Example

In the Crystal Ball Examples folder, there is a Futura Apartments.xls workbook you can use to experiment with the Bootstrap tool. This spreadsheet model forecasts the profit and loss for an apartment complex.

The Bootstrap example includes the following sections:

- "Generating Bootstrap Results" on page 199
- "Interpreting Bootstrap Results" on page 201

Generating Bootstrap Results

- To start the Bootstrap tool:
- 1 In Microsoft Excel with Crystal Ball loaded, open the workbook Futura Apartments.xls.
- Select Run, then More Tools, then Bootstrap. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

The first time you run the Bootstrap tool, the Welcome panel opens.

3 If the Welcome panel opens, click Next.

The Target Forecast panel opens.

Set the target by selecting **Profit Or Loss** from the forecast list.

Instead, you could choose "Or, enter a data range:" and enter a range of spreadsheet data cells to be used for resampling.

Click Next.

The Method panel opens.

Make sure the one-simulation method and the statistics options are selected.

If you choose Percentiles and then choose Custom, you can select 0 or 100 as well as values between these.

Click Next.

The Options panel opens.

Set the following options:


- Number Of Bootstrap Samples = 200
- Number Of Trials Per Sample = 500
- Show Only Target Forecast is selected

Click Run.

The Bootstrap tool displays a forecast chart of the distributions for each statistic and creates a workbook summarizing the data (Figure 72).

Note: If distribution fitting is turned on for the target forecast, it is turned off during the running of simulations within the tool. When simulations finish, distribution fitting is restored.

Figure 72 Bootstrap Example Results



Notice that the forecast certainty is set to 95%, the precision control confidence level that is displayed on the Trials tab of the Run Preferences dialog.

Other statistics are also calculated, although they are not displayed.

Interpreting Bootstrap Results

The Bootstrap tool displays sampling distributions in forecast charts for the following statistics:

- Mean
- Median
- Mode
- Standard deviation
- Variance
- Skewness
- Kurtosis
- Coefficient of variability
- Mean standard error
- Process capability metrics, if these are set to display

When you use the multiple-simulation method, the tool also displays sampling distributions for these statistics:

- Range minimum
- Range maximum
- Range width

For percentiles, the Bootstrap tool displays the percentile sampling distributions on the overlay and trend charts. To display the individual percentile forecast charts, select Analyze, then Forecast Charts.

Note: If you have the Probability Above A Value option selected in the Run Preferences, then Options panel, the percentiles are reversed in meaning, so that the 1st percentile represents the uppermost 1% and the 99th percentile represents the lowest 1%. For more information on this reversal, "Setting Statistics Preferences" on page 82

The forecast charts visually indicate the accuracy of each statistic (Figure 73). A narrow and symmetrical distribution yields more precise statistics estimates than a wide and skewed distribution.





Figure 73 Bootstrap forecast distribution for the mean

The Statistics view (Figure 74) further lets you analyze the statistics' sampling distribution. If the mean standard error or coefficient of variability is very large, the statistic might not be reliable and might require more samples or more trials. This example has a relatively low standard error and coefficient of variability, so the forecast mean is an accurate estimate of the population mean.

Figure 74 Bootstrap forecast statistics for the mean



The results workbook has a correlation matrix showing the correlations between the various statistics. High correlation between certain statistics, such as between the mean and the standard deviation, usually indicates a highly skewed distribution.

You can also use the Bootstrap tool to analyze the distribution of percentiles, but you should run at least 1,000 bootstrap samples and 1,000 trials per sample to obtain good sampling distributions for these statistics (according to Efron and Tibshirani; see Bibliography).

Decision Table Tool

Decision variables are values that you can control, such as how much to charge for a product or how many wells to drill. But, in situations with uncertainty, it is not always obvious what effect changing a decision variable can have on the forecast results.

The Decision Table tool runs multiple simulations to test different values for one or two decision variables. The tool tests values across the range of the decision variables and puts the results in a table that you can analyze using Crystal Ball forecast, trend, or overlay charts.

The Decision Table tool is useful for investigating how changes in the values of a few decision variable affect the forecast results. For models that contain more than a handful of decision variables, or where you are trying to optimize the forecast results, use OptQuest for Crystal Ball, if available.

OptQuest is a wizard-based program that enhances Crystal Ball by automatically finding optimal solutions to simulation models. This program is an option included with Crystal Ball Decision Optimizer.

Table 9 Comparison of Decision Table Tool and OptQuest

Aspect Compared	Decision Table tool	OptQuest		
Process	Runs multiple Crystal Ball simulations for different values of decision variables	Same as Decision Table tool.		
Results	All displayed in a table	Only the best solutions displayed		
Optimization	No	Yes		

Aspect Compared	Decision Table tool	OptQuest	
Number of variables	One or two	Unlimited	
Variable range	Small	Small to large	

The following sections describe and show how to use the Decision Table tool:

- "Using the Decision Table Tool" on page 204
- "Decision Table Example" on page 206

Using the Decision Table Tool

To use the Decision Table tool, start the tool, select a target forecast, select decision variables, set options, and run the tool. The following sections describe these procedures:

- "Starting the Decision Table Tool" on page 204
- "Using the Decision Table Welcome Panel" on page 204
- "Specifying a Target Forecast for Decision Table Analysis" on page 205
- "Selecting Decision Variables for Decision Table Testing" on page 205
- "Setting Decision Table Tool Options" on page 205
- "Running the Decision Table Tool" on page 206

Also see "Decision Table Example" on page 206.

Starting the Decision Table Tool

To start the Decision Table tool, choose **Run** then, **More Tools** then, **Decision Table**. (In Microsoft Excel 2007 or later, select **More Tools** in the **Tools** group.)

The first time you start the Decision Table tool, the Welcome panel opens. Otherwise, the Target Forecast panel opens.

Using the Decision Table Welcome Panel

The Welcome panel opens the first time you use the Decision Table tool. It describes the tool and its use. The controls for this panel include:

- Next Opens the Target Forecast panel for choosing a target forecast.
- Run Runs the Decision Table tool and generates results.
- To continue with the Decision Table tool and define a target forecast, click Next.

The Target Forecast panel opens.

Specifying a Target Forecast for Decision Table Analysis

The Target Forecast panel indicates which forecast to analyze. The controls in this panel include:

- Forecast list Lists all forecast cells in all open spreadsheets. The first forecast is selected
 by default.
- Back Opens the Welcome panel.
- Next Opens the Decision Variables panel.
- Run Runs the Decision Table tool and generates results.
- To open the Decision Variables panel and select which decision variables to test, click Next.

The Decision Variables panel opens.

Selecting Decision Variables for Decision Table Testing

This panel specifies one or two decision variables to test. The controls in this panel are:

- Available Decision Variables Lists all the defined decision variables in the open spreadsheets.
- Chosen Decision Variables Lists one or two decision variables that the tool will test with different values.
- Moves the selected decision variable in the Available Decision Variables list to the Chosen Decision Variables list.
- Moves the selected decision variable in the Chosen Decision Variables list to the Available Decision Variables list.
- Back Returns to the Target Forecast panel.
- Next Opens the Options panel.
- Run Runs the Decision Table tool and generates results.
- To set Decision Table options, click Next

The Options panel opens.

Setting Decision Table Tool Options

The Options panel sets options that control the tool. You can set two kinds of options:

- "Simulation Control Options" on page 206
- "While Running Options" on page 206

Other controls include:

- Back Returns to the Decision Variables panel.
- Run Runs the Decision Table tool.

➤ When options settings are complete, click Run to run the tool.

Simulation Control Options

The Simulation Control options in this panel are:

- Number of test values for each decision variable Sets the number of values the tool will test for each selected decision variable. The tool distributes the number of values evenly across the defined decision variable range. If you have one decision variable, the tool runs a simulation for each test value. For two decision variables, the tool runs a simulation for each combination of values, that is, the product of the two numbers of test values.
- Trials per simulation (maximum) Sets the maximum number of trials to run for each simulation. The default is the number set in the Crystal Ball Run Preferences dialog.

While Running Options

The While Running options include:

- Show Forecasts As Defined Displays a forecast chart for each defined forecast during the simulation.
- Show Only Target Forecast Displays only the forecast chart for the target forecast during the simulation.
- Hide All Forecasts Displays no forecast charts during the simulation.

Running the Decision Table Tool

To run the Decision Table tool, click Run when all settings are complete.

Results are generated as described in "Interpreting Decision Table Results" on page 208.

Decision Table Example

In the Crystal Ball Examples folder, there is an Oil Field Development.xls workbook you can use to experiment with the Decision Table tool. This spreadsheet model predicts how best to develop a new oil field by selecting the optimal number of wells to drill, rate of oil production, and size of the refinery to build that will maximize the net present value of the field.

The Decision Table example description contains the following sections:

- "Generating Decision Table Results" on page 207
- "Interpreting Decision Table Results" on page 208

Generating Decision Table Results

- To run the Decision Table tool:
- 1 In Microsoft Excel with Crystal Ball loaded, open the workbook Oil Field Development.xls.
- 2 In the Run, then Run Preferences, then Sampling dialog:
 - Set Random Number Generation to use the Same Sequence Of Random Numbers and an Initial Seed Value of 999.
 - Select Monte Carlo simulation.

When using this tool, use these options to make the resulting simulations comparable.

- 3 In the Run Preferences dialog, click **OK**.
- 4 Select Run, then More Tools, then Decision Table. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

If this is the first time you have opened the Decision Table tool, the Welcome panel opens.

5 If the Welcome panel opens, click **Next**.

The Target Forecast panel opens.


6 Be sure the NPV forecast is selected and click Next.

The Decision Variables panel opens.

- 7 Move Facility Size and Wells To Drill to the Chosen Decision Variables list.
 - a. Press the Ctrl key and select Facility Size and Wells To Drill in the Available Decision Variables list.
 - b. Click >>>
- 8 Click Next.

The Options panel opens, similar to Figure 75.

Figure 75 Options panel, Decision Table Tool



Set the following options as shown in Figure 75:

- Number of values to test for Facility Size = 7.
- Number of values to test for Wells To Drill = 6.
- Maximum trials per simulation = 500.
- Show Only Target Forecast is selected.

10 Click Run.

The tool runs a simulation for each combination of decision variable values. It compiles the results in a table of forecast cells indexed by the decision variables.

If distribution fitting is turned on for the target forecast, it is turned off during the running of simulations within the tool. When simulations finish, distribution fitting is restored.

Interpreting Decision Table Results


For this example, the Decision Table tool ran 42 simulations, one for each combination of wells to drill and facility sizes. The simulation that resulted in the best mean NPV was the combination of 12 wells and a facility size of 150 mbd (Figure 76).

Figure 76 Decision table for Oil Field Development Results

	Α	В	С	D	Е	F	G	Н	
	Trend Chart	Facility	Facility	Facility	Facility	Facility	Facility	Facility	
	Overlay Chart	lity size	ty size						
	Forecast Chart	ze (50.00)	(100.00)	(150.00)	(200.00)	(250.00)	(300.00)	(350.00)	
1		8	8	8	ĕ	8	8	8	
2	Wells to drill (2)	57.67	-2.33	-52.33	-92.33	-122.33	-142.33	-152.33	1
3	Wells to drill (12)	152.22	296.31	308.35	273.93	243.93	223.93	213.93	2
4	Wells to drill (21)	60.01	222.04	291.94	289.13	263.11	243.17	233.17	3
5	Wells to drill (31)	-43.63	119.87	196.08	204.23	180.93	161.20	151.20	4
6	Wells to drill (40)	-136.92	26.98	103.75	113.09	90.69	71.04	61.04	5
7	Wells to drill (50)	-240.57	-76.64	0.65	10.04	-12.34	-32.00	-42.00	6
8		1	2	3	4	5	6	7	

To view one or more of the forecasts in the decision table, select the cells and click Forecast Chart. To compare one or more forecasts on the same chart, select the cells and click the Trend Chart or Overlay Chart button in column A (Figure 77).

Figure 77 Trend chart of 150 mbd Forecasts



You can create the trend chart shown in Figure 77 by selecting all the forecast cells in the Facility Size (150.00) column of the results table and clicking on Trend Chart. This chart shows that the forecast with the highest mean NPV also has the largest uncertainty compared to other forecasts with smaller NPVs of the same facility size. This indicates a higher risk that you could avoid with a different number of wells (although the lower risk is accompanied with a lower NPV).

Note: If you have the Probability Above A Value option selected in the Run Preferences, then Options panel, the percentiles will be reversed in meaning, so that the 1st percentile represents the uppermost 1% and the 99th percentile represents the lowest 1%. For more information on this reversal, "Setting Statistics Preferences" on page 82.

Scenario Analysis Tool

The Scenario Analysis tool runs a simulation and then sorts and matches all the resulting values of a target forecast with their corresponding assumption values. Then, you can investigate which combination of assumption values gives a particular result.

You can run Scenario Analysis on any Crystal Ball model with at least one assumption and forecast that are not frozen. You select a target forecast to analyze and then the forecast's percentile or value range to examine. The resulting table shows all the values for the target forecast in the designated range, sorted, along with the corresponding assumption values. For instructions, see "Using the Scenario Analysis Tool" on page 210.

Using the Scenario Analysis Tool

To use Scenario Analysis, start the tool, specify a target forecast, specify options, and run the tool. The following sections describe these procedures:

- "Starting Scenario Analysis" on page 210
- "Specifying a Scenario Analysis Target Forecast" on page 210
- "Specifying Scenario Analysis Options" on page 211
- "Running the Scenario Analysis Tool" on page 211

Also see "Scenario Analysis Example" on page 211.

Tip: To help ensure the most accurate Scenario Analysis results, select Run, then Run Preferences, and then Trials and be sure Stop On Calculation Errors is selected before you use Scenario Analysis.

Starting Scenario Analysis

- ➤ To start Scenario Analysis:
- Select Run, then More Tools, and then Scenario Analysis. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

The first time you start Scenario Analysis, the Welcome panel opens. Otherwise, the Target Forecast panel is displayed.

If Welcome is displayed, click Next to display Target Forecast.

Specifying a Scenario Analysis Target Forecast

The Scenario Analysis tool analyzes the corresponding assumptions for a specified forecast. The Target Forecast panel indicates the forecast to use as the target.

- To specify a target forecast for analysis:
- In the Target Forecast panel, select a forecast from the list.
- Click **Next** to open the **Options** panel.

Specifying Scenario Analysis Options

You can use the Options panel to:

- Specify the range of forecast values to analyze
- Indicate which forecast charts to display while running Scenario Analysis
- Set the number of trials to run.
- Include scenarios for non-target forecasts
- To specify Scenario Analysis Options:
- Display the **Options** panel.
- Review the Range of forecast results settings and indicate whether to analyze a range of percentiles or forecast values.
 - All scenarios resulting in a forecast value that falls within the specified range are displayed in the final table, along with their corresponding assumption values. For the percentile range, enter lower and upper percentiles (both numbers between 0 and 100, or 100 to 0 if Probability above a value is selected in the Run Preferences dialog). For the forecast value range, enter the lower and upper limits for the values. The default range is –Infinity to +Infinity.
- In the While running group, specify which forecasts to show while running Scenario Analysis. You can show all the defined forecasts, only the target forecast, or none of the forecasts (the fastest setting).
- Set **Simulation control** to select the maximum number of trials to run.
- Optional: Select Include scenarios for non-target forecasts to include all forecasts in the output table.

Running the Scenario Analysis Tool


To run the Scenario Analysis tool, click Run after a target forecast is selected and appropriate options are set.

Results are displayed, as described in "Interpreting Scenario Analysis Results" on page 213.

Scenario Analysis Example

Crystal Ball includes a Toxic Waste Site example workbook you can use to experiment with the Scenario Analysis tool. This model predicts cancer risk to a nearby population from a toxic waste site. This model has four assumptions and one forecast (Figure 78).

Figure 78 Toxic Waste Site Spreadsheet Model



Generating Scenario Analysis Results

- To run the Scenario Analysis tool on the Toxic Waste Site model:
- In Microsoft Excel with Crystal Ball loaded, select Help, then Crystal Ball, then Examples Guide, and then select Toxic Waste Site in the name list. (For Microsoft Excel 2007 or later, select Resources and then **Examples Guide** in the **Help** group.)
- Select Run, then More Tools, then Scenario Analysis. (In Microsoft Excel 2007 or later, select More **Tools** in the **Tools** group.)
- In the Target Forecast panel, select Risk Assessment.


Note: If Welcome opens, select Next to display Target Forecast.

- Click **Next** to open the **Options** panel.
- 5 In the Range of forecast results group, specify a percentile range from 95 to 100 percent.
- In the While running group, select Show only target forecast.
- In the **Simulation control** group, enter **1000** as the maximum number of trials to run.
- Click Run.

Scenario Analysis creates a table of all the forecast values within the range specified in step 5, along with the corresponding value of each assumption (Figure 79).

Note: Even if you have not selected Stop On Calculation Errors in the Run Preferences dialog, the Scenario Analysis tool stops on calculation errors. In that case, reported results are not fully representative and differ from those produced following a complete, error-free simulation. The original setting is restored when the Scenario Analysis tool closes.

Figure 79 Scenario Analysis Results for Toxic Waste Site



Note: For the next few steps, you might want to arrange the windows in Microsoft Excel vertically, so that you can see the results next to the original workbook.

Select the row with the 98.00% (assuming percentile display is set to the defaults, Probability below a value and 10%, 90%, etc.).

10 Click Paste Selected Scenario.

The scenario of assumption values that produced the 98th percentile value of the target forecast is displayed in the Toxic Waste workbook. Crystal Ball recalculates the workbook and updates the forecast cells.

11 Click Paste Next Scenario.

In the workbook, the assumption and forecast values change to the values for the next scenario (for 98.10%).

12 Click Reset Original Values.

The original assumption and forecast values are displayed in the workbook.

Note: If you have stochastic elements in the model other than Crystal Ball assumptions—for example a RAND() function, a random value returned by a macro, or even a Crystal Ball probability spreadsheet function such as CB.Normal(), the value for that stochastic element is not pasted into the model with the Paste buttons on the Scenario Analysis tool's output spreadsheet. True scenario analysis cannot be performed in this case; if the target forecast is a function of these other elements, the forecast values do not match.


Interpreting Scenario Analysis Results

In this example, the simulation generated 1,000 forecast values. Since you selected to analyze the percentiles from 95 to 100, the resulting table lists 51 forecast values, or the top 5% of the forecast range, including the endpoints. The table sorts the forecast values from lowest to highest value along with the assumption values that Crystal Ball generated for each trial.

If you selected Probability above a value in the Run Preferences dialog, the percentiles are reversed in meaning, so that the 1st percentile represents the uppermost 1% and the 99th percentile represents the lowest 1%.

One way to analyze the Scenario Analysis results is to identify a particular forecast value and see what assumption values created that forecast value.

Another way to analyze the Scenario Analysis results is to generate a scatter chart in Microsoft Excel with the data. For example, you might create a scatter chart comparing the risk assessment with the CPF (Figure 80).



Scatter Chart of Risk Assessment and CPF Figure 80

2D Simulation Tool

Risk analysts must often consider two sources of variation in their models:

- uncertainty Assumptions that are uncertain because you have insufficient information about a true, but unknown, value. Examples of uncertainty include the reserve size of an oil field and the prime interest rate in 12 months. You can describe an uncertainty assumption with a probability distribution. Theoretically, you can eliminate uncertainty by gathering more information. Practically, information can be missing because you haven't gathered it or because it is too costly to gather.
- variability Assumptions that change because they describe a population with different values. Examples of variability include the individual body weights in a population or the daily number of products sold over a year. You can describe a variability assumption with a discrete distribution (or approximate one with a continuous distribution). Variability is inherent in the system, and you cannot eliminate it by gathering more information.

For many types of risk assessments, it is important to clearly distinguish between uncertainty and variability (see Hoffman and Hammonds reference in the Bibliography). Separating these concepts in a simulation lets you more accurately detect the variation in a forecast due to lack of knowledge and the variation caused by natural variability in a measurement or population. In the same way that a one-dimensional simulation is generally better than single-point estimates for showing the true probability of risk, a two-dimensional simulation is generally better than a one-dimensional simulation for characterizing risk.

The 2D Simulation tool runs an outer loop to simulate the uncertainty values, and then freezes the uncertainty values while it runs an inner loop (of the whole model) to simulate the variability. This process repeats for some number of outer simulations, providing a portrait of how the forecast distribution varies due to the uncertainty.

The primary output of this process is a chart depicting a series of cumulative frequency distributions. You can interpret this chart as the range of possible risk curves associated with a population.

Note: When using this tool, set the Seed Value option in the Crystal Ball Run Preferences dialog so that the resulting simulations are more comparable.

The following sections describe and show how to use the 2D Simulation tool and its results:

- "Using the 2D Simulation Tool" on page 215
- "2D Simulation Example" on page 217
- "Second-Order Assumptions" on page 222

Using the 2D Simulation Tool

To use the 2D Simulation tool, start the tool, select a forecast to analyze, sort assumptions into Uncertainty and Variability groups, set options, and run the tool. The following sections describe these procedures:

- "Starting the 2D Simulation Tool" on page 216
- "Using the 2D Simulation Welcome Panel" on page 216
- "Specifying a 2D Simulation Target Forecast" on page 216
- "Sorting Assumptions for 2D Simulation Analysis" on page 216
- "Setting 2D Simulation Options" on page 217
- "Running the 2D Simulation Tool" on page 217

Also see "2D Simulation Example" on page 217.

Starting the 2D Simulation Tool

To start the 2D Simulation tool, select Run, then More Tools, then 2D Simulation. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

The first time you start 2D Simulation, the Welcome panel opens. Otherwise, the Target Forecast panel opens.

Using the 2D Simulation Welcome Panel

The Welcome panel opens the first time you use the 2D Simulation tool. It describes the tool and its use. The controls for this panel include:

- Next Opens the Target Forecast panel for specifying the forecast to analyze.
- Run Runs the 2D Simulation tool if all required settings are complete.
- If the Welcome panel opens, click Next to open the Target Forecast panel.

Specifying a 2D Simulation Target Forecast

The Target Forecast panel of the 2D Simulation tool indicates which forecast to analyze. The controls for this panel include:

- Forecast list Lists all the forecast cells in all open spreadsheets. The first forecast is selected
 by default.
- Back Opens the Welcome panel.
- Next Opens the Assumption Types panel.
- Run Runs the 2D Simulation tool if all required settings are complete.
- When settings are complete on the Target Forecast panel, click Next to open the Assumption Types panel.

Sorting Assumptions for 2D Simulation Analysis

The Assumption Types panel of the 2D Simulation tool separates the assumptions into uncertainty assumptions and variability assumptions. All the assumptions from all open worksheets start in the Uncertainty list by default. You must have at least one assumption of each type. When you save the spreadsheet model, the tool remembers the assumption types for the next time you run the tool. The controls for this panel include:

- >> Moves any selected assumptions in the Uncertainty list to the Variability list.
- << Moves any selected assumptions in the Variability list to the Uncertainty list.
- Back Returns to the Target Forecast panel.
- Next Opens the Options panel.
- Run Runs the 2D Simulation tool if all required settings are complete.

When all assumptions have been sorted into the Uncertainty and Variability lists, click Next to open the Options panel.

Setting 2D Simulation Options

The Options panel of the 2D Simulation tool sets simulation control, display, and report options. This panel includes the following controls:

- Simulation Control Sets the number of trials for the outer (uncertainty) simulation and inner (variability) simulation. The default number of outer trials is 50 and the default number of inner trials is the number you have set on the Trials tab of the Run Preferences dialog.
- While Running Specifies which forecasts to show while you are running the tool. You
 can view the forecasts according to the display setting for each chart, only the target forecast,
 or none of the forecasts.
- Report Options Includes statistics, percentiles, and capability metrics in reports. You can also specify how many simulations to include in accompanying output and overlay charts.
- Back Opens the Assumption Types panel for identifying uncertainty and variability assumptions.
- Run Runs the 2D Simulation tool if all required settings are complete.
- ➤ When 2D Simulation tool settings are complete, click Run to run the tool and generate results, as shown in "Interpreting 2D Simulation Results" on page 219.

Running the 2D Simulation Tool

To run the 2D Simulation tool, confirm that all required settings are complete and then click Run. Results are generated, as described in "Interpreting 2D Simulation Results" on page 219.

2D Simulation Example

In the Crystal Ball Examples folder there is a Toxic Waste Site.xls workbook you can use to experiment with the 2D Simulation tool. This spreadsheet model predicts the cancer risk to the population from a toxic waste site. This spreadsheet has two variability assumptions and two uncertainty assumptions.

The 2D Simulation tool example contains the following sections:

- "Interpreting 2D Simulation Results" on page 219
- "Interpreting 2D Simulation Results" on page 219

Generating 2D Simulation Results

- To run the 2D Simulation tool:
- 1 In Microsoft Excel with Crystal Ball loaded, open the workbook Toxic Waste Site.xls.
- 2 In the Run, then Run Preferences, then Sampling dialog:
 - Set Random Number Generation to use the Same Sequence Of Random Numbers and An Initial Seed Value of 999.
 - Select Monte Carlo simulation.

When using this tool, use these options to make the resulting simulations comparable.

3 Select Run, then More Tools, then 2D Simulation. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

If this is the first time you have opened the 2D Simulation tool, the Welcome panel opens.

4 If the Welcome panel opens, click Next.

The Target Forecast panel opens.

- 5 Be sure the Risk Assessment forecast is selected.
- 6 Click Next.

The Assumption Types panel opens.

- 7 Move Body Weight and Volume Of Water Per Day to the Variability list.
 - a. Select Body Weight.
 - b. Press Ctrl and then click Volume Of Water Per Day.
 - c. Click

This separates the assumptions into the two types: Uncertainty and Variability.

8 Click Next.

The Options panel opens. The default number of trials for the outer (uncertainty) simulation is 50. The default for the inner (variability) simulation is the number of trials currently set on the Trials tab of the Run Preferences dialog.

- 9 Set the following options.
 - Run outer (uncertainty) simulation for 100 trials
 - Run inner (variability) simulation for 1,000 trials
 - Show Only Target Forecast is selected.

Use the default settings in the Report Options group.

10 Click Run.

The simulations start. The tool first single-steps one trial to generate a new set of values for the uncertainty assumptions. Then it freezes these assumptions and runs a simulation for the variability assumptions in the inner loop.

The tool retrieves the Crystal Ball forecast information after each inner loop runs. The tool then resets the simulation and repeats the process until the outer loop has run for the specified number of simulations.

Interpreting 2D Simulation Results

The results of the simulations are displayed in a table containing the forecast means, the uncertainty assumption values, and the statistics (including percentiles) of the forecast distribution for each simulation (Figure 81).




Figure 81 2D Simulation Results Table


The tool also graphs the results of the two-dimensional simulations on an overlay chart and a trend chart.

The overlay chart preferences can be set to show the risk curves for the simulations for different sets of uncertainty assumption values. To do this, set the Chart Type for each series to Line and choose Cumulative Frequency view. It is convenient to use the chart hotkeys — Ctrl+t for the chart type and Ctrl+d for view. Optionally, use Ctrl+n to move or remove the legend and Ctrl+m to cycle through central tendency marker lines.

For this example, Figure 82 shows that most of the risk curves are clustered densely toward the center while a few outlier curves are scattered toward the Cumulative Frequency axis, showing the small probability of having a much greater risk.


In risk analysis literature, the curves are often called the alternate realizations of the population risk assessment.

Figure 82 Overlay Chart of Risk Curves



The trend chart (Figure 83) shows certainty bands for the percentiles of the risk curves. The band width shows the amount of uncertainty at each percentile level for all the distributions.

Figure 83 Trend Chart of Certainty Bands



You can focus in on a particular percentile level, such as the 95th percentile, by viewing the statistics of the 95th percentile forecast, shown in Figure 84. For example, this figure shows 100 trials, the number of 95th percentiles in the forecast.

Porecast: 95% Edit View Forecast Preferences Help Split View 100 Displayed 100 Trials Number of 95% Statistic Forecast values 95th percentiles 100 Trials in the forecast 0.06 Mean 1 45F-04 Median 1.40E-04 0.05 Mode Standard Deviation 4.49E-05 Probability 0.04 2.01E-09 Frequency Variance Skewness 0.3018 Kurtosis 2.26 3 Coeff of Variability 0.3088 0.02 Minimum 7.02E-05 Maximum 2.59E-04 Mean Std. Error 4.49E-06 0.01 0.00 1.00E-04 2.00E-04

Figure 84 95th Percentile Forecast Statistics

-Infinity


If you have the Probability Above A Value option selected in the Run Preferences, then Options panel, the percentiles will be reversed in meaning, so that the 1st percentile represents the uppermost 1% and the 99th percentile represents the lowest 1%. For more information, see "Setting Statistics Preferences" on page 82.

Infinity

%

Compare the results of the two-dimensional simulation to a one-dimensional simulation (with both uncertainty and variability co-mingling together) of the same risk model, as in Figure 85, following.

The mean of the 95th percentiles in Figure 84, 1.45E-4, is lower than the 95th percentile risk of the one-dimensional simulation shown in Figure 85 at 2.06E-4. This indicates the tendency of the one-dimensional simulation results to overestimate the population risk, especially for highly skewed distributions.



Forecast Chart for One-dimensional Simulation Figure 85

Certainty: 100.00

Second-Order Assumptions

Some assumptions contain elements of both uncertainty and variability. For instance, an assumption might describe the distribution of body weights in a population, but the parameters of the distribution might be uncertain. These types of assumptions are called second-order assumptions (also, second-order random variables; see Burmaster and Wilson, 1996, in the Bibliography). You can model these types of assumptions in Crystal Ball by placing the uncertain parameters of the distribution in separate cells and defining these cells as assumptions. You then link the parameters of the variability assumption to the uncertainty assumptions using cell references.

- To illustrate this for the Toxic Waste Site.xls spreadsheet:
- Enter the values 70 and 10 into cells G4 and G5, respectively.

These are the mean and standard deviation of the Body Weight assumption in cell C4, which is defined as a normal distribution.

- Define an assumption for cell G4 using a normal distribution with a mean of 70 and a standard deviation of 2.
- Define an assumption for cell G5 using a normal distribution with a mean of 10 and a standard deviation
- Enter references to these cells in the Body Weight assumption.

For an example, see Figure 86, following.




Figure 86 Assumption Using Cell References for the Mean and Standard Deviation

When you run the tool for second-order assumptions, the uncertainty of the assumptions' parameters is modeled in the outer simulation, and the distribution of the assumption itself is modeled (for different sets of parameters) in the inner simulation.

Note: Often, the parameters of assumptions are correlated. For example, you would correlate a higher mean with a higher standard deviation or a lower mean with a lower standard deviation. Defining correlation coefficients between parameter

distributions can increase the accuracy of the two-dimensional simulation. With data available, as in sample body weights of a population, you can use the Bootstrap tool to estimate the sampling distributions of the parameters and the correlations between them.

Data Analysis Tool

The Data Analysis tool imports and analyzes data in Crystal Ball. The data is imported directly into Crystal Ball forecasts, one for each data series. Then, you can analyze it using any of Crystal Ball's features.

To use the Data Analysis tool, the data series must be contiguous (in adjacent rows or columns) in either rows or columns.

The following sections describe and show how to use the Data Analysis tool:

- "Using the Data Analysis Tool" on page 223
- "Data Analysis Example" on page 225

Using the Data Analysis Tool

The following sections describe procedures for starting and using the Data Analysis tool:

- "Starting the Data Analysis tool" on page 223
- "Specifying Data Analysis Input Data" on page 224
- "Setting Data Analysis Options" on page 224
- "Running the Data Analysis Tool" on page 225

Starting the Data Analysis tool

To start the Data Analysis tool, choose Run, then More Tools, then Data Analysis. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

The first time you start the Data Analysis tool, the Welcome panel opens. Otherwise, the Input Data panel opens.

Using the Data Analysis Welcome Panel

The Welcome panel opens the first time you use the Data Analysis tool. It describes the tool and its use. The controls for this panel include:

- Next Opens the Input Data panel for specifying the location of the data series.
- Run Runs the Data Analysis tool.

➤ If the Welcome panel opens, click Next to open the Input Data panel.

Specifying Data Analysis Input Data

The Input Data panel of the Data Analysis tool indicates the location of the data that to analyze. You can also set input-related options. The Data Analysis data selector chooses possible data for fitting. This information is displayed in the Location Of Data Series field and the illustration. You can choose different data if necessary. The controls in this panel include:

- Location Of Data Series Indicates the cells that contain data to analyze. If the data has
 headers or labels at the beginning of the rows or columns of data, include them in the
 selection and select the appropriate Headers setting(s). The data must be in adjacent rows
 or columns.
- Orientation Sets whether the data is in rows or columns: Data In Rows indicates that data
 is in horizontal rows; Data In Columns indicates that data is in vertical columns.
- Headers Indicates whether the data has headers and/or labels and whether they are located in the first row or the first column (varies with orientation). Checked items are used in the output: Top Row Has Headers/Labels includes text in the top (first) row in the selection; Left Column Has Labels/Headers includes text in the left (first) column in the selection.
- Back Returns to the Welcome panel.
- Next Opens the Options panel.
- Run Runs the Data Analysis tool, automatically generating forecasts for each selected data series.
- Click Next to open the Options panel and set Data Analysis options.

Setting Data Analysis Options

The Options panel of the Data Analysis tool sets a variety of preferences for displaying generated forecast charts, fitting probability distributions to generated forecast data, generating correlations between forecast data series, and running simulations on opened models. The Options panel includes these controls:

- Automatically Open Forecast Charts When checked, automatically opens forecast charts when the Data Analysis tool runs.
- View Indicates which forecast chart view to use, similar to the forecast chart View menu commands.
- Split View When checked, shows a chart in the first pane and statistics in the second pane.
- Fit A Probability Distribution To The Data When checked, calculates and plots a curve for the distribution that best fits the data in each series. Select Fit Options to review or change current settings in the Fit Options panel.
- Generate Correlation Matrix Between Data Series When checked, plots the rank correlation between pairs of forecasts. You can click the Scatter Chart button in the results

- worksheet to show the forecast relationships graphically, along with lines of fit and correlation coefficients.
- Run Simulation On Opened Models (To Compare Data Simulation Results) When checked, this setting can be used for validating models. A simulation is run on all open workbooks at the same time the selected data is analyzed. In this case, forecast charts for all opened models are displayed with those generated from the data selected for analysis.
- Back Returns to the Input Data panel.
- Run Runs the tool, automatically generating forecasts for each selected data series.
- ➤ When all settings are complete, click Run to perform the Data Analysis import and generate forecasts.

Running the Data Analysis Tool

To run the Data Analysis tool, confirm that all required settings are complete and then click Run.

For more information, see "Generating Data Analysis Results" on page 226.

Data Analysis Example

You can use the Magazine Sales.xls workbook, provided in the Crystal Ball Examples folder, to experiment with the Batch Fit tool. This model shows the estimated gross profit resulting from newsstand sales of four of the company's most popular magazines (Figure 87).




Figure 87 Magazine Sales workbook

This model shows sales figures for four magazines. The accompanying Sales Data worksheet contains historical data for each magazine. For this example, you will import the data into the Data Analysis tool, automatically create a forecast for each magazine, run a simulation, display the simulated data as forecast charts, correlate the forecasts for each magazine, and produce other charts using buttons in the new DataAnalysisOutput worksheet generated by the tool.

Generating Data Analysis Results

- To run the Data Analysis tool:
- 1 In Microsoft Excel with Crystal Ball loaded, open the workbook Magazine Sales.xls.
- 2 Click the Sales Data tab to select and display it.
- 3 Choose Run, then More Tools, then Data Analysis. (In Microsoft Excel 2007 or later, select More Tools in the Tools group.)

If this is the first time you have opened the Data Analysis tool, the Welcome panel opens.

4 If the Welcome panel opens, click Next.


The Input Data panel opens.

The Data Analysis data selector attempts to select relevant data and indicate its orientation.

The Input Data panel shows a cell range with four data series and 360 rows of data. The tool determined that the data is in columns with headers in the first row. There are no labels in the first column.

- 5 Examine the Sales Data tab of the Magazine Sales workbook to confirm that this information is correct.
- 6 Click Next to display the Options panel, similar to Figure 88.

Figure 88 Options panel, Data Analysis Tool



- 7 Confirm that the settings are as shown in Figure 88. Change them if necessary:
 - Check "Automatically open forecast charts."
 - Set View to Frequency.
 - Check "Fit a probability distribution to the data."
 - Check "Generate correlation matrix between data series."
 - Check "Run simulation on opened models..."
- Click Fit Options to specify which distributions and ranking statistics to use for fitting.

The Fit Options panel opens.

9 Choose AutoSelect for both Distributions To Fit and Rank By Goodness-of-Fit Statistic.

Note: If you know location, shape, or other parameter values that might help create a more accurate fit with certain distributions, check Lock Parameters and enter appropriate values in the Lock Parameters dialog. For details, see "Locking Parameters When Fitting Distributions" on page 49.

- 10 Click OK to return to the Options panel.
- 11 Click Run.

When the Data Analysis tool runs, it creates:

- A series of forecast charts.
- A new workbook with data and buttons on a worksheet named DataAnalysisOutput.
 The new worksheet is similar to Figure 89.




Figure 89 Data Analysis Output worksheet

- Cells B2 through E2 contain forecasts, one for each series of magazine data.
- Below that is a correlation matrix, showing the relationship of each forecast to the other three.
- Cell A1 contains four buttons you can use to display forecast, trend, overlay, and scatter charts.
- 12 Select all four forecast cells in row 2 and click the Forecast Chart button.

A chart is displayed for each forecast, in Frequency view. You can customize them as usual, changing the view or setting other chart preferences as described in Chapter 6.

To see which type of distribution fit best, select a chart and choose View, then Goodness of Fit.

13 Experiment with the other chart buttons as well. Select two or more forecasts and click a button. You can create new charts of the selected type just as if you used the Analyze command or clicked a chart icon.



Selecting and Using Probability Distributions

In This Appendix

Introduction	229
Understanding Probability Distributions	229
Selecting Probability Distributions	233
Probability Distribution Descriptions	235
Using the Custom Distribution	270
Truncating Distributions	285
Comparing Distributions	286
Distribution Parameter Summary	288
Using Probability Functions	290
Sequential Sampling with Custom Distributions	291

Introduction

This appendix explains probability and probability distributions. Understanding these concepts will help you select the most appropriate probability distribution for the spreadsheet model. This section describes in detail the distribution types Crystal Ball uses and demonstrates their use with real-world examples.

Understanding Probability Distributions

For each uncertain variable in a simulation, you define the possible values with a *probability distribution*. The type of distribution you select depends on the conditions surrounding the variable. For example, some common distribution types are shown in Figure 90: normal, triangular, uniform, and lognormal

Figure 90 Common Distribution Types









During a simulation, the value to use for each variable is selected randomly from the defined possibilities.

A simulation calculates numerous scenarios of a model by repeatedly picking values from the probability distribution for the uncertain variables and using those values for the cell. Commonly, a Crystal Ball simulation calculates hundreds or thousands of scenarios in just a few seconds. The following section, "A Probability Example" on page 230, shows how a probability distribution relates to a simple set of employment data.

Crystal Ball works with two types of distributions, described in "Continuous and Discrete Probability Distributions" on page 232. For suggestions about using the best distribution when defining an assumption, see "Selecting Probability Distributions" on page 233. "Probability Distribution Descriptions" on page 235 describes the properties and uses of each distribution available in Crystal Ball.

A Probability Example

To begin to understand probability, consider this example: You want to look at the distribution of non-exempt wages within one department of a large company. First, you gather raw data, in this case the wages of each non-exempt employee in the department. Second, you organize the data into a meaningful format and plot the data as a frequency distribution on a chart. To create a frequency distribution, you divide the wages into groups (also called intervals or bins) and list these intervals on the chart's horizontal axis. Then you list the number or frequency of employees in each interval on the chart's vertical axis. Now you can easily see the distribution of non-exempt wages within the department.

A glance at the chart illustrated in Figure 91 reveals that the most common wage range is \$12.00 to \$15.00.

Approximately 60 employees (out of a total of 180) earn from \$12 to \$15.00 per hour.





Figure 91 Raw frequency data for a probability distribution

You can chart this data as a probability distribution. A probability distribution shows the number of employees in each interval as a fraction of the total number of employees. To create a probability distribution, you divide the number of employees in each interval by the total number of employees and list the results on the chart's vertical axis.


The chart illustrated in Figure 92 shows you the number of employees in each wage group as a fraction of all employees; you can estimate the likelihood or probability that an employee drawn at random from the whole group earns a wage within a given interval. For example, assuming the same conditions exist at the time the sample was taken, the probability is 0.33 (a 1 in 3 chance) that an employee drawn at random from the whole group earns between \$12 and \$15 an hour.

Figure 92 Probability distribution of wages



Compare the probability distribution in the previous example to the probability distributions in Crystal Ball (Figure 93).

Figure 93 Distribution Gallery Dialog



The probability distribution in the example in Figure 92 has a shape similar to many of the distributions in the Distribution Gallery. This process of plotting data as a frequency distribution and converting it to a probability distribution provides one starting point for selecting a Crystal Ball distribution. Select the distributions in the gallery that is displayed similar to the probability distribution, then read about those distributions in this chapter to find the correct distribution.

For information about the similarities between distributions, see "Comparing Distributions" on page 286. For a complete discussion of probability distributions, refer to the sources listed in the bibliography.

Continuous and Discrete Probability Distributions

Notice that the Distribution Gallery shows whether the probability distributions are continuous or discrete.

Continuous probability distributions, such as the normal distribution, describe values over a range or scale and are shown as solid figures in the Distribution Gallery. Continuous distributions are actually mathematical abstractions because they assume the existence of every possible intermediate value between two numbers. That is, a continuous distribution assumes there is an infinite number of values between any two points in the distribution.

Discrete probability distributions describe distinct values, usually integers, with no intermediate values and are shown as a series of vertical columns, such as the binomial distribution at the bottom of Figure 93. A discrete distribution, for example, might describe the number of heads in four flips of a coin as 0, 1, 2, 3, or 4.

However, in many situations, you can effectively use a continuous distribution to approximate a discrete distribution even though the continuous model does not necessarily describe the situation exactly.

In the dialogs for the discrete distributions, Crystal Ball displays the values of the variable on the horizontal axis and the associated probabilities on the vertical axis. For the continuous distributions, Crystal Ball does not display values on the vertical axis since, in this case, probability can only be associated with areas under the curve and not with single values.

Initially, the precision and format of the displayed numbers in the probability and frequency distributions come from the cell itself. To change the format, see "Customizing Chart Axes and Axis Labels" on page 123.

The following sections list continuous and discrete distributions available in Crystal Ball:

- "Continuous Probability Distributions" on page 232
- "Discrete Probability Distributions" on page 233

Note: Custom distributions can be defined as continuous, discrete, or both. For more information, see "Custom Distribution" on page 241.

Continuous Probability Distributions

The following sections describe continuous distributions available in Crystal Ball:

- "Beta Distribution" on page 235
- "BetaPERT Distribution" on page 237
- "Exponential Distribution" on page 243
- "Gamma Distribution" on page 245
- "Logistic Distribution" on page 251
- "Lognormal Distribution" on page 252
- "Maximum Extreme Distribution" on page 254
- "Minimum Extreme Distribution" on page 255
- "Normal Distribution" on page 258
- "Pareto Distribution" on page 260
- "Student's t Distribution" on page 263
- "Triangular Distribution" on page 264
- "Uniform Distribution" on page 266
- "Weibull Distribution" on page 267

Discrete Probability Distributions

The following sections describe discrete distributions available in Crystal Ball:

- "Binomial Distribution" on page 239
- "Discrete Uniform Distribution" on page 241
- "Geometric Distribution" on page 247
- "Hypergeometric Distribution" on page 249
- "Negative Binomial Distribution" on page 256
- "Poisson Distribution" on page 261
- "Yes-No Distribution" on page 269
- "Triangular Distribution" on page 264
- "Uniform Distribution" on page 266
- "Weibull Distribution" on page 267

Selecting Probability Distributions

Plotting data is one guide to selecting a probability distribution. The following steps provide another process for selecting probability distributions that best describe the uncertain variables in the spreadsheets.

To select the correct probability distribution:

1. Look at the variable in question. List everything you know about the conditions surrounding this variable.

You might be able to gather valuable information about the uncertain variable from historical data. If historical data are not available, use the own judgment, based on experience, to list everything you know about the uncertain variable.

For example, look at the variable "patients cured" that is discussed in "Tutorial 2 — Vision Research" on page 315. The company plans to test 100 patients. You know that the patients will either be cured or not cured. And, you know that the drug has shown a cure rate of around 0.25 (25%). These facts are the conditions surrounding the variable.

Review the descriptions of the probability distributions.

"Probability Distribution Descriptions" on page 235 describes each distribution in detail, outlining the conditions underlying the distribution and providing real-world examples of each distribution type. As you review the descriptions, look for a distribution that features the conditions you have listed for this variable.

Select the distribution that characterizes this variable.

A distribution characterizes a variable when the conditions of the distribution match those of the variable.

The conditions of the variable describe the values for the parameters of the distribution in Crystal Ball. Each distribution type has its own set of parameters, which are explained in the following descriptions.

For example, look at the conditions of the binomial distribution, as described in "Binomial Distribution" on page 239:

- For each trial, only two outcomes are possible: success or failure.
- The trials are independent. What happens on the first trial does not affect the second trial, and so on.
- The probability of success remains the same from trial to trial.
 - Now check the "patients cured" variable in Tutorial 2 in the Crystal Ball Getting Started Guide against the conditions of the binomial distribution:
- There are two possible outcomes: the patient is either cured or not cured.
- The trials (100) are independent of each other. What happens to the first patient does not affect the second patient.
- The probability of curing a patient 0.25 (25%) remains the same each time a patient is tested.
 - Since the conditions of the variable match the conditions of the binomial distribution, the binomial distribution would be the correct distribution type for the variable in question.
- 4. If historical data are available, use distribution fitting to select the distribution that best describes the data.

Crystal Ball can automatically select the probability distribution that most closely approximates the data's distribution. The feature is described in detail in "Fitting Distributions to Data" on page 46. You can also populate a custom distribution with the historical data.

After you select a distribution type, determine the parameter values for the distribution. Each distribution type has its own set of parameters. For example, there are two parameters for the binomial distribution: trials and probability. The conditions of a variable contain the values for the parameters. In the example used, the conditions show 100 trials and 0.25 (25%) probability of success.

In addition to the standard parameter set, each continuous distribution (except uniform) also lets you choose from alternate parameter sets, which substitute percentiles for one or more of the standard parameters. For more information on alternate parameters, see "Alternate Parameter Sets" on page 45.

Probability Distribution Descriptions

This section contains descriptions of all the Crystal Ball probability distributions, listed in alphabetical order.

The following sections list continuous and discrete distributions described below:

- "Continuous Probability Distributions" on page 232
- "Discrete Probability Distributions" on page 233

See "Custom Distribution" on page 241 for a description of the Custom distribution, which can be continuous, discrete, or both

As you work with the Crystal Ball probability distributions, you can use the Parameters menu found in the distribution menubar to specify different combinations of parameters. For more information, see "Alternate Parameter Sets" on page 45.

Beta Distribution



The beta distribution is continuous. It is commonly used to represent variability over a fixed range. It can represent uncertainty in the probability of occurrence of an event. It is also used to describe empirical data and predict the random behavior of percentages and fractions and can be used to represent the reliability of a company's devices.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Beta Distribution Parameters" on page 236
- "Beta Distribution Conditions" on page 236
- "Beta Distribution Example" on page 236

Beta Distribution Parameters

Minimum, Maximum, Alpha, Beta

Beta Distribution Conditions

The beta distribution is used under these conditions:

- Minimum and maximum range is between 0 and a positive value.
- Shape can be specified with two positive values, alpha and beta. If the parameters are equal, the distribution is symmetrical. If either parameter is 1 and the other parameter is greater than 1, the distribution is J-shaped. If alpha is less than beta, the distribution is said to be positively skewed (most of the values are near the minimum value). If alpha is greater than beta, the distribution is negatively skewed (most of the values are near the maximum value). Because the beta distribution is complex, the methods for determining the parameters of the distribution are beyond the scope of this manual. For more information about the beta distribution and Bayesian statistics, refer to the texts in the Bibliography.

Beta Distribution Example

A company that manufactures electrical devices for custom orders wants to model the reliability of devices it produces. After analyzing the empirical data, the company knows that it can use the beta distribution to describe the reliability of the devices if the parameters are alpha = 10 and beta = 2.


The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the beta distribution:

- The reliability rate is a random value somewhere between 0 and 1.
- The shape of the distribution can be specified using two positive values: 10 and 2.

These conditions match those of the beta distribution.

Figure 94 shows the beta distribution with the alpha parameter set to 10, the beta parameter set to 2, and Minimum and Maximum set to 0 and 1. The reliability rate of the devices will be x.

Figure 94 Beta distribution



Models that use beta distributions will run more slowly because of the inverse CDF and alternate parameter calculations that take place when random numbers are handled as part of beta distributions.

BetaPERT Distribution



BetaPERT

The betaPERT distribution is continuous. It describes a situation where you know the minimum, maximum, and most likely values to occur. It is useful with limited data. For example, you could describe the number of cars sold per week when past sales show the minimum, maximum, and usual number of cars sold. It is similar to the triangular distribution, described in "Triangular Distribution" on page 264, except the curve is smoothed to reduce the peak. The betaPERT distribution is often used in project management models to estimate task and project durations.

The following sections describe the parameters, conditions, and other features of this distribution:

- "BetaPERT Parameters" on page 237
- "BetaPERT Conditions" on page 238
- "BetaPERT Example" on page 238

BetaPERT Parameters

Minimum, Likeliest, Maximum

BetaPERT Conditions

The betaPERT distributions is used under these conditions:

- Minimum and maximum are fixed.
- It has a most likely value in this range, which forms a triangle with the minimum and maximum; betaPERT forms a smoothed curve on the underlying triangle.

BetaPERT Example

A project manager wants to estimate the time (number of days) required for completion of a project. From the manager's past experience, similar projects typically take 7 days to finish, but can be finished in 5 days given favorable conditions, and can take as long as 12 days if things do not happen as expected. The project manager wants to estimate the probability of finishing within 9 days.

The first step in selecting a probability distribution is matching available data with a distribution's conditions. Checking the betaPERT distribution for this project:

- The minimum number of days for completion is 5.
- The maximum number of days for completion is 12.
- The most likely number of days for completion is 7, which is between 5 and 12.

These conditions match those of the betaPERT distribution shown in Figure 95.





Figure 95 BetaPERT distribution

If that distribution is located in cell A1 and a forecast with formula =A1 is created, simulation results show there is about an 87% probability of the project completing within 9 days.

If the same forecast is calculated using a triangular distribution instead of a betaPERT, the probability of completing within 9 days is about 74% (Figure 96).

Figure 96 Project duration based on betaPERT distribution



Binomial Distribution



The binomial distribution is discrete. It describes the number of times a particular event occurs or fails to occur in a fixed number of trials, such as the number of heads in 10 flips of a coin or the number of defective items in 50 items. It can also used for Boolean logic (true/false or on/ off).

The following sections describe the parameters, conditions, and other features of this distribution:

- "Binomial Parameters" on page 239
- "Binomial Conditions" on page 239
- "Binomial Example 1" on page 240
- "Binomial Example 2" on page 240

Binomial Parameters

Probability, Trials

Binomial Conditions

The binomial distribution is used under these conditions:

- For each trial, only two outcomes are possible, such as success or failure.
- Trials are independent. Probability is the same from trial to trial.
- The Yes-No distribution is equivalent to the Binomial distribution with one trial.

Binomial Example 1

You want to describe the number of defective items in a total of 50 manufactured items, 7% of which (on the average) were found to be defective during preliminary testing.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the binomial distribution:

- There are only two possible outcomes: the manufactured item is either good or defective.
- The trials (50) are independent of one another. Any given manufactured item is either defective or not, independent of the condition of any of the others.
- The probability of a defective item (7%) is the same each time an item is tested.

These conditions match those of the binomial distribution.

The parameters for the binomial distribution are Probability and Trials. In example one, the values for these parameters are 50 (Trials) and 0.07 (7% Probability of producing defective items). You would enter these values to specify the parameters of the binomial distribution in Crystal Ball.

The distribution illustrated in Figure 97 shows the probability of producing x number of defective items.




Figure 97 **Binomial Distribution**

Binomial Example 2

A company's sales manager wants to describe the number of people who prefer the company's product. The manager conducted a survey of 100 consumers and determined that 60% prefer the company's product over the competitor's.

Again, the conditions fit the binomial distribution with two important values: 100 (trials) and 0.6 (60% probability of success). These values specify the parameters of the binomial distribution in Crystal Ball. The result would be a distribution of the probability that x number of people prefer the company's product.

Custom Distribution



You can use the custom distribution in Crystal Ball to represent a unique situation that cannot be described using other distribution types. It can describe a series of single values, discrete ranges, or continuous ranges.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Custom Parameters" on page 241
- "Custom Conditions" on page 241

Also see "Using the Custom Distribution" on page 270.

Custom Parameters

Variable, see "Using the Custom Distribution" on page 270.

Custom Conditions

The custom distribution is used under these conditions:

- It is a flexible distribution, used to represent a situation you cannot describe with other distribution types.
- It can be continuous, discrete, or a combination of both and can be used to input an entire set of data points from a range of cells.

For an example of a custom distribution, see the ClearView tutorial ("Growth Rate Assumption: Custom Distribution" on page 321). Also see "Using the Custom Distribution" on page 270.

Discrete Uniform Distribution



In the discrete uniform distribution, you know the minimum and maximum values and you know that all non-continuous values between the minimum and maximum are equally likely to occur. It can be used to describe a real estate appraisal or a leak in a pipe. It is the discrete equivalent of the continuous uniform distribution ("Uniform Distribution" on page 266).

The following sections describe the parameters, conditions, and other features of this distribution:

- "Discrete Uniform Parameters" on page 242
- "Discrete Uniform Conditions" on page 242

"Discrete Uniform Example" on page 242

Discrete Uniform Parameters

Minimum, Maximum

Discrete Uniform Conditions

The discrete uniform distribution is used under these conditions:

- Minimum is fixed.
- Maximum is fixed.
- All values in range are equally likely to occur.
- Discrete Uniform is the discrete equivalent of the Uniform distribution.

Discrete Uniform Example


A manufacturer determines that he must receive 10% over production costs—or a minimum of \$5 per unit—to make the manufacturing effort worthwhile. He also wants to set the maximum price for the product at \$15 per unit, so that he can gain a sales advantage by offering the product for less than his nearest competitor. All values between \$5 and \$15 per unit have the same likelihood of being the actual product price, however he wants to limit the price to whole dollars.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the uniform distribution:

- The minimum value is \$5 per unit.
- The maximum value is \$15 per unit.
- All integer values between \$5 and \$15 are equally possible.

You would enter these values in Crystal Ball to produce a discrete uniform distribution showing that all whole dollar values from \$5 to \$15 occur with equal likelihood. Figure 98 illustrates this scenario.

Figure 98 Discrete Uniform Distribution



Exponential Distribution



The exponential distribution is continuous. It is widely used to describe events recurring at random points in time or space, such as the time between failures of electronic equipment, the time between arrivals at a service booth, incoming phone calls, or repairs needed on a certain stretch of highway. It is related to the Poisson distribution, which describes the number of occurrences of an event in a given interval of time or space.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Exponential Parameter" on page 243
- "Exponential Conditions" on page 243
- "Exponential Example 1" on page 244
- "Exponential Example 2" on page 244

Exponential Parameter

Rate

Exponential Conditions

The exponential distribution is used under these conditions:

- Distribution describes the time between occurrences.
- Distribution is not affected by previous events.

Exponential Example 1

A travel agency wants to describe the time between incoming telephone calls when the calls are averaging about 35 every 10 minutes. This same example was used for the Poisson distribution to describe the number of calls arriving every 10 minutes.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the exponential distribution:

- The travel agency wants to describe the time between successive telephone calls.
- The phone calls are not affected by previous history. The probability of receiving 35 calls every 10 minutes remains the same.

The conditions in this example match those of the exponential distribution.

The exponential distribution has only one parameter: rate. The conditions outlined in this example include the value for this parameter: 35 (calls) every 10 minutes or a rate of 35. Enter this value to set the parameter of the exponential distribution in Crystal Ball.

The distribution in Figure 99 shows the probability that x number of time units (10 minutes in this case) will pass between calls.




Figure 99 **Exponential Distribution**

Exponential Example 2

A car dealer needs to know the amount of time between customer drop-ins at his dealership so that he can staff the sales floor more efficiently. The car dealer knows an average of 6 customers visit the dealership every hour.

Checking the exponential distribution:

- The car dealer wants to describe the time between successive customer drop-ins.
- The probabilities of customer drop-ins remain the same from hour to hour.

These conditions fit the exponential distribution.

The resulting distribution would show the probability that x number of hours will pass between customer visits.

Gamma Distribution



The gamma distribution is continuous. It applies to a wide range of physical quantities and is related to other distributions: lognormal, exponential, Pascal, Erlang, Poisson, and chi-squared. It is used in meteorological processes to represent pollutant concentrations and precipitation quantities. The gamma distribution is also used to measure the time between the occurrence of events when the event process is not completely random. Other applications of the gamma distribution include inventory control (such as the demand for an expected number of units sold during the lead time), economics theory, and insurance risk theory.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Gamma Parameters" on page 245
- "Gamma Conditions" on page 245
- "Gamma Example 1" on page 245
- "Gamma Example 2" on page 246
- "Chi-square and Erlang Distributions" on page 247

Gamma Parameters

Location, Scale, Shape

Gamma Conditions

- Possible occurrences in any unit of measurement are not limited.
- Occurrences are independent.
- Average numbers of occurrences are constant from unit to unit.

Gamma Example 1

A computer dealership knows that the lead time for re-ordering their most popular computer system is 4 weeks. Based upon an average demand of 1 unit per day, the dealership wants to model the number of business days it will take to sell 20 systems.

Checking the conditions of the gamma distribution:

- The number of possible customers demanding to buy the computer system is unlimited.
- The decisions of customers to buy the system are independent.

The demand remains constant from week to week.

These conditions match those of the gamma distribution. (Note that in this example the dealership has made several simplifying assumptions about the conditions. In reality, the total number of computer purchasers is finite, and some might have influenced the purchasing decisions of others.)

The shape parameter is used to specify the r th occurrence of the Poisson event. In this example, you would enter 20 for the shape parameter (5 units per week times 4 weeks). The result is a distribution showing the probability that x number of business days will pass until the 20th system is sold.

Figure 100 illustrates the gamma distribution.




Figure 100 Gamma distribution

Gamma Example 2

Suppose a particular mechanical system fails after receiving exactly 5 shocks to it from an external source. The total time to system failure, defined as the random time occurrence of the 5th shock, follows a gamma distribution with a shape parameter of 5.

Some characteristics of the gamma distribution:

- When shape = 1, gamma becomes a scalable exponential distribution.
- The sum of any two gamma-distributed variables is a gamma variable.
- If you have historical data that you believe fits the conditions of a gamma distribution, computing the parameters of the distribution is easy. First, compute the mean and variance of the historical data:

Parameter	Symbol
mean =	(X)

Parameter	Symbol
variance =	(s ²)

Then compute the distribution's parameters:

Parameter	Symbol
shape =	\bar{x}^2/s^2
scale =	s^2/\bar{x}

Chi-square and Erlang Distributions

You can model two additional probability distributions, the chi-square and Erlang distributions, by adjusting the parameters entered in the Gamma Distribution dialog. To model the chi-square distributions with parameters N and S, where N = number of degrees of freedom and S = scale, set the parameters as follows:

Parameter	Symbol
shape =	$\frac{N}{2}$
scale =	2S ²

The chi-square distribution is the sum of the squares of *N* normal variates.

The Erlang distribution is identical to the gamma distribution, except the shape parameter is restricted to integer values. Mathematically, the Erlang distribution is a summation of N exponential distributions.

Geometric Distribution



The geometric distribution is discrete. It describes the number of trials until the first successful occurrence, such as the number of times you need to spin a roulette wheel before you win or how many wells to drill before you strike oil.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Geometric Parameter" on page 248
- "Geometric Conditions" on page 248
- "Geometric Example 1" on page 248

"Geometric Example 2" on page 249

Geometric Parameter

Probability

Geometric Conditions

The geometric distribution is used under these conditions:

- Number of trials is not fixed.
- Trials continue until the first success.
- Probability of success is the same from trial to trial; if there is a 10% probability, this is entered 0.10.

Geometric Example 1

If you are drilling for oil and want to describe the number of dry wells you would drill before the next producing well, you would use the geometric distribution. Assume that in the past you have hit oil about 10% of the time.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the geometric distribution:


- The number of trials (dry wells) is unlimited.
- You continue to drill wells until you hit the next producing well.
- The probability of success (10%) is the same each time you drill a well.

These conditions match those of the geometric distribution.

The geometric distribution has only one parameter: Probability. In this example, the value for this parameter is 0.10, representing the 10% probability of discovering oil. You would enter this value as the parameter of the geometric distribution in Crystal Ball.

The distribution illustrated in Figure 101 shows the probability of x number of wells drilled before the next producing well.

Figure 101 Geometric Distribution



Geometric Example 2

An insurance company wants to describe the number of claims received until a "major" claim arrives. Records show that 6% of the submitted claims are equal in dollar amount to all the other claims combined.

Again, identify and enter the parameter values for the geometric distribution in Crystal Ball. In this example, the conditions show one important value: a 0.06 (6%) probability of receiving that "major" claim. The result would be a distribution showing the probability of x number of claims occurring between "major" claims.

Hypergeometric Distribution



The hypergeometric distribution is discrete. It is similar to the binomial distribution. Both describe the number of times a particular event occurs in a fixed number of trials. However, binomial distribution trials are independent, while hypergeometric distribution trials change the success rate for each subsequent trial and are called "trials without replacement." The hypergeometric distribution can be used for sampling problems such as the chance of picking a defective part from a box (without returning parts to the box for the next trial).

The following sections describe the parameters, conditions, and other features of this distribution:

- "Hypergeometric Parameters" on page 250
- "Hypergeometric Conditions" on page 250
- "Hypergeometric Example 1" on page 250
- "Hypergeometric Example 2" on page 251

Hypergeometric Parameters

Success, Trials, Population

Hypergeometric Conditions

The hypergeometric distribution is used under these conditions:

- Total number of items (population) is fixed.
- Sample size (number of trials) is a portion of the population.
- Probability of success changes after each trial.

Hypergeometric Example 1

You want to describe the number of consumers in a fixed population who prefer Brand X. You are dealing with a total population of 40 consumers, of which 30 prefer Brand X and 10 prefer Brand Y. You survey 20 of those consumers.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the hypergeometric distribution:

- The population size (40) is fixed.
- The sample size (20 consumers) represents a portion of the population.
- Initially, 30 of 40 consumers preferred Brand X so the initial success rate is 30. This rate changes each time you question one of the 20 consumers, depending on the preference of the previous consumer.


The conditions in this example match those of the hypergeometric distribution.

Note: If you have a probability from a different-sized sample instead of a success rate, you can estimate initial success by multiplying the population size by the probability of success. In this example, the probability of success is 75% ($.75 \times 40 = 30$ and 30/40 = .75).

The three parameters of this distribution are initial Success, number of Trials, and Population size. The conditions outlined in this example contain the values for these parameters: a population Size of 40, sample size (Trials) of 20, and initial Success of 30 (30 of 40 consumers will prefer Brand X). You would enter these values as the parameters of the hypergeometric distribution in Crystal Ball.

The distribution illustrated in Figure 102 shows the probability that x number of consumers prefer Brand X.

Figure 102 Hypergeometric distribution



Hypergeometric Example 2

The U.S. Department of the Interior wants to describe the movement of wild horses in Nevada. Researchers in the department travel to a particular area in Nevada to tag 100 horses in a total population of 1,000. Six months later the researchers return to the same area to find out how many horses remained in the area. The researchers look for tagged horses in a sample of 200.

Check the data against the conditions of the hypergeometric distribution. The parameter values for the hypergeometric distribution in Crystal Ball are the population size of 1,000, sample size of 200, and an initial success rate of 100 out of 1,000 (or a probability of 10% — 0.1 — of finding tagged horses. The result would be a distribution showing the probability of observing *x* number of tagged horses.

If you used this distribution in a model created in Crystal Ball 2000.x, you might notice slight data changes when running that model in the current version of Crystal Ball. This is because some rounding might occur when converting the probability parameter used in previous releases to the success parameter used in this version of Crystal Ball.

Logistic Distribution



The logistic distribution is continuous. It is commonly used to describe growth (the size of a population expressed as a function of a time variable). It can also be used to describe chemical reactions and the course of growth for a population or individual.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Logistic Parameters" on page 252
- "Logistic Conditions" on page 252

"Calculating Logistic Parameters" on page 252

Logistic Parameters

Mean, Scale

Logistic Conditions

The logistic distribution is used under these conditions:

Conditions and parameters are complex. See: Fishman, G. Springer Series in Operations Research. NY: Springer-Verlag, 1996.

Note: The mean parameter is the average value, which for this distribution is the same as the mode, since this is a symmetrical distribution. After you select the mean parameter, you can estimate the scale parameter. The scale parameter is a number greater than 0. The larger the scale parameter, the greater the variance.

Calculating Logistic Parameters

There are two standard parameters for the logistic distribution: mean and scale. The mean parameter is the average value, which for this distribution is the same as the mode, since this is a symmetrical distribution.

After you select the mean parameter, you can estimate the scale parameter. The scale parameter is a number greater than 0. The larger the scale parameter, the greater the variance. The scale parameter's symbol is:

To calculate a more exact scale, you can estimate the variance and use the equation:

$$\alpha = \sqrt{\frac{3 \cdot \text{variance}}{\pi^2}}$$

Lognormal Distribution



The lognormal distribution is continuous. It is widely used in situations where values are positively skewed, for example, for determining stock prices, real estate prices, pay scales, and oil reservoir size.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Lognormal Parameters" on page 253
- "Lognormal Conditions" on page 253

- "Lognormal Example" on page 253
- "Lognormal Parameter Sets" on page 254

Lognormal Parameters

Location, Mean, Standard Deviation

Lognormal Conditions

The lognormal distribution is used under these conditions:

- Upper and lower limits are unlimited, but the uncertain variable cannot fall below the value of the location parameter.
- Distribution is positively skewed, with most values near lower limit.
- Natural logarithm of the distribution is a normal distribution.

Lognormal Example

The lognormal distribution can be used to model the price of a particular stock. You purchase a stock today at \$50. You expect that the stock will be worth \$70 at the end of the year. If the stock price drops at the end of the year, rather than appreciating, you know that the lowest value it can drop to is \$0.

On the other hand, the stock could end up with a price much higher than expected, thus implying no upper limit on the rate of return. In summary, the losses are limited to the original investment, but the gains are unlimited.

Using historical data, you can determine that the standard deviation of the stock's price is \$12.


Note: If you have historical data available with which to define a lognormal distribution, it is important to calculate the mean and standard deviation of the logarithms of the data and then enter these log parameters using the Parameters menu (Location, Log Mean, and Log Standard Deviation). Calculating the mean and standard deviation directly on the raw data will not give you the correct lognormal distribution. Alternatively, use the distribution fitting feature described in "Fitting Distributions to Data" on page 46.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the lognormal distribution:

- The price of the stock is unlimited at the upper end but cannot drop below \$0.
- The distribution of the stock price is positively skewed.
- The natural logarithm of the stock price yields a normal distribution.

These conditions match those of the lognormal distribution (Figure 103).

Figure 103 Lognormal distribution



In the lognormal distribution, the mean parameter is set at \$70.00 and the standard deviation set at \$12.00. The default location is 0, which works for this example. This distribution shows the probability that the stock price will be \$x.

Lognormal Parameter Sets

By default, the lognormal distribution uses the arithmetic mean and standard deviation. For applications where historical data are available, it is more appropriate to use the logarithmic mean and logarithmic standard deviation or the geometric mean and geometric standard deviation. These options are available from the Parameters menu in the menubar. Note that the location parameter is always in the arithmetic space.

For more information on these alternate parameters, see "Lognormal distribution" in the "Equations and Methods" chapter of the online Oracle Crystal Ball Statistical Guide. For more information about this menu, see "Alternate Parameter Sets" on page 45.

Maximum Extreme Distribution



The maximum extreme distribution is continuous. It is commonly used to describe the largest value of a response over a period of time: for example, in flood flows, rainfall, and earthquakes. Other applications include the breaking strengths of materials, construction design, and aircraft loads and tolerances. This distribution is also known as the Gumbel distribution, and is closely related to the minimum extreme distribution, its "mirror image."

The following sections describe the parameters, conditions, and other features of this distribution:

- "Maximum Extreme Parameters" on page 255
- "Maximum Extreme Conditions" on page 255

"Calculating Maximum Extreme Parameters" on page 255

Maximum Extreme Parameters

Likeliest, Scale

Maximum Extreme Conditions

The maximum extreme distribution is used under these conditions:

Conditions and parameters are complex. See: Castillo, Enrique. Extreme Value Theory in Engineering. London: Academic Press, 1988.

Note: After you select the Likeliest parameter, you can estimate the Scale parameter. The Scale parameter is a number greater than 0. The larger the Scale parameter, the greater the variance.

Calculating Maximum Extreme Parameters

There are two standard parameters for the maximum extreme value distribution: Likeliest and Scale. The Likeliest parameter is the most likely value for the variable (the highest point on the probability distribution, or mode).

After you select the Likeliest parameter, you can estimate the Scale parameter. The Scale parameter is a number greater than 0. The larger the Scale parameter, the greater the variance.

The Scale parameter symbol is:

α

To calculate a more exact scale, you can estimate the mean and use the equation:

$$\alpha = \left| \frac{\text{mean} - \text{mode}}{0.57721} \right|$$

Or estimate the variance and use the equation:

$$\alpha = \sqrt{\frac{6 \cdot \text{variance}}{\pi^2}}$$

Minimum Extreme Distribution



The minimum extreme distribution is continuous. It is commonly used to describe the smallest value of a response over a period of time: for example, rainfall during a drought. This distribution is closely related to the maximum extreme distribution.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Minimum Extreme Parameters" on page 256
- "Minimum Extreme Conditions" on page 256
- "Calculating Minimum Extreme Parameters" on page 256

Minimum Extreme Parameters

Likeliest, Scale

Minimum Extreme Conditions

The minimum extreme distribution is used under these conditions:

Conditions and parameters are complex. See: Castillo, Enrique. Extreme Value Theory in Engineering. London: Academic Press, 1988.

Note: After you select the Likeliest parameter, you can estimate the Scale parameter. The Scale parameter is a number greater than 0. The larger the Scale parameter, the greater the variance.

Calculating Minimum Extreme Parameters

There are two standard parameters for the minimum extreme value distribution: Likeliest and Scale. The Likeliest parameter is the most likely value for the variable (the highest point on the probability distribution, or mode).

After you select the Likeliest parameter, you can estimate the Scale parameter. The Scale parameter is a number greater than 0. The larger the Scale parameter, the greater the variance.

The Scale parameter symbol is:

To calculate a more exact scale, you can estimate the mean and use the equation:

$$\alpha = \left| \frac{\text{mean} - \text{mode}}{0.57721} \right|$$

Or estimate the variance and use the equation:

$$\alpha = \sqrt{\frac{6 \cdot \text{variance}}{\pi^2}}$$

Negative Binomial Distribution

The negative binomial distribution is discrete. It is useful for modeling the distribution of the number of trials until the rth successful occurrence, such as the number of sales calls you need to make to close ten orders. It is essentially a super-distribution of the geometric distribution.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Negative Binomial Parameters" on page 257
- "Negative Binomial Conditions" on page 257
- "Negative Binomial Example" on page 257

Negative Binomial Parameters

Probability, Shape

Negative Binomial Conditions

The negative binomial distribution is used under these conditions:

- Number of trials is not fixed.
- Trials continue to the *r*th success (trials never less than *r*).
- Probability of success is the same from trial to trial.

Negative Binomial Example

A manufacturer of jet engine turbines has an order to produce 50 turbines. Since about 20% of the turbines do not make it past the high-velocity spin test, the manufacturer will actually have to produce more than 50 turbines.


Matching these conditions with the negative binomial distribution:

- The number of turbines to produce (trials) is not fixed.
- The manufacturer will continue to produce turbines until the 50th one has passed the spin
- The probability of success (80%) is the same for each test.

These conditions match those of the negative binomial distribution.

The negative binomial distribution has two parameters: Probability and Shape. The Shape parameter specifies the r th successful occurrence. In this example you would enter 0.8 for the Probability parameter (80% success rate of the spin test) and 50 for the Shape parameter (Figure 104).

Figure 104 Negative Binomial Distribution



Some characteristics of the negative binomial distribution:

- When Shape = 1, the negative binomial distribution becomes the geometric distribution.
- The sum of any two negative binomial distributed variables is a negative binomial variable.
- Another form of the negative binomial distribution, sometimes found in textbooks, considers only the total number of failures until the r th successful occurrence, not the total number of trials. To model this form of the distribution, subtract out r (the value of the shape parameter) from the assumption value using a formula in the worksheet.

Normal Distribution



The normal distribution is continuous It is the most important distribution in probability theory because it describes many natural phenomena, such as people's IQs or heights and the reproductive rates of animals. Decision-makers can use the normal distribution to describe uncertain variables such as the inflation rate or the future price of gasoline.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Normal Parameters" on page 258
- "Normal Conditions" on page 259
- "Normal Example" on page 259

Normal Parameters

Mean, Standard Deviation

Normal Conditions

The normal distribution is used under these conditions:

- Mean value is most likely.
- It is symmetrical about the mean.
- More likely to be close to the mean than far away.

Note: Of the values of a normal distribution, approximately 68% are within 1 standard deviation on either side of the mean. The standard deviation is the square root of the average squared distance of values from the mean.

Normal Example

The normal distribution can be used to describe future inflation. You believe that 4% is the most likely rate. You are willing to bet that the inflation rate could as likely be above 4% as it could be below. You are also willing to bet that the inflation rate has a 68% chance of falling somewhere within 2% of the 4% rate. That is, you estimate there is approximately a two-thirds chance that the rate of inflation will be between 2% and 6%.


The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the normal distribution:

- The mean inflation rate is 4%.
- The inflation rate could as likely be above or below 4%.
- The inflation rate is more likely to be close to 4% than far away. In fact, there is approximately a 68% chance that the rate will lie within 2% of the mean rate of 4%.

These conditions match those of the normal distribution.

The normal distribution uses two parameters: Mean and Standard Deviation. Figure 105 shows the values from the example entered as parameters of the normal distribution in Crystal Ball: a mean of 0.04 (4%) and a standard deviation of 0.02 (2%). The distribution shows the probability of the inflation rate being a particular percentage.

Figure 105 Normal Distribution



Pareto Distribution



The Pareto distribution is continuous. It is widely used for the investigation of other distributions associated with such empirical phenomena as city population sizes, the occurrence of natural resources, the size of companies, personal incomes, stock price fluctuations, and error clustering in communication circuits.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Pareto Parameters" on page 260
- "Pareto Conditions" on page 260
- "Calculating Pareto Parameters" on page 261

Pareto Parameters

Location, Shape

Pareto Conditions

The Pareto distribution is used under these conditions:

Conditions and parameters are complex. See: Fishman, G. Springer Series in Operations Research. NY: Springer-Verlag, 1996.

Note: Note: The location parameter is the lower bound for the variable. After you select the Location parameter, you can estimate the shape parameter. The shape parameter is

a number greater than 0, usually greater than 1. The larger the shape parameter, the smaller the variance and the thicker the right tail of the distribution.

Calculating Pareto Parameters

There are two standard parameters for the Pareto distribution: Location and Shape. The Location parameter is the lower bound for the variable.

After you select the Location parameter, you can estimate the Shape parameter. The Shape parameter is a number greater than 0, usually greater than 1. The larger the Shape parameter, the smaller the variance and the thicker the right tail of the distribution.

The Shape parameter symbol is (beta).

The Location parameter symbol is L.

To calculate a more exact shape, you can estimate the mean and use the equation (for shapes greater than 1):

$$mean = \frac{\beta \cdot L}{\beta - 1}$$

Or estimate the variance and use the following equation (for shapes greater than 2):

variance =
$$\frac{\beta \cdot L^2}{(\beta - 2)(\beta - 1)^2}$$

Poisson Distribution



The Poisson distribution is discrete. It describes the number of times an event occurs in a given interval (usually time), such as the number of telephone calls per minute, the number of errors per page in a document, or the number of defects per 100 yards of material.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Poisson Parameter" on page 261
- "Poisson Conditions" on page 262
- "Poisson Example 1" on page 262
- "Poisson Example 2" on page 263

Poisson Parameter

Rate

Poisson Conditions

The Poisson distribution is used under these conditions:

- Number of possible occurrences is not limited.
- Occurrences are independent.
- Average number of occurrences is the same from unit to unit.

Poisson Example 1

An aerospace company wants to determine the number of defects per 100 square yards of carbon fiber material when the defects occur an average of 8 times per 100 square yards.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the Poisson distribution:

- Any number of defects is possible within 100 square yards.
- The occurrences are independent of one another. The number of defects in the first 100 square yards does not affect the number of defects in the second 100 square yards.
- The average number of defects (8) remains the same for each 100 square yards.

These conditions match those of the Poisson distribution.

The Poisson distribution has only one parameter: Rate. In this example, the value for this parameter is 8 (defects). You would enter this value to specify the parameter of the Poisson distribution in Crystal Ball.

The distribution illustrated in Figure 106 shows the probability of observing x number of defects in 100 square yards of the carbon fiber material.




Figure 106 **Poisson Distribution**

The size of the interval to which the rate applies, 100 square yards in this example, has no bearing on the probability distribution; the rate is the only key factor. If needed for modeling a situation, information on the size of the interval must be encoded in the spreadsheet formulas.

Poisson Example 2

A travel agency wants to describe the number of calls it receives in 10 minutes. The average number of calls in 10 minutes is about 35.

Again, you begin by identifying and entering the values to set the parameters of the Poisson distribution in Crystal Ball. In this example, the conditions show one important value: 35 calls or a rate of 35. The result would be a distribution showing the probability of receiving x number of calls in 10 minutes.

Student's t Distribution



The Student's *t* distribution is continuous. It is used to describe small sets of empirical data that resemble a normal curve, but with thicker tails (more outliers). It is often used for econometric data and exchange rates. For sets of data larger than 30, you can use the normal distribution instead.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Student's t Parameters" on page 263
- "Student's t Conditions" on page 263
- "Student's t Example" on page 264

Student's t Parameters

Midpoint, Scale, Degrees of Freedom

Student's t Conditions

The Student's *t* distribution is used under these conditions:

- Midpoint value is most likely.
- It is symmetrical about the mean.
- Approximates the Normal distribution when degrees of freedom are equal to or greater than 30.

Note: Note: The Midpoint parameter is the central location of the distribution (also mode), the x-axis value where you want to place the peak of the distribution. The Degrees of Freedom parameter controls the shape of the distribution. Smaller values result in thicker tails and less mass in the center. The Scale parameter affects the width of the distribution by increasing the variance without affecting the overall shape and proportions of the curve. Scale can be used to widen the curve for easier reading and interpretation. For example, if the midpoint were a large number, say 5000, the scale could be proportionately larger than if the midpoint were 500.

Student's t Example

For examples, see "Normal Distribution" on page 258. The uses are the same except that the sample degrees of freedom will be < 30 for the Student's t distribution.

Triangular Distribution



The triangular distribution is continuous. It describes a situation where you know the minimum, maximum, and most likely values to occur. It is useful with limited data in situations such as sales estimates, the number of cars sold in a week, inventory numbers, and marketing costs. For example, you could describe the number of cars sold per week when past sales show the minimum, maximum, and usual number of cars sold.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Triangular Parameters" on page 264
- "Triangular Conditions" on page 264
- "Triangular Example 1" on page 264
- "Triangular Example 2" on page 265

Triangular Parameters

Minimum, Likeliest, Maximum

Triangular Conditions

The triangular distribution is used under these conditions:

- Minimum and maximum are fixed.
- It has a most likely value in this range, which forms a triangle with the minimum and maximum.

Triangular Example 1

An owner needs to describe the amount of gasoline sold per week by his filling station. Past sales records show that a minimum of 3,000 gallons to a maximum of 7,000 gallons are sold per week, with most weeks showing sales of 5,000 gallons.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the triangular distribution:

- The minimum number of gallons is 3,000.
- The maximum number of gallons is 7,000.
- The most likely number of gallons (5,000) falls between 3,000 and 7,000, forming a triangle.

These conditions match those of the triangular distribution.

The triangular distribution has three parameters: Minimum, Likeliest, and Maximum. The conditions outlined in this example contain the values for these parameters: 3,000 (Minimum), 5,000 (Likeliest), and 7,000 (Maximum). You would enter these values as the parameters of the triangular distribution in Crystal Ball.

The following triangular distribution (Figure 107 on page 265shows the probability of x number of gallons being sold per week.




Figure 107 Triangular Distribution

Triangular Example 2

The triangular distribution also could be used to approximate a computer-controlled inventory situation. The computer is programmed to keep an ideal supply of 25 items on the shelf, always keep inventory at 10 items or more, and prevent inventory from ever rising above 30 items.

Check the triangular distribution conditions:

- The minimum inventory is 10 items.
- The maximum inventory is 30 items.
- The ideal level most frequently on the shelf is 25 items.

These conditions match those of the triangular distribution.

The result would be a distribution showing the probability of x number of items in inventory.

Uniform Distribution



The uniform distribution is continuous. In the uniform distribution, you know the range between the minimum and maximum values and you know that all values in the range are equally likely to occur. It can be used to describe a real estate appraisal or a leak in a pipe.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Uniform Parameters" on page 266
- "Uniform Conditions" on page 266
- "Uniform Example 1" on page 266
- "Uniform Example 2" on page 267

Uniform Parameters

Minimum, Maximum

Uniform Conditions

The uniform distribution is used under these conditions:

- Minimum is fixed.
- Maximum is fixed.
- All values in range are equally likely to occur.
- Discrete uniform is the discrete equivalent of the uniform distribution.

Uniform Example 1


An investment company interested in purchasing a parcel of prime commercial real estate wants to describe the appraised value of the property. The company expects an appraisal of at least \$500,000 but not more than \$900,000. They believe that all values between \$500,000 and \$900,000 have the same likelihood of being the actual appraised value.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. In this case:

- The minimum value is \$500,000.
- The maximum value is \$900,000.
- All values between \$500,000 and \$900,000 are equally possible.

These conditions match those of the uniform distribution. The uniform distribution has two parameters: the Minimum (\$500,000) and the Maximum (\$900,000). You would enter these values as the parameters of the uniform distribution in Crystal Ball (Figure 108).

Figure 108 Uniform Distribution



The distribution in Figure 108 shows that all values between \$500,000 and \$900,000 are equally possible.

Uniform Example 2

A manufacturer determines that he must receive 10% over production costs—or a minimum of \$3 per unit—to make the manufacturing effort worthwhile. He also wants to set the maximum price for the product at \$6 per unit, so that he can gain a sales advantage by offering the product for less than his nearest competitor. All values between \$3 and \$6 per unit have the same likelihood of being the actual product price.

The first step in selecting a probability distribution is matching the data with a distribution's conditions. Checking the uniform distribution:

- The minimum value is \$3 per unit.
- The maximum value is \$6 per unit.
- All values between \$3 and \$6 are equally possible.

You would enter these values in Crystal Ball to produce a uniform distribution showing that all values from \$3 to \$6 occur with equal likelihood.

Weibull Distribution



The Weibull distribution is continuous. It describes data resulting from fatigue tests and can be used to describe failure time in reliability studies or the breaking strengths of materials in reliability and quality control tests. Weibull distributions are also used to represent various physical quantities, such as wind speed.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Weibull Parameters" on page 268
- "Weibull Conditions" on page 268
- "Weibull Example" on page 268
- "Calculating Weibull Parameters" on page 268

Weibull Parameters

Location, Scale, Shape

Weibull Conditions

The Weibull distribution is used under these conditions:

- This flexible distribution can assume the properties of other distributions. When the shape parameter is equal to 1.0, the Weibull distribution is identical to the exponential distribution. The location parameter lets you set up an exponential distribution to start at a location other than 0.0. When the shape parameter is less than 1.0, the Weibull distribution becomes a steeply declining curve. A manufacturer might find this effect useful in describing part failures during a burn-in period.
- When shape parameters equal 1, it is identical to Exponential; when equal to 2, it is identical to Rayleigh.

Weibull Example

A lawn mower company is testing its gas-powered, self-propelled lawn mowers. They run 20 mowers, and keep track of how many hours each mower runs until its first breakdown. They use a Weibull distribution to describe the number of hours until the first failure.

Calculating Weibull Parameters

There are three standard parameters for the Weibull distribution: Location, Scale, and Shape. The Location parameter is the lower bound for the variable.

The Shape parameter is a number greater than 0, usually a small number less than 10. When the Shape parameter is less than 3, the distribution becomes more and more positively skewed until it starts to resemble an exponential distribution (shape < 1). At a shape of 3.25, the distribution is symmetrical, and above that value, the distribution becomes more narrow and negatively skewed.

After you select the Location and Shape parameter, you can estimate the Scale parameter. The larger the scale, the larger the width of the distribution.

To calculate a more exact scale, estimate the mean and use the equation:

$$\alpha \ = \ \frac{mean - L}{\Gamma\!\!\left(1 + \frac{1}{\beta}\right)}$$

where (alpha) is the scale, (beta) is the shape, L is the location, and (gamma) is the gamma function. You can use the Microsoft Excel GAMMALN function and Microsoft Excel Solver to help you calculate this parameter.

For this distribution, there is a 63% probability that x falls between a and a+L.

Or estimate the mode and use the equation:

$$\alpha = \frac{\text{mode} - L}{\left(1 - \frac{1}{\beta}\right)^{\frac{1}{\beta}}}$$

where (alpha) is the scale, (beta) is the shape, and L is the location.

Yes-No Distribution



The yes-no distribution, also called the Bernoulli distribution, is a discrete distribution that describes a set of observations that can have only one of two values, such as yes or no, success or failure, true or false, or heads or tails.

The following sections describe the parameters, conditions, and other features of this distribution:

- "Yes-No Parameters" on page 269
- "Yes-No Parameters" on page 269
- "Yes-No Example" on page 270

Yes-No Parameters

Probability of Yes (1)

Yes-No Conditions

The yes-no distribution is used under these conditions:


- For each trial, only 2 outcomes are possible, such as success or failure; the random variable can have only one of two values, for example, 0 and 1.
- The mean is p, or probability (0 .

- Trials are independent. Probability is the same from trial to trial.
- The Yes-No distribution is equivalent to the Binomial distribution with one trial.

Yes-No Example

A machine shop produces complex high-tolerance parts with a .02 probability of failure and a . 98 probability of success. If a single part is pulled from the line, Figure 109 shows the probability that the part is good.

Figure 109 Probability of Pulling a Good Part



Using the Custom Distribution



If none of the provided distributions fits the data, you can use the custom distribution to define the own. For example, a custom distribution can be especially helpful if different ranges of values have specific probabilities. You can create a distribution of one shape for one range of values and a different distribution for another range. You can describe a series of single values, discrete ranges, or continuous ranges. This section uses real-world examples to describe the custom distribution.

Since it is easier to understand how the custom distribution works with a hands-on example, you may want to start Crystal Ball and use it to follow the examples. To follow the custom distribution examples, first create a new Microsoft Excel workbook then select cells as specified.

The following sections show how to use the custom distribution:

- "Custom Distribution Example 1" on page 271
- "Custom Distribution Example 2" on page 273
- "Custom Distribution Example 3 Loading Data" on page 276

- "Entering Tables of Data into Custom Distributions" on page 279
- "Changes from Crystal Ball 2000.x (5.x)" on page 284
- "Other Important Custom Distribution Notes" on page 284

For summaries of the data entry rules used in the examples plus additional rules, see "Entering Tables of Data into Custom Distributions" on page 279 and "Other Important Custom Distribution Notes" on page 284.

Custom Distribution Example 1


- ➤ Before beginning example one, open the Custom Distribution dialog as follows:
- Click cell **D11**.
- Select **Define**, then **Define Assumption**.

The Distribution Gallery dialog opens.

- Click the All category to select it.
- Scroll to find the custom distribution, then click it.
- Click OK.

Crystal Ball displays the Define Assumption dialog.

Figure 110 Define Assumption Dialog for Custom Distributions



Using the custom distribution, a company can describe the probable retail cost of a new product. The company decides the cost could be \$5, \$8, or \$10. In this example, you will use the custom distribution to describe a series of single values.

- To enter the parameters of this custom distribution:
- Type 5 in the Value field and click Enter.

Since you do not specify a probability, Crystal Ball defaults to a relative probability of 1.00 for the value 5. A single value bar displays the value 5.00.

Relative probability means that the sum of the probabilities does not have to add up to 1. So the probability for a given value is meaningless by itself; it makes sense only in relation to the total relative probability. For example, if the total relative probability is 3 and the relative probability for a given value is 1, the value has a probability of 0.33.

Type 8 in the Value field.

Click Enter.

Since you did not specify a probability, Crystal Ball defaults to a relative probability of 1.00 (displayed on the Probability scale along the first vertical axis) for the value 8. A second value bar represents the value 8.

Type 10 in the Value field.

Click Enter.

Crystal Ball displays a relative probability of 1.00 for the value 10. A third single value bar represents the value 10.

Figure 111 shows the value bars for the values 5, 8, and 10, each with a relative probability of 1.00.




Figure 111 Single Values

Now, each value has a probability of 1. However, when you run the simulation, their total relative probability becomes 1.00 and the probability of each value is reset to .3333.

- To reset their probabilities before you run the simulation, follow these steps:
- Click the bar with a value of 5.00.

Its value is displayed in the Value field.


Type the probability as the formula =1/3 in the Probability field and click **Enter**.

You could also enter a decimal — for example, 0.333333 — but the formula is more exact.

Follow steps 6 and 7 for the other two bars.

Crystal Ball rescales each value to a relative probability of 0.33 on the Relative Probability scale along the first vertical axis (Figure 112).

Figure 112 Single Values with Adjusted Probabilities



Custom Distribution Example 2

- **>** Before beginning example two, clear the values entered in example one as follows:
- Right-click in the chart and choose Clear Distribution from the menu.


In this example, you will use the custom distribution to describe a continuous range of values, since the unit cost can take on any value within the specified intervals.

- Choose Parameters, then Continuous Ranges to enter value ranges.
- Enter the first range of values:
 - Type 5 in the Minimum field.
 - Type 15 in the Maximum field.
 - Type .75 in the Probability field. This represents the total probability of all values within the range.

Click Enter.

Crystal Ball displays a continuous value bar for the range 5.00 to 15.00, as in Figure 113, and returns the Crystal Ball cursor to the Minimum field. Notice that the height of the range is 0.075. This represents the total probability divided by the width (number of units) in the range, 10 (Figure 113).

Figure 113 A Continuous Custom Distribution




Enter the second range of values:

- Type 16 in the Minimum field.
- Type 21 in the Maximum field.
- Type .25 in the Probability field.
- Click Enter.

Crystal Ball displays a continuous value bar for the range 16.00 to 21.00. Its height is . 050, equal to .25 divided by 5, the number of units in the range. Both ranges now are displayed in the Custom Distribution dialog (Figure 114).

Figure 114 Custom Distribution with Two Continuous Ranges



- You can change the probability and slope of a continuous range, as described in the following steps:
- Click anywhere on the value bar for the range 16 to 21.

The value bar changes to a lighter shade.

Choose Parameters, then Sloping Ranges.

Additional parameters are displayed in the Custom Distribution dialog (Figure 115).

Figure 115 Sloping Range Parameters, Custom Distribution Dialog

Enter one or more continuous or discrete sloping ranges:							
Minimum	16.00	Maximum	21.00				
Height of Min.	0.25	Height of Max.		Step			

Set the Height of Min. and Height of Max. equal to what currently is displayed in the chart, 0.05.

This can be an approximate value. The Height of Min. is the height of the range Minimum and the Height of Max. is the height of the range Maximum.


Click Enter.

The range returns to its original color and its height is displayed unchanged.

Click in the range again to select it and set the Height of Max. to 0.025. Then, click Enter.

The second side of the range drops to half the height of the first, as shown in Figure 116. The range is selected to show its parameters after the change.

Figure 116 Sloping Continuous Value Range




Optional: Change the range from continuous to discrete values by adding a step value. Type .5 in the Step field and click Enter.

The sloped range is now discrete. Separate bars are displayed at the beginning and end of the range and every half unit in between (16, 16.5, 17, 17.5 and so on until 21), as shown in Figure 117. If the discrete range represented money, it could only include whole dollars and 50-cent increments.

You can enter any positive number in the Step field. If you entered 1 in this example, the steps would fall on consecutive integers, such as whole dollars. Leave the Step parameter blank for continuous ranges.

Figure 117 A Sloped Discrete Range with Steps of .5



Although the bars have spaces between them, their heights and the width of the range they cover are equal to the previous continuous sloped range and the total probability is the same.

While a second continuous range could have extended from 15 to 20, the second range in this example starts at 16 rather than 15 to illustrate a discrete range because, unlike continuous ranges, discrete ranges cannot touch other ranges.


With Crystal Ball, you can enter single values, discrete ranges, or continuous ranges individually. You also can enter any combination of these three types in the same Custom Distribution dialog as long as you follow these guidelines: ranges and single values cannot overlap one another; however, the ending value of one continuous range can be the starting value of another continuous range.

Custom Distribution Example 3 — Loading Data

This example describes a special feature on the Custom Distribution dialog: the Load Data button, which lets you pull in numbers from a specified cell range (grouped data) on the worksheet. This example is not a hands-on exercise, but the illustrations will guide you through the procedure. After you read this section, you can experiment with the own data by pulling in numbers from specified cell ranges on the worksheet.

In this example, the same company decides that the unit cost of the new product can vary widely. The company feels it has a 20% chance of being any number between \$10 and \$20, a 10% chance of being any number between \$20 and \$30, a 30% chance of being any number between \$40 and \$50, a 30% chance of being a whole dollar amount between \$60 and \$80, and there is a 5% chance the value will be either \$90 or \$100. All the values have been entered on the worksheet in this order: range minimum value, range maximum value (for all but Single Value ranges), total probability, and step (for the Discrete Range only) as shown in Figure 118.

Figure 118 Four-column Custom Data Range




In this case, discrete ranges have the most parameters. So, you can create an assumption, choose Custom Distribution, and then choose Parameters, then Discrete Ranges before loading the data.

If the data also included discrete sloping ranges, you could choose Parameters, then Sloping Ranges before loading the data. The data table would then have five columns and could accommodate all data types.

- To complete the data load after the Parameters setting has been made:
- Click the More button beside the Name field.

The Custom Distribution dialog expands to include a data table, as shown in Figure 119.

Figure 119 Custom Distribution with Data Table




A column is displayed for each parameter in the current set (selected using the Parameters menu). Parameters, then Discrete Ranges was set before viewing the data table, so there is a column in the data table for each discrete range parameter. Because the single value and continuous ranges have subsets of the same group of parameters, their parameters will also fit into the table.

Since the values are already on the worksheet, you can click Load Data to enter them into the Custom Distribution dialog.

The Load Data dialog opens, as shown in Figure 120.

Figure 120 Load Data Dialog, Custom Distribution




The default settings are appropriate for most purposes, but the following other options are available:

- When loading unlinked data, you can choose to replace the current distribution with the new data or append new data to the existing distribution.
- If probabilities are entered cumulatively into the spreadsheet you are loading, you can check Probabilities Are Cumulative. Then, Crystal Ball determines the probabilities for each range by subtracting the previous probability from the one entered for the current range. You will need to choose View, then Cumulative Probability to display the data cumulatively in the assumption chart.
- Enter a location range for the data, in this case A2:D7. If the range has a name, you can enter the name preceded by an = sign.
- When all settings are correct, click **OK**.

Crystal Ball enters the values from the specified range into the custom distribution and plots the specified ranges, as shown in Figure 121.

Figure 121 Custom Data from Worksheet



Entering Tables of Data into Custom Distributions

Follow the rules in this section for loading data. If a data range has been named, you can enter the name (preceded by the = sign) in the range field of the Load Data dialog.

The following sections show how to enter tables of data into custom distributions:

- "Unweighted Values" on page 279
- "Weighted Values" on page 280
- "Mixed Single Values, Continuous Ranges, and Discrete Ranges" on page 281
- "Mixed Ranges, Including Sloping Ranges" on page 281
- "Connected Series of Ranges (Sloping)" on page 282
- "Connected Series of Continuous Uniform Ranges (Cumulative)" on page 283
- "Other Data Load Notes" on page 283

Unweighted Values


Single values are values that do not define a range. Each value stands alone. For a series of single values with the same probabilities (unweighted values), use a one-column format or more than five columns (Figure 122). The values go in each cell and the relative probabilities are all assumed to be 1.0. Choose Parameters, then Unweighted Values to enter these.

Figure 122 Single Values with the Same Probability

	Α	В	С	D	Е	F	G
1	2	8	9	22	24	32	
2	5	10	12	26	34	27	
3	7	11	30	14	39		
4							

Figure 123 shows the custom distribution created by loading the unweighted values illustrated in Figure 122.

Figure 123 Unweighted Values Loaded in a Custom Distribution



Weighted Values

For a series of single values all with different probabilities, use a two-column format. The first column contains single values, the second column contains the probability of each value (Figure 124).

Single Values with Different Probabilities (Weighted Values) Figure 124





Figure 125 shows the custom distribution created by loading the values illustrated in Figure 124.

Figure 125 Weighted Values Loaded in a Custom Distribution



Mixed Single Values, Continuous Ranges, and Discrete Ranges

For any mixture of single values and continuous ranges, use a three-column format, obtained by choosing Parameters, then Continuous ranges. The three-column format is the same as using the first three columns shown in Figure 119, Figure 120, and Figure 121.

If the mix includes uniform (non-sloping) discrete ranges, use a four-column format, as in the first four columns of Figure 126 and Figure 127. To obtain four columns, choose Parameters, then Discrete Ranges.

Mixed Ranges, Including Sloping Ranges

If sloping ranges are included in a mix of ranges, choose Parameters, then Sloping Ranges to display a five-column data table. The first column contains the range Minimum value, the second column contains the range Maximum value, the third column contains Height of Min. (the relative probability — height — at the Minimum value), the fourth column contains Height of Max. (the relative probability at the Maximum value), and the fifth column contains the Step value for discrete sloping ranges. For continuous sloping ranges the fifth column (Step) is left blank.


Note that if there are uniform discrete ranges, their first three columns contain the Minimum, Maximum, and Probability as in a four-column format but the fourth column is left blank and Step is entered in the fifth column (Figure 126).

Figure 126 Mixed Ranges, Including Sloping Ranges

	A	В	С	D	E	F
14	Custom Value Description	Minimum	Maximum	Height of Min.	Height of Max.	Step
15	Discrete Value = 0.5, Probability = 1.0	0.5				
16	Continuous Range = 1-2, Uniform Probability = 1.0	1	2			
17	Discrete Value = 2.5, Probability = 5	2.5		5		
18	Continuous Range = 3-5, Uniform Probability = 4	3	5	4		
19	Continuous Range = 5-7, Sloping Probability = 3-6	5	7	3	6	
20	Continuous Connected Range = 7-10, Sloping Probability = 6-8		10		8	
21	Discrete Range = 12-14, Uniform Probability = 2, Step = 0.5	12	14	2		0.5
22	Discrete Range = 15-17, Sloping Probability = 1-4, Step = 1	15	17	1	4	1

Figure 127 shows the custom distribution created by loading the values illustrated in Figure 126.

Mixed Ranges Loaded in a Custom Distribution



Connected Series of Ranges (Sloping)

For a connected series of sloping continuous ranges, choose Parameters, then Sloping Ranges to use a five-column format. The first column contains the lowest Minimum value of the highestvalue (right-most) range, the second column contains the Maximum value of each connected range, the third column contains the Height of Min. (relative probability of the Minimum value) if it differs from the previous Height of Max. (otherwise it is left empty), and the fourth column contains Height of Max. (relative probability of the Maximum value) for that range. The fifth column is left blank for continuous ranges but a fifth column is necessary to indicate that these are sloping ranges.

For example, row 20 in Figure 126 shows a connected continuous sloping range. The Minimum cell is blank because the Minimum value is equal to 7, the previous Maximum. The Height of Min. is blank because it is equal to 6, the previous Height of Max.

Connected Series of Continuous Uniform Ranges (Cumulative)

For a connected series of continuous uniform ranges specified using cumulative probabilities, use a three-column format with the common endpoints of the ranges in the second column and the cumulative probabilities in the third column. The first column is left blank except for the minimum value of the first range, beside the maximum in the second column (Figure 128). Be sure to check Probabilities Are Cumulative in the Load Data dialog.

Figure 128 Connected Continuous Uniform Ranges

	Α	В	С	
		Maximums of		
	Minimum of	Continuous	Cumulative	
1	First Range	Ranges	Probabilities	
2	5	10	0.25	
3		15	0.50	
4		25	0.75	
5		30	0.90	

Figure 129 shows the custom distribution created by loading the data illustrated in Figure 128.




Figure 129 Connected Continuous Uniform Ranges After Loading

Other Data Load Notes

You can load each type of range separately or you can specify the range type with the greatest number of parameters and load all types together. Other rules are.

- Cumulative probabilities are supported for all but sloping ranges.
- Blank probabilities are interpreted as a relative probability of 1.0.

- For continuous connected ranges, for either endpoint values or probabilities, if the starting cell is blank, the previous end value is used as the start for this range.
- When you load a discrete value that exists in the table already, its probability is incremented by 1. For continuous ranges, this is not allowed; an error message about overlapping ranges opens.

Note: In versions of Crystal Ball earlier than 11.1.1.3.00, ranges or values with 0 probabilities were removed. Sloping ranges with Height of Min. and Height of Max. equal to 0 were also removed. The current rules are as follows:

In the simple view for weighted values, if users enter a probability of 0 for a value, the users are prompted about deleting the value. For other types of values, 0 works as before to delete the range or value.

In advanced (table) view for weighted values, if users enter a probability of 0 for a value, it is ignored (or added, which is the same thing). For other types of values, 0 works as before to delete the range or value.

If a custom distribution was created with unlinked data in a version of Crystal Ball earlier than version 11.1.1.3.00, the frequency distribution does not change when the model is loaded in Crystal Ball version 11.1.1.3.00 or later. However, if the input range was linked for a custom distribution, the frequency distribution is updated according to current rules. All new custom distributions, linked and unlinked, are evaluated according to current rules.

Changes from Crystal Ball 2000.x (5.x)

In versions of Crystal Ball earlier than version 7.0, discrete values with the same probability could be entered in ranges with five columns or more. Now, they cannot be entered in ranges with five columns but can only be entered in single columns or ranges with six or more columns (to distinguish them from sloping ranges).

In versions of Crystal Ball earlier than version 7.0, continuous uniform ranges with cumulative probabilities could be entered in a two-column format. Now a three-column format is required, discussed in "Connected Series of Continuous Uniform Ranges (Cumulative)" on page 283.

The three-column sloping range format used in previous versions of Crystal Ball has been replaced by a five-column format, described in "Mixed Ranges, Including Sloping Ranges" on page 281 and the section that follows it, "Connected Series of Ranges (Sloping)" on page 282

Other Important Custom Distribution Notes

Even if you do not load data from the spreadsheet into the Custom Distribution dialog, you can still add and edit data using the data table. To do this, click the More button, 🦭, to display the data table. Then, you can:

Enter a different value in the data table and click Enter to change the data.

- Type the minimum, maximum, probability, and step (if discrete data) into a blank row and click Enter to add new data.
- To delete a single range of data, select that row of data, right-click, and choose Delete Row.
- To clear all data rows, right-click within the data table and choose Clear Distribution.
- To delete a single range of data without using the data table, click the range to select it and either:
 - Set the Probability or Height of Min. and Height of Max. to 0, or 0
 - Choose Edit, then Delete Row or right-click and choose Delete Row.
- Statistics for custom distributions are approximate.


Truncating Distributions

You can change the bounds or limits of each distribution, except the custom distribution, by dragging the truncation grabbers or by entering different numeric endpoints for the truncation grabbers. This "truncates" the distribution. You can also exclude a middle area of a distribution by crossing over the truncation grabbers to "white out" the portion to exclude.

Note: To display the truncation grabbers, open an assumption in the Define Assumption dialog and click the More button beside the assumption name field.

For example, suppose you want to describe the selling price of a house up for auction after foreclosure. The bank that holds the mortgage will not sell for less than \$80,000. They expect the bids to be normally distributed around \$100,000 with a standard deviation of \$15,000. In Crystal Bally, ou would specify the mean as 100,000 and the standard deviation as 15,000 and then move the first (left) grabber to set the limit of 80,000. The grabber "whites out" the portion to exclude, as shown in Figure 130.

Figure 130 Truncated distribution example



Be Aware...

Each adjustment changes the characteristics of the probability distribution. For example, the truncated normal distribution in Figure 130 will no longer have an actual mean of \$100,000 and standard deviation of \$15,000. Also, statistics values will be approximate for truncated distributions.

When using alternate percentile parameters, the actual percentiles calculated for a truncated distribution will differ from the specified parameter values. For example, a normal distribution specified with 10th/90th percentiles and truncated on either side of the distribution will have actual 10th/90th percentiles greater or less than the specified percentiles.

Showing the mean line of the distribution is useful when truncating distributions. However, the mean line value might differ from the Mean parameter field. The mean line shows the actual mean of the truncated distribution while the Mean parameter field shows the mean of the complete distribution.

Comparing Distributions

Many of the distributions discussed in this chapter are related to one another in various ways. For example, the geometric distribution is related to the binomial distribution. The geometric distribution represents the number of trials until the next success while the binomial represents the number of successes in a fixed number of trials. Similarly, the Poisson distribution is related to the exponential distribution. The exponential distribution represents the amount of time until the next occurrence of an event while the Poisson distribution represents the number of times an event occurs within a given period of time.

In some situations, as when the number of trials for the binomial distribution becomes very large, the normal and binomial distributions become similar. For these two distributions, as the number of binomial trials approaches infinity, the probabilities become identical for any given interval. For this reason, you can use the normal distribution to approximate the binomial distribution when the number of trials becomes too large for Crystal Ball to handle (more than 1000 trials). You also can use the Poisson distribution to approximate the binomial distribution when the number of trials is large, but there is little advantage to this since Crystal Ball takes a comparable amount of time to compute both distributions.

Likewise, the normal and Student's t distributions are related. With Degrees of Freedom, then 30, Student's t closely approximates the normal distribution.

The binomial and hypergeometric distributions are also closely related. As the number of trials and the population size increase, the hypergeometric trials tend to become independent like the binomial trials: the outcome of a single trial has a negligible effect on the probabilities of successive observations. The differences between these two types of distributions become important only when you are analyzing samples from relatively small populations. As with the Poisson and binomial distributions, Crystal Ball requires a similar amount of time to compute both the binomial and hypergeometric distributions.

The yes-no distribution is simply the binomial distribution with Trials = 1.

The Weibull distribution is flexible. Actually, it consists of a family of distributions that can assume the properties of several distributions. When the Weibull shape parameter is 1.0, the Weibull distribution is identical to the exponential distribution. The Weibull location parameter lets you set up an exponential distribution to start at a location other than 0.0. When the shape parameter is less than 1.0, the Weibull distribution becomes a steeply declining curve. A manufacturer might find this effect useful in describing part failures during a burn-in period. When the shape parameter is equal to 2.0, a special form of the Weibull distribution, called the Rayleigh distribution, results. A researcher might find the Rayleigh distribution useful for analyzing noise problems in communication systems or for use in reliability studies. When the shape parameter is set to 3.25, the Weibull distribution approximates the shape of the normal distribution; however, for applications when the normal distribution is appropriate, us it instead of the Weibull distribution.

The gamma distribution is also a flexible family of distributions. When the shape parameter is 1.0, the gamma distribution is identical to the exponential distribution. When the shape parameter is an integer greater than one, a special form of the gamma distribution, called the Erlang distribution, results. The Erlang distribution is especially useful in the areas of inventory control and queueing theory, where events tend to follow Poisson processes. Finally, when the shape parameter is an integer plus one half (e.g., 1.5, 2.5, etc.), the result is a chi-square distribution, useful for modeling the effects between the observed and expected outcomes of a random sampling.

When no other distribution seems to fit the historical data or accurately describes an uncertain variable, you can use the custom distribution to simulate almost any distribution. The Load Data button on the Custom Distribution dialog lets you read a series of data points or ranges from value cells in the worksheet. If you like, you can use the mouse to individually alter the probabilities and shapes of the data points and ranges so that they more accurately reflect the uncertain variable.

Distribution Parameter Summary

The following table lists valid parameter values for each Crystal Ball distribution. Distributions are listed alphabetically by type (continuous or discrete). The Oracle Crystal Ball Statistical Guide lists defaults for each parameter value.

Table 10 Distributions and Their Parameters

Distribution		Туре	Parameter 1	Parameter 2	Parameter 3	Parameter 4
Beta	Beta	continuous	alpha (greater than 0.3, alpha + beta must be less than 1000)	beta (greater than 0.3, alpha + beta must be less than 1000)	maximum value	minimum value
BetaPERT	BetaPERT	continuous	minimum value	likeliest value	maximum value	N/A
Exponential	Exponential	continuous	rate (greater than 0)	N/A	N/A	N/A
Gamma	Gamma	continuous	location	scale (greater than 0)	shape (greater than 0.05 and less than 1000)	N/A
Logistic	Logistic	continuous	mean value	scale (greater than 0)	N/A	N/A
Lognormal	Lognormal	continuous	location	mean value	standard deviation value	N/A
Maximum extreme	Max Extreme	continuous	likeliest	scale (greater than 0)	N/A	N/A
Minimum extreme	Min Extreme	continuous	likeliest	scale (greater than 0)	N/A	N/A
Normal	Normal	continuous	mean value	standard deviation value	N/A	N/A
Pareto	Pareto	continuous	location (greater than 0)	shape (greater than 0. 05)	N/A	N/A

Distribution		Туре	Parameter 1	Parameter 2	Parameter 3	Parameter 4
Student's t	Student's t	continuous	midpoint	scale (greater than 0)	degrees of freedom (integer between 1 and 30, inclusive)	N/A
Triangular	Triangular	continuous	minimum value	likeliest value	maximum value	N/A
Uniform	Uniform	continuous	minimum value	maximum value	N/A	N/A
Weibull	Weibull	continuous	location	scale (greater than 0)	shape (greater than 0.05)	N/A
Binomial	Binomial	discrete	probability (between 0 and 1)	trials (a whole number greater than 0 and less than 1000)	N/A	N/A
Discrete Uniform	Discrete Uniform	discrete	minimum (integer)	maximum (integer)	N/A	N/A
Geometric	Geometric	discrete	probability (between 0 and 1)	N/A	N/A	N/A
Hyper- geometric	Hypergeometric	discrete	success	trials (whole number less than population)	population (whole number greater than 0 and less than 1000)	N/A
Negative binomial	Neg Binomial	discrete	probability (between 0 and 1)	shape (whole number greater than 0 and less than 1000)	N/A	N/A
Poisson	Poisson	discrete	rate value (between 0 and 1000)	N/A	N/A	N/A
Yes-No	Yes-No	discrete	probability (between 0 and 1)	N/A	N/A	N/A
Custom	Custom	custom	See Appendix A of the current <i>Oracle</i> <i>Crystal Ball User's</i> <i>Guide</i> .	N/A	N/A	N/A

Using Probability Functions


For each of the Crystal Ball distributions, there is an equivalent Microsoft Excel function. You can enter these functions in the spreadsheet directly instead of defining distributions using the Define Assumption command. Be aware, though, that there are a number of limitations in using these functions. For details, see "Limitations of Probability Functions" on page 290.

Also see "Probability Functions and Random Seeds" on page 291 for information about setting a random seed so probability functions can have comparable values.

To view these functions and their parameters, choose **Insert**, then **Function** in Microsoft Excel, and then be sure the category is set to Crystal Ball (Figure 131, following).

(In Microsoft Excel 2007 or later, choose Formulas, then Insert Function.)

Figure 131 Crystal Ball Functions in Microsoft Excel



Parameters and a brief description are displayed below the list of functions. The Cutoff parameters indicate truncation values, while NameOf is the assumption name. For parameter descriptions and details on each distribution, see the entry for that distribution earlier in this appendix.

Note: The beta distribution is now different from versions earlier than Crystal Ball 7.0. Both the original and revised functions are displayed for compatibility. CB.Beta has three parameters but CB.Beta2 is the current Crystal Ball version with Minimum and Maximum instead of Scale.

Limitations of Probability Functions

Distributions defined with probability functions differ from those entered with the Define Assumption command in these ways:

- You cannot correlate them.
- You cannot view charts or statistics on them.

- You cannot extract data from them or include them in reports.
- They are not included in sensitivity analyses or charts.
- Latin Hypercube sampling is not supported.

Probability Functions and Random Seeds

"Setting Sampling Preferences" on page 79 describes how you can use the Sampling tab of the Run Preferences dialog to use the same sequence of random numbers for each simulation.

If you use Define, then Define Assumption or the Define Assumption toolbar button to define assumptions, the same sequence of random numbers is used for each simulation, even if you switch from Extreme speed to Normal speed or back to Extreme speed. If you use the probability functions to define assumptions, one sequence of random numbers is used for Extreme speed and a different sequence is used for Normal speed.

Sequential Sampling with Custom Distributions

The probability distributions supplied with Crystal Ball are useful in a variety of modeling situations. Organizations might still want to prepare their own libraries of distributions based on data specific to their applications and situations. One such system involves libraries of stochastic information packets (SIPs), an approach set forth in the article "Probability Management" (see "Sequential Sampling with SIPs" on page 414; Savage, S. et al, 2006).

A SIP is a list of time- or order-sensitive values for a particular variable. These values are sampled as sequential trials during a Monte Carlo simulation. SIPs are used to preserve the correlation structure between SIP variables without having to explicitly compute and define a matrix of correlation coefficients.

SIPs can be represented by custom distributions in Crystal Ball and can then be published and shared by organizations using Crystal Ball's Publish and Subscribe features in the Distribution Gallery.

For details, see:

- "Creating Custom SIP Distributions" on page 291
- "Running Simulations with SIPs" on page 292

Creating Custom SIP Distributions

The easiest way to load a SIP or similar set of values into Crystal Ball is to organize the data in a single column in Microsoft Excel.


- Then, define an assumption in Crystal Ball using a custom distribution:
- Select a cell and choose Define, then Define Assumption.
- Looking at the list of All distributions in the Distribution Gallery, scroll to the Custom distribution and double-click it.

- In the Define Assumption dialog, choose Parameters, then Sample Sequentially. This also switches the parameter set to Unweighted Values.
- Click the More button, \bigotimes , to expand the dialog.
- 5 Click Load Data and enter the range of data in the Location of Data field.

You can type in the range or a range name, or click the cell selector icon to select the range. Be sure Sample Sequentially (instead of randomly) is still checked and click OK.

The distribution is displayed in the Define Assumption dialog, as shown in Figure 132, following.





Now, you can choose Edit, then Add to Gallery to add the SIP to a custom library (Category) in the Distribution Gallery. You can also use Publish and Subscribe on the Distribution Gallery Categories menu to share the library of SIPs with others.

Running Simulations with SIPs

To run a simulation with one or more SIPs, define assumptions in the model as usual, pulling in SIPs from a library (Category) in the Distribution Gallery as needed. When the simulation runs, Crystal Ball samples the SIP values sequentially instead of randomly. If the simulation has more trials than there are values in the SIP, sampling wraps around to the first value of the distribution. If you are using several SIPs with different numbers of values, a warning is displayed. You can either continue, knowing that mismatched values will be sampled together after sampling wraps for one SIP but not the other, or you can cancel the simulation and correct the problem. You can also run fewer trials so that sampling doesn't need to wrap for any SIP.

Correlations are not allowed for assumptions defined as SIPs, since the correlation structures are implicitly defined for them. Warnings are displayed if correlations are defined for them.

Two Crystal Ball Developer Kit calls, CB.DefineAssumND and CB.GetAssum, support the new parameter for sequential sampling. The Developer Kit constant is called cbAsmIsSequential, a Boolean type.



Using the Extreme Speed Feature

In This Appendix

Overview	293
Compatibility Issues	294
Numerical Differences and the Compare Run Modes Tool	300
Other Important Differences	302
Maximizing the Benefits of Extreme Speed	304

Overview

Extreme speed, available in Crystal Ball Decision Optimizer, runs simulations up to 100 times faster than Normal speed. With Extreme speed, you won't need to wait for simulation results or compromise on the number of trials you run. Extreme speed makes it more practical to use Crystal Ball tools that run multi-simulations or to seek optimal solutions with OptQuest in a reasonable length of time.

The Extreme Speed feature uses PSI Technology, a high-speed, Microsoft Excel-compatible Polymorphic Spreadsheet Interpreter for running simulations on workbooks. This technology was developed by Frontline Systems, the maker of Microsoft Excel's Solver add-in. PSI Technology supports nearly all of the 320 standard Microsoft Excel functions, including the financial, statistical, and engineering functions that are part of the Analysis Toolpak.

By default, Crystal Ball Decision Optimizer is set to use Extreme speed when initially loaded. If a model is not compatible with Extreme speed, a dialog offers the opportunity to downshift to Normal speed temporarily for that simulation. "Compatibility Issues" on page 294 below, explains conditions that cause a model to be incompatible with Extreme speed. The simulation speed can be changed using the Speed tab in the Run Preferences dialogs (see "Setting Speed Preferences" on page 80 for details).

This section contains the following main topics:

- "Compatibility Issues" on page 294
- "Numerical Differences and the Compare Run Modes Tool" on page 300
- "Other Important Differences" on page 302
- "Maximizing the Benefits of Extreme Speed" on page 304

Note: Because the Extreme Speed feature uses the Polymorphic Spreadsheet Interpreter, certain Extreme Speed functions may return slightly different results than the corresponding Microsoft Excel function for extreme values of any input arguments. For example, this can occur with statistical and inverse distribution functions.

Compatibility Issues

While Extreme speed can greatly decrease simulation run times, not all models are compatible with Extreme speed. When you start a simulation, Crystal Ball detects if the spreadsheet is compatible with Extreme speed and warns you if it is incompatible. If you choose, you can run the simulation in Normal speed using standard Microsoft Excel, or you can change the spreadsheet model to correct the incompatibility.

This section lists functions and formula constructs that are not compatible with Extreme speed and suggests workarounds. Besides the issues listed below, some differences may also exist due to undocumented Microsoft Excel behaviors, changes in the newest versions of Microsoft Excel, and so on. Also, small differences in the last few decimal places of certain built-in function values should be expected, due to minor algorithmic differences in the way formulas might be computed.

It is important to note that incompatibilities in functions and formula constructs only concern the cells involved in the calculation of a forecast cell. If there are incompatibilities in cells that are not part of that calculation path, these are not detected and the simulation is allowed to run.

This section discusses the following compatibility issues:

- "Multiple-Workbook Models" on page 294
- "Circular References" on page 295
- "Crystal Ball Microsoft Excel Functions" on page 295
- "User-Defined Functions" on page 296
- "Running User-Defined Macros" on page 298
- "Special Functions" on page 298
- "Undocumented Behavior of Standard Functions" on page 298
- "Incompatible Range Constructs" on page 299
- "Data Tables" on page 300

Multiple-Workbook Models

Extreme speed can now run simulations on multiple workbooks. If you are running in Extreme speed and the workbook contains external references to cells in other closed workbooks, Crystal Ball obtains the current value from those workbooks. References to cells in other open workbooks are dynamically updated if those cells depend on one or more assumptions. If the external reference is part of a formula (not a simple external reference), this is not compatible with Extreme speed:

- Example message: "Unable to interpret formula at cell location [Book1.xls] Sheet1!A1. (Code #5524 - Complex external reference)"
- Workaround: If possible, consolidate all of the variables and formulas from a multiworkbook model containing Crystal Ball data into a single workbook.

Circular References

Circular references within a model are supported as long as Iteration is checked on the Tools, then Options, then Calculation tab in Microsoft Excel.

(In Microsoft Excel 2007 or later, click the Office button and choose Microsoft Excel Options, then Formulas, and then check Enable Iterative Calculation in the Calculation Options group.)

If Crystal Ball detects a circular reference and Iteration is not checked, this error is displayed:

- Example message: "Circular reference detected at cell [Book1.xls] Sheet1!A1. (Code #5523)"
- Workaround: Stop the simulation and check Iteration on the Tools, then Options, then Calculation tab.

In Extreme speed, circular references with short Iteration setting may not match Microsoft Excel's values because of differences in calculation algorithms. For most consistent results, set Iteration to at least 1,000.

However, if a circular reference is non-converging, its results can differ greatly when run at both Extreme and Normal speed, regardless of the Iteration setting. If a circular reference is nonconverging, this error message is displayed:

- Example message: "Unable to run in Extreme speed for the following reason: Circular references do not converge, results cannot be guaranteed to match Normal speed. To bypass this message, turn off "Stop on calculation errors" in the Run Preferences dialog. (Code #5545)"
- Workaround: There is no workaround. Check the formulas in the workbook that have created this circular reference and look for a problem that keeps this circular reference from converging to a single value.

Simulations with circular references run in non-vectorized mode. For this reason, they will probably run more slowly than simulations without circular references.

Crystal Ball Microsoft Excel Functions

The following Crystal Ball spreadsheet functions are handled normally:

- CB.IterationsFN
- distribution functions (such as CB.Binomial)

These functions are not supported in Extreme speed during a simulation:

CB.GetForeStatFN

- CB.GetForePercentFN
- CB.GetRunPrefsFN
- CB.GetAssumPercentFN
- CB.GetCertaintyFN

While running in Extreme speed, all values for these functions return #VALUE. At the end of the simulation, Crystal Ball performs a final recalculation on the model so these functions are evaluated properly. Normally, this should not present a problem unless one of these functions was defined as a forecast and you are expecting a valid value to be computed during the simulation. If one of these "Get" functions feeds into a forecast during a simulation, this is not compatible with Extreme speed:

- Example message: "Unsupported Microsoft Excel or Crystal Ball function at cell location [Book1.xls]Sheet1!A1. (Code #5539)"
- Workaround: Defining forecasts on statistical functions that are dependent on other forecasts is generally not a good modeling practice. If you need to have a forecast cell defined on a statistical result from another forecast, use the Auto Extract feature for the dependent forecast instead of using one of the above Crystal Ball functions.

User-Defined Functions

Calls to user-defined or third-party functions are allowed. The functions may be written in Visual Basic or they may be in XLL or COM Automation DLL libraries that have been opened in Microsoft Excel.

This section includes the following topics:

- "Pure Functions" on page 296
- "Range Arguments" on page 297
- "Volatile Functions and Array Arguments" on page 297

Pure Functions

To be compatible with Extreme Speed, user-defined functions must be "pure." A "pure function" is one that computes its value solely on the basis of values passed to it as arguments. A function that is not "pure" might reference global data not passed as an argument. For example, it might get the value of a worksheet cell or a defined name and use this as an input to compute the function value. If the global data depends on the assumptions — if, for example, it is a worksheet cell with a formula computed from the assumptions — it will have a distribution of values in Normal Speed, but it will appear deterministic (have a single value) in Extreme Speed. This is because the worksheet cells change on every trial in Normal Speed, but they do not change in Extreme Speed.

Range Arguments

Range arguments in user-defined functions are only compatible with Extreme Speed when they are handled as Variant types. For example, for a function called on the worksheet as =MyFunc(A1:E4, 5, 4):

Volatile Functions and Array Arguments

User-defined functions whose arguments are static (their values do not change during a simulation) are not called by Extreme Speed unless the Volatile property of the function has been set.

When a user-defined function is encountered in Extreme Speed, Crystal Ball first checks to see if the function is Volatile. If Crystal Ball is denied access to the VBA project and the user-defined function is not passing array or cell range arguments, Crystal Ball treats the function as Volatile.

If access to the project is denied and the call is passing array or cell range arguments, the following message is displayed:

Cannot interpret a user-defined function with array arguments. You must first check the 'Trust Access to Visual Basic Project' checkbox in Microsoft Excel's macro security settings dialog. See Appendix C in the User's Guide for more information about this error.

Cause: The Microsoft Excel Volatile property has not been set.

- In that case, in order for Extreme Speed to determine the state of the Volatile property in Microsoft Excel 2003, you must first make the following setting in Microsoft Excel:
- 1 Go to the Tools, then Options panel.
- 2 Click the Security tab.
- 3 Click the Macro Security button.
- 4 Click the Trusted Publishers tab.
- 5 Check Trust Access To Visual Basic Project.

In Microsoft Excel 2007 or later, you should not encounter this specific problem with array arguments or any others involving the Analysis ToolPak, since the ToolPak is embedded directly into Microsoft Excel 2007 or later.

- ➤ However, to set the "Trust access..." property in Microsoft Excel 2007 or later to ensure that the Volatile property is set correctly, follow these steps:
- Click the Office button.

- Click the following sequence of buttons and links: Microsoft Excel Options, then Trust Center, then Trust Center Settings, then Macro Settings.
- On the Macro Settings page, under Developer Macro Settings, select Trust access to the VBA project object model.

Running User-Defined Macros

In Extreme speed, it is not possible to run user-defined macros — such as CBBeforeTrial, CBAfterTrial, and CBAfterRecalc — during a simulation. Before- and after-simulation macros are allowed, such as CBBeforeSimulation and CBAfterSimulation. If the former group of simulation macros are present, they will be flagged as incompatible:

- Example message: "One or more user-defined simulation macros were detected (e.g., CBBeforeTrial, CBAfterTrial, etc.). (Code #5701)"
- Workaround: Use Normal speed to run these user-defined macros during a simulation.

Special Functions

A small group of Microsoft Excel functions are not supported in Extreme speed: CALL, CELL, GETPIVOTDATA, INFO, HYPERLINK, REGISTER.ID, and the CUBE functions added in Microsoft Excel 2007 or later (CUBEMEMBER, CUBEVALUE, CUBESET, CUBESETCOUNT, CUBERANKEDMEMBER, CUBEMEMBERPROPERTY, CUBEKPIMEMBER). Any forecast formulas that contain one or more of these functions will be flagged as incompatible:

- Example message: "Unsupported Microsoft Excel or Crystal Ball function at cell location [Book1.xls]Sheet1!A1. (Code #5539)
- Workaround: Avoid using these functions if you want to run the model in Extreme speed.

Undocumented Behavior of Standard Functions

Microsoft Excel allows certain undocumented argument constructs for standard functions. Some of these constructs that were previously unsupported in Extreme speed are now allowed, for example:

```
=SUMPRODUCT(A1:A10*B1:B10)
=SUMPRODUCT(A1:A10/B1:B10)
```

However, some of these argument constructs are still not supported in Extreme speed and will be flagged as incompatible:

- Example message: "Unsupported Microsoft Excel or Crystal Ball function at cell location [Book1.xls]Sheet1!A1. (Code #5539)"
- Workaround: For best results, always use standard syntax with fully valid arguments.

Incompatible Range Constructs

The following sections discuss Microsoft Excel range constructs that are not supported in Extreme speed:

- "Dynamic Ranges" on page 299
- "Labels in Formulas That Are Not Defined Names" on page 299
- "Multiple Area References" on page 299
- "3-D References" on page 300

Dynamic Ranges

Extreme speed does not support dynamic ranges, where the OFFSET function is used on one or both sides of the range ":" constructor. For example,

=AVERAGE(Cellname1:OFFSET(Cellname2, x, y)).

- Example message: "Unable to interpret formula at cell location [Book1.xls] Sheet1!A1. (Code #5504 - Invalid token)"
- Workaround: Avoid using the OFFSET function to construct dynamic ranges.

Labels in Formulas That Are Not Defined Names

Extreme speed supports defined names and their use in formulas, but not the Microsoft Excel "Accept labels in formulas" option which allows cell labels to be used in formulas without defining them as names:

- Example message: "Unable to interpret formula at cell location [Book1.xls] Sheet1!A1. (Code #5514 - Undefined identifier)"
- Workaround: For best results, use defined names instead of cell labels in formulas.

Multiple Area References

Extreme speed does not support multiple-area references such as (A1:A5,B1,C1:E1) except as used in standard functions such as SUM that accept a variable-length argument list of cell ranges. The only supported use of the comma is as the separator in the argument list of a standard function, not as a cell-range "union operator." A defined name whose value is a multiple-area reference is not accepted:

- Example message: "Multi-area reference not supported at cell location [Book1.xls] Sheet1! A1. (Code #5525)"
- Workaround: For best results, avoid using multiple-area references when defining names or as arguments to any functions except those that accept a variable-length, commaseparated list of cell ranges.

3-D References

Extreme speed does not support 3-D references, where a cell range – used as an argument in a function call, for example – spans multiple workbooks:

- Example message: "Unable to interpret formula at cell location [Book1.xls]Sheet1!A1. ("Code #5514 - Undefined identifier")
- Workaround: For best results, avoid using 3-D cell references.

Data Tables

While previous versions of Crystal Ball did not support the Microsoft Excel TABLE function, this data table feature is now supported in Extreme speed. You can use one- and two-variable data tables in models that will run in Extreme speed.

Numerical Differences and the Compare Run Modes Tool


Ordinarily, there may be slight numerical differences in simulation results between Extreme and Normal speed modes. These differences are primarily due to rounding errors and are usually in the relative range of 1e-12 to 1e-15. In certain circumstances, as for extremely large models containing thousands of formulas, these slight differences can compound and grow in magnitude.

If you are concerned about possible differences in the own model, Crystal Ball provides a special Compare Run Modes tool that you can use to compare the results between the two run modes.

- To use this tool:
- Open and click the model to test.
- Choose Run, then More Tools, and then Compare Run Modes.

The Compare Run Modes dialog opens.

Figure 133 Compare Run Modes dialog



- If you are using Crystal Ball with Microsoft Excel 2007 or later, choose whether to compare results between Extreme and Normal speed or between single-threaded calculation and multi-threaded calculation.
- Indicate the amount of difference to detect, whether that difference is absolute or relative, and the number of trials to run. Check the box To compare assumptions as well as forecasts.
 - Depending on the size of the model, differences might not become obvious until after a fairly large number of trials have run. You might want to run 5,000 or more trials for the test.
- When you are ready, click **Run Comparison**.

If you are comparing speed results, the simulation runs once in Extreme speed and once in Normal speed. Otherwise, the simulation runs once with single-threaded and once with multi-threaded calculations. Results are displayed in a new workbook. The comparison Summary tab is displayed when the comparison is complete.

Figure 134 shows comparison results for the Tolerance Analysis.xls example file with 5,000 trials. In this case there were no differences in results and the model ran 28.8 times faster in Extreme speed.

Comparison of Run Modes Workbook: Tolerance Analysis.xls Results comparison: Total differences detected 20096 Total cells checked: 7 8 9 10 0.00% Percentage differences to total: Average difference (RMSE): 0 (Relative) Maximum difference 12 13 Speed comparison: 14 Extreme speed time: 3.27 seconds 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 94.27 seconds Normal speed time: Speed difference: 28.8 x **Speed Comparison** Extreme 0.00 20.00 40.00 60.00 80.00 100.00 seconds

Figure 134 Comparison results for Tolerance Analysis.xls, 5,000 trials

Because of variations in the random number seeds, you might see differences in comparison results if you use Microsoft Excel's RAND or Crystal Ball's probability functions (such as CB.Uniform) in the model.

Other Important Differences

■ Normal ■ Extreme

This section discusses several other differences between Extreme speed and Normal speed:

- "OptQuest and other Crystal Ball Tools" on page 302
- "Precision Control and Cell Error Checking" on page 303
- "Spreadsheet Updating" on page 303
- "Very Large Models" on page 303
- "Memory Usage" on page 304
- "Spreadsheets with No Crystal Ball Data" on page 304

OptQuest and other Crystal Ball Tools

OptQuest and certain other Crystal Ball tools run multiple simulations of the spreadsheet model back-to-back. Since Extreme speed enhances the performance of each simulation, you will see

35

an overall speed improvement when using one of these tools. All of the compatibility and other differences apply to these tools as well.

Precision Control and Cell Error Checking

Extreme speed uses a bursting feature to optimize the balance between speed and system responsiveness. Since bursts are usually several hundreds or thousands of trials at a time, this feature affects a number of features like precision control, confidence testing, and checking for cell errors.

Crystal Ball tests for precision control (and performs confidence testing in OptQuest) after each burst of several hundred or thousands of trials. In Normal speed, these checks are performed every 50 trials by default. As a result, you may see the simulations run more trials than usual when these features are turned on. Generally speaking, these features are not as critical to simulation performance in Extreme speed as they are in Normal speed.

Because of bursting, Crystal Ball cannot check cell errors after every trial as it does in Normal speed, only at the end of a burst. If cell errors occur in forecast cells (such as "Divide by zero" or "#VALUE") and the Run Preferences, then Stop On Calculation Errors option is selected, the simulation stops at the end of the next burst.


Spreadsheet Updating

While running simulations in Extreme speed, cells in the spreadsheet model appear to update every second or so after each burst. This is normal and does not affect simulation results. Singlestepping mode still updates cell values for every step.

Very Large Models

Very large models (with greater than 10,000 formulas, for example), might take Crystal Ball many seconds or minutes in Extreme speed to analyze the model and convert it into a form for high-speed simulation. When this happens, the following progress dialog opens.

Figure 135 Extreme Speed Progress Dialog Box



If you choose, you can cancel the simulation by pressing the Ctrl+Break keys or wait for the process to finish.

Memory Usage

Running simulations in Extreme speed is extremely memory-intensive. Depending on the machine, you may encounter issues with memory usage when running long simulations or working with large models.

During a long simulation (greater than 100,000 trials), you might also see that the machine's virtual memory is being used. For example, you might notice that the hard drive is being accessed continually and the simulation starts to slow down. No message is given when this condition occurs. If such a slow-down occurs and it is absolutely necessary that additional trials be run, you can try to disable the saving of assumption values (see the Run Preferences, then Options tab) and then restart the simulation.

If you run an Extreme Speed simulation multiple times on a large or complex model, you might find that you receive the "Sufficient memory may not..." warning dialog. This dialog lets you disable the Store Assumption Values run preference and change the number of trials to run. Notice that these changes in the warning dialog reset those preferences in the Run Preferences dialog on the Trials and Options tabs. They are then applied to other models unless you reset them after the large model has run.

Spreadsheets with No Crystal Ball Data

Spreadsheets without Crystal Ball data always run in Normal speed even if Extreme speed is selected in the Run Preferences dialog.

This could happen if you have Crystal Ball spreadsheet functions entered in the model but have not defined assumptions, forecasts, or decision variables.

Maximizing the Benefits of Extreme Speed

To get maximum benefit from running Crystal Ball in Extreme speed, you should be aware of several issues in the models that could potentially cause simulations to run more slowly than expected.

When you first start a simulation in Extreme speed, the model is analyzed and converted into a special form for high-speed processing. This form involves the use of a technique called vectorization to process batches of formula results together for maximum efficiency. Certain functions and formula constructs in the model can interfere with this vectorization process. Some of the more common cases are discussed in the following sections:

- "String Intermediate Results in Formulas" on page 305
- "Calls to User-Defined Functions" on page 305
- "Dynamic Assumptions" on page 305
- "Microsoft Excel Functions" on page 306

If one or more cases are detected, Crystal Ball Decision Optimizer displays the message "nonvectored" on the Microsoft Excel status bar while the simulation is running. Also, additional

information is usually written to the Crystal Ball log file, which can be found with the suffix CB Log and extension .txt in:

C:\Documents and Settings\username\Application Data\Oracle\ Crystal Ball\11.0\Logs\CB $Log_x.txt$

Note: If you have a very large number of trials, the virtual memory feature of Microsoft Windows is activated. This can slow down processing. Too, if you have large numbers of forecasts and assumptions, managing these large arrays of values can take extra time. If you think array management could be an issue, try unchecking Store Assumption Values For Sensitivity Analysis on the Options tab of the Run Preferences dialog.

String Intermediate Results in Formulas

In Microsoft Excel, it is possible to construct formulas that produce an intermediate character string result, which is used in later formulas to produce a numeric result. For maximum speed gains, you should avoid the use of strings in the formulas. Strings in formulas will not be flagged as incompatible, but will cause simulations to run much slower depending on the number of string formulas in the model.

Workaround: Some simple strings like "Yes" and "No" can be replaced with 1 and 0 values, or True and False Boolean values.

Calls to User-Defined Functions

Calls to user-defined or third-party functions cannot be vectorized and will cause simulations to run slower.

Workaround: There is no workaround for this case other than minimizing the use of these functions.

Dynamic Assumptions

Assumptions that contain cell references to parameter values that can change during a simulation are called dynamic assumptions. Crystal Ball needs to perform additional validation checks on the parameters for these assumption types during a simulation, causing simulations to run more slowly.

Workaround: If the dynamic assumption feature is not needed, make sure that parameter cell references do not point to other assumption cells in the model.

Cell references in custom distributions are treated as static instead of dynamic when a model is run at Extreme speed.

Microsoft Excel Functions

Crystal Ball supports LOOKUP and OFFSET functions in both Normal and Extreme speed. However, if you use more than a thousand in a model, this can slow down the initial parsing time and cause the simulation to run slowly.

C

Crystal Ball Tutorials

In This Appendix

Introduction	307
Tutorial 1 — Futura Apartments	307
Tutorial 2 — Vision Research	315
Improving Process Quality – Tutorial 3	334
Packaging Pump Design — Tutorial 4	348

Introduction

This appendix presents the basics you need to understand how Crystal Ball can be used to analyze uncertainty in two financial settings, a process improvement setting, and a product design setting.

- "Tutorial 1 Futura Apartments" on page 307 is ready to run so you can quickly see how Crystal Ball works. If you work regularly with statistics and forecasting techniques, this might be all the introduction you need before running the own spreadsheets with Crystal Ball.
- "Tutorial 2 Vision Research" on page 315 teaches more about defining and running simulations and interpreting their results.
- "Improving Process Quality Tutorial 3" on page 334 shows howCrystal Ball supports process improvement initiatives similar to Six Sigma.
- "Packaging Pump Design Tutorial 4" on page 348 shows how Crystal Ball supports
 quality initiatives for product and system design, such as Design for Six Sigma (DFSS).

Tutorial 1 — Futura Apartments

This tutorial contains the following sections:

- "Starting Crystal Ball" on page 308
- "Opening the Example Model" on page 309
- "Running Simulations" on page 310
- "Results Analysis Determining Profit" on page 311
- "Take a Look Behind the Scenes" on page 312
- "Resetting and Single-Stepping" on page 313

"Tutorial Review" on page 314

Starting Crystal Ball

With Microsoft Excel closed or open, choose Start, then All Programs, then Oracle Crystal Ball, then Oracle Crystal Ball.

If the Crystal BallWelcome screen opens, click Use Crystal Ball.

Crystal Ball opens and launches Microsoft Excel. If Microsoft Excel is already open, Crystal Ball opens in a new Microsoft Excel window.

The following sections describe the Crystal Ball menus and toolbar:

- "Crystal Ball Menus" on page 308
- "Crystal Ball Toolbar" on page 308

Crystal Ball Menus

When you load Crystal Ball with Microsoft Excel, some new menus are displayed in the Microsoft Excel menubar: Define, Run, and Analyze.

You can use these menus to define, run, and analyze simulations in Crystal Ball. You can also use commands on Microsoft Excel's Help menu to display online help, documentation, and example models for Crystal Ball.

Crystal Ball Toolbar

The Crystal Ball toolbar provides instant access to the most commonly used menu commands. Each section of the toolbar corresponds to a menu. When you move the cursor over a toolbar button, the command name is displayed.

Note: If you are using Crystal Ball in Microsoft Excel 2007 or later, you will see a single Crystal Ball tab and then a ribbon instead of a toolbar. For more information, see Appendix B, "Crystal Ball Ribbon in Microsoft Excel 2007 or Later" on page 402.

Figure 136 Crystal Ball Toolbar



To turn the Crystal Ball toolbar off for the current session, choose View, then Toolbars, then Oracle Crystal Ball.

For more information about the menus and toolbar, see Appendix F, "Accessibility".

(If you are using Microsoft Excel 2007 or later, see "Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2007 or Later" on page 402.)

Opening the Example Model

Open the Futura Apartments workbook (Futura Apartments.xls) from the Crystal Ball Examples folder.

You can find this folder by choosing Help, then Crystal Ball, and then Examples Guide. You can also choose Start, then All Programs, then Oracle Crystal Ball, then Examples in the Windows taskbar.

(In Microsoft Excel 2007 or later, choose Resources, then Examples Guide in the Help group.)

The workbook opens as shown in Figure 137.




Figure 137 **Futura Apartments workbook**

All example models included with Crystal Ball include these two worksheets:

- A Model tab with the spreadsheet model
- A Description tab with information about the model

For a list of the models included with Crystal Ball, choose Help, then Crystal Ball, then Examples Guide in the Microsoft Excel menubar or choose Start, then All Programs, then Oracle Crystal Ball, then Examples and open the Examples Guide. (In Microsoft Excel 2007 or later, choose Resources, then Examples Guide.)

The Futura Apartments Model Scenario

In this example, you are a potential purchaser of the Futura Apartments complex. You have created Futura Apartments.xls to reflect these assumptions:

- \$500 per month is the going rent for the area.
- The number of units rented during any given month will be somewhere between 30 and 40.
- Operating costs will average around \$15,000 per month for the entire complex, but might vary slightly from month to month.

You want to know how profitable the apartment complex will be for various combinations of rented units and operating costs. This is difficult to determine with a traditional spreadsheet model, but with Crystal Ball, this kind of analysis is easy.

For this tutorial, the simulation is already set up for you. You just need to run it using the Crystal Ball default settings.

Running Simulations

To run the simulation, choose Run, then Start [Simulation].



Crystal Ball runs a simulation for the situation in the Futura Apartments workbook and displays a forecast chart as it calculates the results.




By default, the simulation stops automatically after it has run for 1,000 trials. For larger models, you can use the Stop button or choose Run, then Stop [Simulation] if it is necessary to stop the simulation before all trials have run.

When the simulation stops, the forecast window opens, as in Figure 138. The numbers will vary slightly each time the simulation is run, but the forecast window should look similar to this illustration.

The forecast chart shows the total range of profit or loss outcomes predicted for the Futura Apartments scenario. Each bar on the chart represents the likelihood, or probability, of earning a given income. The cluster of columns near the center indicates that the most likely income level is between \$2000 and \$4000 per month. Note that there is also a small chance of losing almost \$2000 per month (the lower-value end of the display range) and a small chance of making about a \$7,000 gain.

Figure 138 Futura Apartments profit/loss forecast



Notice that the probability, or certainty, of a value falling within the range of negative infinity and positive infinity is 100 percent. Notice too that the chart shows 1000 trials were run but only 998 are displayed. The excluded values, if any, are extreme values that are included in calculations but are not included in the forecast chart.

Note: If the forecast window disappears behind Microsoft Excel's window during a simulation, you can bring it back to the front by clicking the Crystal Ball icon in the Windows task Crystal Ball In Microsoft Excel 2007 or later, you can choose Analyze, then View Charts, then Forecast Charts.

Results Analysis — Determining Profit

Now, you can use Crystal Ball to determine the likelihood of making a profit.


- To determine the statistical likelihood of making a profit:
- Select the first (left) certainty field in the forecast window.
- Type **0** in the field.



Press Enter.

The value in the Certainty field changes to reflect the probability of making a profit reaching an income level ranging from \$0 to positive infinity. This information puts you in a much better position to decide whether to purchase the Futura Apartments. Figure 139 shows the chance of making a profit is about 90%.

Figure 139 Chance of profit




Take a Look Behind the Scenes

While powerful results are displayed in this example with virtually no effort, obviously there must be some drivers in the process. Crystal Ball cannot generate the same results for any typical spreadsheet without some help.

The key is using Crystal Ball to define certain input cells of the spreadsheet as assumptions and certain output cells of interest as forecasts.

After these cells are defined, Crystal Ball uses Monte Carlo simulation to model the complexity of a real-world scenario.

For each trial of a simulation, Crystal Ball repeats the following three steps:



For every assumption cell, a random number is generated according to the range you defined and then is placed into the spreadsheet.

- 2. The spreadsheet is recalculated.
- 3. A value is retrieved from every forecast cell and added to the chart in the forecast windows.

This is an iterative process that continues until either:

- The simulation reaches a stopping criterion
- You stop the simulation manually

The final forecast chart reflects the combined uncertainty of the assumption cells on the model's output. Keep in mind that Monte Carlo simulation can only approximate a real-world situation. When you build and simulate the own spreadsheet models, you need to carefully examine the nature of the problem and continually refine the models until they approximate the situation as closely as possible. See "Crystal Ball Cells in the Example Model" on page 313 for an example.

Crystal Ball Cells in the Example Model

The Futura Apartments model has two assumption cells and a forecast cell. They were already defined before you ran the simulation:

- Cell C5 defines the assumption about occupancy that units rented each month will vary between 30 and 40.
- Cell C7 defines the assumption about operating costs that they will average around \$15,000 per month but might vary slightly.
- Cell C9 defines the simulation forecast (the results). If you highlight cell C9, you can see it contains a formula that references cells C5 and C7.

By default, assumption cells are green and forecast cells are blue. For each trial of the simulation, the values within these cells change as the worksheet is recalculated.

To see this process close-up, reset the model and run it again in single-step mode. You can use the Crystal Ball Control Panel for these procedures.

Resetting and Single-Stepping

When you first run a simulation, the Crystal Ball Control Panel opens. After it opens, you will find it convenient to use for managing simulations and analyzing results.

Note: If the Control Panel or any other Crystal Ball windows disappear behind Microsoft Excel, click the Crystal Ball icon in the Windows task bar to display them again,



Figure 140 Crystal Ball Control Panel



For information about the Crystal Ball Control Panel menus, see "The Crystal Ball Control Panel menubar" in Chapter 4 of the Oracle Crystal Ball User's Guide.

 \triangleright To reset the simulation and clear all previous calculations, click the Reset button, \blacktriangleleft .



To single-step through the simulation one trial at a time, click the Single-step button,



Notice that the values in the assumption and forecast cells change each time you click the Singlestep button.

Closing Crystal Ball

This completes Tutorial 1. You can save and close Crystal Ball models the same as any other Microsoft Excel workbook.

If you want, you can click the Reset button or choose Run, then Reset [Simulation] to reset the model before you close Crystal Ball.

- To close Crystal Ball, either:
- Right-click the Crystal Ball icon in the Windows taskbar and choose Close, or
- Close Microsoft Excel.

Tutorial Review

In this tutorial, you have:

- Opened Crystal Ball.
- Used the Crystal Ball menus, toolbar, and Control Panel to run an example model.
- Observed how Crystal Ball assumption and forecast cells change while a simulation runs.
- Closed Crystal Ball.

For background information about risk, risk analysis, models, and Monte Carlo simulation, see Chapter 2.

"Tutorial 2 — Vision Research" on page 315, shows how to define assumption and forecast cells and gives more suggestions for analyzing the results.

Tutorial 2 — Vision Research

This tutorial contains the following sections for review and instruction:

- "Starting Crystal Ball" on page 315
- "Opening the Example Model" on page 315
- "Reviewing the Vision Research Scenario" on page 316
- "Defining Assumptions" on page 316
- "Defining Forecasts" on page 326
- "Running Simulations" on page 328
- "Interpreting the Results" on page 329
- "Closing Crystal Ball" on page 334

Starting Crystal Ball

If Crystal Ball is not already started, follow the instructions in "Starting Crystal Ball" on page 308.

Opening the Example Model


Den the Vision Research workbook (Vision Research.xls) from the Crystal Ball Examples folder.

You can find this folder by choosing Help, then Crystal Ball, then Examples Guide. You can also choose Start, then All Programs, then Oracle Crystal Ball, then Examples in the Windows taskbar.

(In Microsoft Excel 2007 or later, choose Resources, then Examples Guide.)

The Vision Research workbook for the ClearView project opens, as in Figure 141.

Figure 141 Vision Research's ClearView Project Workbook



This spreadsheet models the problem that Vision Research is trying to solve.

Reviewing the Vision Research Scenario

The Vision Research spreadsheet models a business situation filled with uncertainty. Vision Research has completed preliminary development of a new drug, code-named ClearView, that corrects nearsightedness. This revolutionary new product could be completely developed and tested in time for release next year if the FDA approves the product. Although the drug works well for some patients, the overall success rate is marginal, and Vision Research is uncertain whether the FDA will approve the product.

You begin the analysis by defining assumption cells to support this scenario.

Defining Assumptions

In Crystal Ball, you define an assumption for a value cell by choosing a probability distribution that describes the uncertainty of the data in the cell. To accomplish this, you choose among the distribution types in the Distribution Gallery (see Figure 142).

This part of the tutorial will help you understand how to choose a distribution type. For more information about choosing distributions, see Appendix A, "Selecting and Using Probability Distributions."

You need to define or review these assumptions:

"Testing Costs Assumption: Uniform Distribution" on page 317

- "Marketing Costs Assumption: Triangular Distribution" on page 319
- "Patients Cured Assumption: Binomial Distribution" on page 320
- "Growth Rate Assumption: Custom Distribution" on page 321
- "Market Penetration Assumption: Normal Distribution" on page 324

Testing Costs Assumption: Uniform Distribution

So far, Vision Research has spent \$10,000,000 developing ClearView and expects to spend an additional \$3,000,000 to \$5,000,000 to test it based on the cost of previous tests. For this variable, "testing costs," Vision Research thinks that any value between \$3,000,000 and \$5,000,000 has an equal chance of being the actual cost of testing.

The uniform distribution describes a situation where all values between the minimum and maximum values are equally likely to occur, so this distribution best describes the cost of testing ClearView.


- To define the assumption cell for testing costs:
- Click cell C5.
- Choose Define, then Define Assumption.



Because an assumption has not yet been defined in cell C5, the Distribution Gallery dialog opens as shown in Figure 142.

By default, the Basic distributions are displayed. These are six of the most frequently used continuous and discrete distributions. When you click a distribution to choose it, information about that distribution is displayed at the bottom of the Distribution Gallery.

Figure 142 Distribution Gallery Dialog




Note: In Microsoft Excel 2007 or later, if you click the upper half of the Define Assumption icon or if the assumption has already been defined, the Distribution Gallery opens. If you click the lower half of the Define Assumption icon, a list of All, Basic, or Favorite distributions opens, depending on the active category in the Distribution Gallery.

- 3 Click the Uniform distribution.
- 4 Click OK.

The Uniform Distribution dialog opens (Figure 143).

Since cell C5 already has label text beside it on the worksheet, that text is displayed in the Assumption Name field. Use this name, rather than typing a new one. Also, notice that Crystal Ball assigns default values to the distribution parameters, Minimum and Maximum.

Figure 143 Uniform Distribution for C5



Vision Research expects to spend a minimum of \$3,000,000 and a maximum of \$5,000,000 on testing. Use these values in place of the defaults to specify the parameters of the uniform distribution in Crystal Ball, as described in the following steps.

- To specify parameters:
- 1 Type 3 in the Minimum field (remember that the numbers on the worksheet represent millions of dollars).

 This represents \$3,000,000, the minimum amount Vision Research estimates for testing costs.
- 2 Press Tab.
- 3 Type 5 in the Maximum field.

This represents \$5,000,000, the maximum estimate for testing costs.

4 Click Enter.

The distribution changes to reflect the values you entered, as shown in Figure 144.

Define Assumption: Cell C5 Edit View Parameters Preferences Help Name: Testing Costs ×. (8) Uniform Distribution \$3.0 \$3.2 \$3.4 \$3.6 \$3.8 \$4.0 \$4.2 \$4.4 \$4.6 \$4.8 \$5.0 κ. ٧, Maximum \$5.0 Minimum \$3.0 Cancel Enter OK Gallery Correlate. Help

Figure 144 Changed distribution values

With the values from step 1 and step 3 entered correctly, the distribution looks like Figure 144. Later, when you run the simulation, Crystal Ball generates random values for cell C5 that are evenly spread between 3 and 5 million dollars.

Click OK to return to the worksheet.

The assumption cell is now green.

Marketing Costs Assumption: Triangular Distribution


Vision Research plans to spend a sizeable amount marketing ClearView if the FDA approves it. Including sales commissions and advertising costs, Vision Research expects to spend between \$12,000,000 and \$18,000,000, with a most likely amount of \$16,000,000.

Vision Research chooses the triangular distribution to describe marketing costs because the triangular distribution describes a situation where you can estimate the minimum, maximum, and most likely values to occur. This assumption is already defined for you.

- To examine the assumption cell for marketing costs:
- 1 Click cell C6.
- Choose **Define**, then **Define Assumption**.

The Triangular Distribution dialog (Figure 145) opens for cell C6.

Figure 145 Triangular Distribution for Cell C6



The triangular distribution has three parameters — Minimum (\$12 million), Likeliest (\$16 million), and Maximum (\$18 million).

When you run the simulation, Crystal Ball generates random values that fall around 16, with fewer values near 12 and 18.

Click OK to return to the worksheet.

Patients Cured Assumption: Binomial Distribution

For FDA approval of ClearView, Vision Research must conduct a controlled test on a sample of 100 patients for one year. Vision Research expects that the FDA will grant an approval if 20% or more of the patients tested are cured (show corrected vision) after taking ClearView for one year. Vision Research is encouraged by their preliminary testing, which shows a success rate of around 25%.


Vision Research chooses the binomial distribution to describe the uncertainties in this situation because the binomial distribution describes the random number of successes (25) in a fixed number of trials (100).

This assumption is already defined.

- To examine the assumption cell for patients cured, use the following steps:
- Click cell C10.
- Choose Define, then Define Assumption.

The Binomial Distribution dialog opens as shown in Figure 146.

Figure 146 Binomial Distribution Dialog



The binomial distribution has two parameters: Probability and Trials. Because Vision Research experienced a 25% success rate during preliminary testing, the Probability parameter is set to 0.25 to show the likelihood of success.

Note: You can express probabilities either as decimals between 0 and 1, such as 0.03, or as whole numbers followed by the percent sign, such as 3%.

Because the FDA expects Vision Research to test 100 people, the Trials parameter is set to 100. When you run the simulation, Crystal Ball generates random integers between 0 and 100, simulating the number of patients that would be cured in the FDA test.

Click OK to return to the worksheet.

Growth Rate Assumption: Custom Distribution

Vision Research has determined that nearsightedness afflicts nearly 40 million people in the United States, and an additional 0% to 5% of these people will develop this condition during the year in which ClearView is tested.

However, the marketing department has learned that a 25% chance exists that a competing product will be released on the market soon. This product would decrease ClearView's potential market by 5% to 15%.

Since the uncertainties in this situation require a unique approach, Vision Research chooses Crystal Ball's custom distribution to define the growth rate.

The method for specifying parameters in the custom distribution is quite unlike the other distribution types, so follow the directions carefully. If you make a mistake, click Gallery to return to the distribution gallery, then start again at step 4.

Use the custom distribution to plot both the potential increase and decrease of ClearView's market.

- To define the assumption cell for the growth rate of nearsightedness:
- Click cell C15.
- Choose Define, then Define Assumption.


(In Microsoft Excel 2007 or later, click the upper half of the Define Assumption icon.) The Distribution Gallery dialog opens.

- Click All in the navigation pane of the Distribution Gallery to show all distributions shipped with Crystal
- Scroll down to the end of the Distribution Gallery and click the Custom distribution.
- 5 Click OK.

The Custom Distribution dialog opens.

Notice in Figure 147 that the chart area remains empty until you specify the Parameters type and enter the values for the distribution.


Figure 147 Custom Distribution Dialog



You know that you will be working with two distribution ranges: one showing growth in nearsightedness and one showing the effects of competition. Both ranges are continuous.

Open the Parameters menu (Figure 148).

Figure 148 Custom Distribution Parameters Menu



Choose Continuous Ranges in the Parameters menu.


The Custom Distribution dialog now has three parameters: Minimum, Maximum, and Probability.

- 8 Enter the first range of values to show the growth of nearsightedness with low probability of competitive effects:
 - Type 0% in the Minimum field.
 - This represents a 0% increase in the potential market.
 - b. Type 5% in the Maximum field.
 - This represents a 5% increase in the potential market.
 - c. Type 75% or .75 in the Probability field.
 - This represents the 75% chance that Vision Research's competitor will not enter the market and reduce Vision Research's share.
 - d. Click Enter.

A uniform distribution for the range 0% to 5% is displayed (Figure 149.

Figure 149 **Uniform Distribution Range**

Uniform distribution for the first range of values.



Notice that the total area of the range is equal to the probability: 5% wide by 15 units high equals 75%.

Now, enter a second range of values to show the effect of competition:

a. Type -15% in the Minimum field.

This represents a 15% decrease in the potential market.

b. Type -5% in the Maximum field.

This represents a 5% decrease in the potential market.


c. Type 25% in the Probability field.

This represents the 25% chance that Vision Research's competitor will enter the market place and decrease Vision Research's share by 5% to 15%.

d. Click Enter.

A uniform distribution for the range -15% to -5% is displayed. Both ranges are now displayed in the Custom Distribution dialog (Figure 150).

Figure 150 Customized Uniform Distribution



Notice that the area of the second range is also equal to its probability: $2.5 \times 10\% = 25\%$.

10 Click OK to return to the worksheet.

When you run the simulation, Crystal Ball generates random values within the two ranges according to the probabilities you specified.

Market Penetration Assumption: Normal Distribution

The marketing department estimates that Vision Research's eventual share of the total market for the product will be normally distributed around a mean value of 8% with a standard deviation of 2%. "Normally distributed" means that Vision Research expects to see the familiar bell-shaped curve with about 68% of all possible values for market penetration falling between one standard deviation below the mean value and one standard deviation above the mean value, or between 6% and 10%.

In addition, the marketing department estimates a minimum market of 5%, given the interest shown in the product during preliminary testing.

Vision Research chooses the normal distribution to describe the variable "market penetration."


- To define the assumption cell for market penetration:
- Click cell C19.
- Choose Define, then Define Assumption.
- In the Distribution Gallery, click the normal distribution. 3

(Scroll up to the top of the All category or click Basic to immediately display the normal distribution.)

Click OK.

The Normal Distribution dialog opens (Figure 151).

Normal Distribution for Cell C19 Figure 151



- Specify the parameters for the normal distribution: the mean and the standard deviation.
 - a. If the Mean field doesn't contain 8.00%, type 8% in the Mean field. This represents an estimated average for market penetration of 8%.
 - b. Type 2% in the Std. Dev. field. This represents an estimated 2% standard deviation from the mean.
- Click Enter.

The normal distribution scales to fit the chart area, so the shape of the distribution does not change. However, the scale of percentages on the chart axis does change.

Click the More button, , to display additional fields (Figure 152).

Figure 152 Assumption Truncation Fields



These fields, marked by gray arrows, display the minimum and maximum values of the assumption range. If values are entered into them, they cut or truncate the range. These fields are then called the truncation minimum and maximum.

Type 5% in the minimum truncation field (the first or left field).

This represents 5%, the minimum market for the product.

Click Enter.

The distribution changes to reflect the values you entered (Figure 153).




Figure 153 **Changed Distribution for the Truncated Values**

When you run the simulation, Crystal Ball generates random values that follow a normal distribution around the mean value of 8%, and with no values generated below the 5% minimum limit.

10 Click OK to return to the worksheet.

Defining Forecasts

Now that you have defined the assumption cells in the model, you are ready to define the forecast cells. Forecast cells contain formulas that refer to one or more assumption cells.

The president of Vision Research would like to know both the likelihood of achieving a profit on the product and the most likely profit, regardless of cost. These forecasts are displayed in the gross profit (cell C21) and net profit (cell C23) for the ClearView project.

You can define both the gross profit and net profit formulas as forecast cells, described in the following sections:

- "Gross Profit Forecast" on page 327
- "Net Profit Forecast" on page 327

Gross Profit Forecast

First, look at the contents of the cell for gross profit:

Click cell C21.

The cell contents are displayed in the formula bar near the top of the worksheet. The contents are C16*C19*C20. Crystal Ball uses this formula to calculate gross profit by multiplying Persons With Nearsightedness After One Year (C16) by Market Penetration (C19) by Profit Per Customer (C20).

Now that you understand the gross profit formula, you are ready to define the forecast cell for gross profit.

To define this forecast cell:

2 Choose Define, then Define Forecast.



The Define Forecast dialog opens as shown in Figure 154. You can enter a name for the forecast. By default, the forecast cell label is displayed as the forecast name.

Figure 154 Define Forecast dialog—Gross Profit If Approved



Use the forecast name that is displayed, rather than typing a new name.

- 3 Since the spreadsheet model involves millions of dollars, type Millions in the Units field.
- Click **OK** to return to the worksheet.

Net Profit Forecast

- ▶ Before defining the forecast cell formula for net profit, look at the contents of the cell for net profit:
- 1 Click cell C23.

The contents are displayed in the Microsoft Excel formula bar. The contents are IF(C11,C21-C7,-C4-C5).

The formula translates as follows:

If the FDA approves the drug (C11 is true), then calculate net profit by subtracting total costs (C7) from gross profit (C21). However, if the FDA does not approve the drug, (C11 is false), then calculate net profit by deducting both development costs (C4) and testing costs (C5) incurred to date.

To define the forecast cell for net profit:

Choose Define, then Define Forecast.



The Define Forecast dialog opens.

Again, use the forecast name that is displayed in the Forecast Name field and specify millions in the Units field.

Click OK to return to the worksheet.

You have defined assumptions and forecast cells for the Vision Research spreadsheet, and are now ready to run a simulation.

Setting Run Preferences

- To specify the number of trials and initial seed value (so the charts will look like this tutorial):
- Choose Run, then Run Preferences, then Trials.



The Run Preferences Trials dialog opens. For this example, we will run more trials to get more accurate forecast results.

- In the Number Of Trials To Run field, type 5000.
- Click Sampling.
- Choose Use Same Sequence Of Random Numbers.
- In the Initial Seed Value field, type 999.
- Notice there are two choices in the Sampling Methods group: Monte Carlo and Latin Hypercube. Latin Hypercube is less random and produces a smoother, more even results chart. For now, though, choose the default - Monte Carlo.
- Click OK.

Running Simulations

When you run a simulation in Crystal Ball, you have the freedom to stop and then continue the simulation at any time. Buttons for Run, Stop, and Continue are displayed on the Crystal Ball toolbar and, after you have started to run a simulation, they are displayed on the Crystal Ball Control Panel.

To run the simulation, click Run,



Interpreting the Results

Now that you have run the simulation, you are ready to interpret the forecast results. Should Vision Research scrap the ClearView project or proceed to develop and market this revolutionary new drug? For the answer, review the forecast charts.

The following sections describe how to analyze results for this scenario:

- "Reviewing the Net Profit Forecast Chart" on page 329
- "Determining the Net Profit Certainty Level" on page 330
- "Customizing Forecast Charts" on page 332

Note: Crystal Ball windows are separate from Microsoft Excel windows. If Crystal Ball's windows or charts disappear from the screen, they are usually simply behind the main Microsoft Excel window. To bring them to the front, click the Crystal Ball icon in the Windows taskbar or press Alt+Tab and choose Crystal Ball, C Crystal Ball


Note: To display forecast charts in Microsoft Excel 2007 or later, choose Analyze, then View Charts, then Forecast Charts.

Reviewing the Net Profit Forecast Chart

Forecast charts are frequency distributions. They show the number or frequency of values occurring in a given bin or group interval and show how the frequencies are distributed. In Figure 155, the bin that contains the most values has a frequency of about 280.

Crystal Ball forecasts the entire range of results for the Vision Research project. However, the forecast charts do not display the most extreme values. Here, the display range includes values from approximately -\$15 million to \$38 million.

Figure 155 The Net Profit Forecast



The forecast chart also shows the certainty range for the forecast. By default, the certainty range includes all values from negative infinity to positive infinity.

Crystal Ball compares the number of values in the certainty range with the number of values in the entire range to calculate the certainty level.

The previous example shows a certainty level of 100%, since the initial certainty range includes all possible values. Remember that the certainty level is an approximation, since a simulation can only approximate the elements of the real world.

Determining the Net Profit Certainty Level


The Vision Research president wants to know how certain Vision Research can be of achieving a profit and what are the chances of a loss.

- To determine the certainty level of a specific value range:
- In the Net Profit forecast chart, type 0 in the certainty range minimum field.
- 2 Press Enter.

Crystal Ball moves the lower-value (left) certainty grabber to the break-even value of \$0.0 and recalculates the certainty level.

Analyzing the Net Profit forecast chart again (Figure 156), you can see that the value range between the certainty grabbers shows a certainty level of about 79%. That means that Vision Research can be 79% certain of achieving a net profit. You can therefore calculate a 21% chance of suffering a net loss (100% minus 79%).

Figure 156 Net Profit Forecast-\$0 Minimum




Now, the president of Vision Research would like to know the certainty of achieving a minimum profit of \$2,000,000. With Crystal Ball, you can easily answer this question.

- To do so:
- Type 2 in the certainty range minimum field.
- Press Enter.

As Figure 157 shows, Crystal Ball moves the lower-value (left) certainty grabber to \$2.0 and recalculates the certainty level. Vision Research can be about 72% certain of achieving a minimum net profit of \$2,000,000.

Figure 157 **Recalculated Certainty Level**




Vision Research is encouraged by the forecast result. The president now wants to know how certain Vision Research can be of achieving a minimum net profit of \$4,000,000. If Crystal Ball shows that Vision Research can be at least two-thirds certain of a \$4,000,000 net profit, the president is ready to go ahead with the ClearView project.

- Again, Crystal Ball can easily answer this question:
- Type 4 in the range minimum field.
- Press Enter.

Crystal Ball moves the lower-value (left) certainty grabber to \$4.0 and recalculates the certainty level.

The Net Profit forecast chart in Figure 158 shows a certainty level of almost 66%. With virtually two-thirds certainty of a minimum net profit of \$4,000,000, Vision Research decides to go ahead with the ClearView project and proceed to develop and market this revolutionary new drug.



Net Profit Forecast-\$4.0 Million Minimum Figure 158

You can work with the Gross Profit chart in similar ways.

Customizing Forecast Charts

Crystal Ball charts are useful for presenting results as well as analyzing them. A variety of chart preferences are available to show different chart views, types, colors, and more.

To display the chart preferences, choose Preferences, then Chart in the forecast chart window. You can click the Help button to learn more about the preferences on each tab.

You can also use keyboard shortcuts to customize the appearance of charts without using the Chart Preferences dialog.

- Select a forecast chart and try these shortcuts now:
- Press Ctrl+t to cycle through chart types (area, line, and column).
- Press Ctrl+d to change the chart views (frequency, cumulative frequency, reverse cumulative frequency).
- Press Ctrl+m to cycle through a series of markers that show the mean and other measures of central tendency.
- Press Ctrl+p to cycle through a series of percentile markers.
- Press Ctrl+b to change the density of the chart by varying the number of bins.
- Press Ctrl+w to display the chart in 3D.
- Drag the edges of the chart window until the proportions and size suit the presentation plans.

For example, Figure 163 shows the ClearView gross profit forecast presented as a 3-D area chart with a mean marker, stretched to show more detail in the x-axis. The chart is also set for 50% transparency. The certainty fields have been set to show that the certainty of a gross profit greater than \$32 million is about 78%.





Figure 159 Customized Gross Profit Chart

You can choose Edit, then Copy Chart in the chart menubar to copy a chart to the clipboard for pasting into Microsoft Excel or another application. Figure 160 shows the ClearView gross profit chart pasted into a presentation slide.

Figure 160 Gross Profit Chart Pasted into a Presentation Slide



Closing Crystal Ball

To close Crystal Ball, exit Microsoft Excel.

Summary

In this tutorial, you started Crystal Ball, opened the tutorial example model, defined assumptions and forecasts, ran the simulation, and reviewed and customized forecast charts. By experimenting with certainty values, you explored a few questions that Vision Research executives might ask as they analyzed the results of the simulation.

As you read through the Oracle Crystal Ball User's Guide, you will learn more about customizing forecast charts, creating other kinds of charts, interpreting the descriptive statistics, and printing comprehensive reports for any simulation. Crystal Ball provides these features so that you can quantify the risk inherent in the assumptions.

Crystal Ball keeps the assumption and forecast definitions (but not the forecast values) with the spreadsheet. When you save the spreadsheet, the definitions are saved with it. To learn about saving and restoring forecast results, see "Saving and Restoring Simulation Results" on page 88.

The following section, "Improving Process Quality — Tutorial 3" on page 334, shows how to use Crystal Ball in the Analyze and Improve phases of Six Sigma and other process quality improvement methodologies.

Improving Process Quality — Tutorial 3

Six Sigma and other process improvement methodologies strive to save costs and time by minimizing process variation and reducing rework. Six Sigma has five phases for gathering information, analyzing it, and implementing suggested improvements: Define, Measure, Analyze, Improve, Control. The following sections in this tutorial show how Crystal Ball can be used in the Analyze and Improve phases to discover inconsistencies and reduce variation in a bank's loan processing system:

- "Overall Approach" on page 335
- "Using the Loan Processing Model" on page 337

"Packaging Pump Design — Tutorial 4" on page 348, shows how to use Crystal Ball in design quality programs such as Design for Six Sigma (DFSS).

Overall Approach

Suppose you are assigned the task of improving loan processing quality at a bank. You want to minimize average processing time, keeping below the bank's target. At the same time, you want to evaluate and improve processing time consistency using the following techniques:

- Define You plan to talk with stakeholders in various departments to learn how they are currently working, how they view of the problem, and what outcome they hope to see.
- Measure You will determine current processing times and the amount of variation within them.
- Analyze You plan to use Crystal Ball to help you discover process capability, or how closely the processing times map to process requirements.
- Improve You will use information gathered in the previous phases to design and implement improvements in processing time and consistency.
- Control You need to ensure permanent improvements by setting up a system of feedback loops to continuously measure process results and change the output if measures exceed desired control targets and control limits.

In the Define and Measure phases, you obtain and use information to build a Crystal Ball model. Then, you use Crystal Ball to simulate real-world variations that occur in processing loans. Finally, you can use sensitivity analysis and further optimization to arrive at the best balance between time and consistency.

The following sections give details of each phase:

- "Define Phase" on page 335
- "Measure Phase" on page 336
- "Analyze Phase" on page 336
- "Improve Phase" on page 336
- "Control Phase" on page 337

Define Phase

The bank's loan processing procedures involve six steps. You interview staff involved with each step and learn that marketing and management groups are hoping for completion of all six steps within 96 working hours. Based on these interviews, you establish both the target and upper

specification limit (USL) at 96 hours. There is no lower specification limit (LSL), although consistency (the sigma level) seems to be an issue. You decide to obtain more concrete data.

Measure Phase

You document time and time variation data for each processing step, as follows:

- Customer Inquiry involves the creation of an initial rate quote following the first customer contact by phone, in-person visit, or the Internet. Historical data indicates this cycle time is lognormally distributed with a mean of one hour and a standard deviation of 0.25 hours.
- Loan Application involves the completion of forms by customers, with or without staff assistance. This step takes from one to five eight-hour days, but usually takes three days.
- Document Verification and processing occurs in the bank, where a loan specialist reviews credit data, presents loan alternatives to customers, and independently verifies information. This step usually takes between two and four eight-hour days. However, about two out of ten times, the loan is suspended and takes between four and six days for additional reviews and documentation.
- Underwriting, or loan approval, usually takes from one hour to 8 hours. All lengths of time have the same likelihood of occurrence.
- Loan Closing includes the preparation of the final documentation, locking in an interest rate, and arranging where to deposit funds. Historical data show that this step has a mean of two days, but you don't yet know how the data are distributed.
- Loan Disbursement, moving the funds to the customer's bank, usually takes 16 hours, with some variation. Two-thirds of disbursement times fall within 12 and 20 hours.

After gathering this information, you decide to load it into Microsoft Excel and use Crystal Ball to analyze it.

Analyze Phase

During the Analyze phase, you create a Microsoft Excel workbook that includes historical data on loan closing times. Then, you can capture the time and variation information for each step as a Crystal Ball assumption. Finally, you can define overall processing time as a Crystal Ball forecast. When the simulation runs, you can view capability metrics to gain insight into the degree of process variation, or inconsistency. You can also generate a sensitivity chart to learn more about the main sources of inconsistency.

"Using the Loan Processing Model" on page 337, shows how to use a Crystal Ball model in the Analyze and Improve phases to accomplish the quality goals.

Improve Phase

You can experiment with the effects of changing the main sources of inconsistency. If you have OptQuest, you can also define decision variables and optimize for the shortest processing time and least time variation.

Control Phase

You set up a system of measuring cycle times for each step and analyzing them on a regular basis. You run a Crystal Ball simulation on real data once a month with a program of corrective actions to take whenever results exceed desired limits by a specified amount.

Using the Loan Processing Model

This section shows how to construct a model that effectively simulates each step of the loan process with dual goals of minimizing processing time and maximizing time consistency. Then, you run the model and analyze the output using Crystal Ball's process capability features.

The following sections describe how to work with the loan processing model and analyze data from it:

- "Introduction" on page 307
- "Activating the Process Capability Features" on page 337
- "Setting the Sampling Seed Value" on page 338
- "Opening the Example Model" on page 315
- "Reviewing the Vision Research Scenario" on page 316
- "Reviewing the Assumptions" on page 339
- "Reviewing the Forecast" on page 341
- "Running Simulations" on page 342
- "Analyzing Simulations" on page 342
- "Investigating Improvement Possibilities" on page 344

Starting Crystal Ball

If Crystal Ball is not already started, start it as usual ("Starting Crystal Ball" on page 308). If the Crystal Ball Welcome screen opens, check Quality and then click Use Crystal Ball.

Activating the Process Capability Features

- To confirm that the process capability features have been activated and activate them, if necessary:
- Choose Run, then Run Preferences.



- Click the Statistics tab in the Run Preferences dialog.
- **Check Calculate Capability Metrics.**
- Click the Options button to confirm the settings.

In the Capability Options panel, set the options to use short-term metrics with a Z-score shift value of 1.5 (the default Short Vs. Long Term settings) and use the default Calculation Method settings.

Click OK in both dialogs to accept the current settings and close them.

Setting the Sampling Seed Value

To more closely reproduce the results in this tutorial, you might want to set the sampling seed value on the Sampling tab of the Run Preferences dialog ("Setting Run Preferences" on page 328).

Opening the Example Model

Open Loan Process.xls from the Crystal Ball Examples folder. To choose this example from a list, choose Help, then Crystal Ball, then Examples Guide or choose Start, then All Programs, then Oracle Crystal Ball, then Examples in the Windows taskbar.

(In Microsoft Excel 2007 or later, choose Resources, then Examples Guide in the Help group.) The workbook opens, as shown in Figure 161.




Figure 161 Loan Processing Example Model

Reviewing the Parts of the Model

The loan processing spreadsheet models a process improvement effort with two basic goals: meeting the cycle time target as closely as possible and maximizing the sigma level. The model has these parts:

- The flowchart at the top of the model illustrates the process with approximate measured mean times for each step.
- The blue box beneath the flowchart highlights the cycle time target.
- The green boxes in the Simulated Cycle Time column are Crystal Ball assumptions that represent each step of the loan processing effort.
- The assumption parameters beside each assumption show the distribution and the parameters entered for each assumption. They are entered in each assumption as cell references relative to cells in column C.
- The blue box at the bottom of the Simulated Cycle Time column is a Crystal Ball forecast that represents the total cycle time.

Reviewing the Assumptions

The assumption definitions are based on the measurements obtained in the second Six Sigma phase, described in "Measure Phase" on page 336. For example, Step 1, Customer Inquiry, has a mean of 1 hour and a standard deviation of 0.25 hour as discussed in "Measure Phase" on page 336. You can compare the other assumptions with the descriptions in "Measure Phase" on page 336 to see how they match.

Notice that Step 3 is defined with a custom distribution. We will examine it in more detail shortly and then define the distribution in Step 5.

Examining the Custom Distribution

Because the measurement for Step 3 involves two separate sets of values, it is defined as a custom distribution.

- To explore this:
- Select cell C18, the assumption for Step 3.
- Choose Define, then Define Assumption to display the Define Assumption dialog for that assumption.




Figure 162 shows the expanded Define Assumption dialog for Step 3, Document



Verification. The dialog has been expanded by clicking the More button.

This dialog reflects the information from "Measure Phase" on page 336 that 80% of the time, cycle time for this step ranges from two to four days (16 to 32 hours) but 20% of the time, cycle time ranges from four to six days (32 to 48 hours).

Figure 162 Custom Assumption for Step 3, Document Verification



Click OK or Cancel to close the dialog.

Fitting a Distribution for Step 5

When you first open Loan Process.xls, no assumption has been defined for Step 5, Loan Closing. The description in "Measure Phase" on page 336 states that it takes an average of two days (16 hours) to close loans but the distribution is not known.

- To create an assumption for Step 5 and fit a distribution to it:
- 1 Select cell C22.
- Choose Define, then Define Assumption to display the Distribution Gallery dialog for that cell.



Click the Fit button.

The Fit Distribution dialog opens.


Type the range A2:A101 or click the cell selector for the Range edit box. Then, click the Data tab of the workbook and select cells A2 through A101.

When you accept the selection, the Fit Distribution dialog opens with the range entered.

Leave the other default settings as shown and click OK.

The Comparison Chart dialog opens (Figure 163. The normal distribution appears to be the best fit.

Figure 163 Comparison Chart Dialog, Cell C22




6 Click Accept to accept the normal distribution.

The Define Assumption dialog opens.

Figure 164 shows that the new assumption for Step 5 is a normal distribution with a mean of 16 and a standard deviation of 4.

Figure 164 Define Assumption Dialog, Step 5, Cell C22



7 Rename the assumption to "Loan Closing" and click OK.

Now, all assumptions for the Loan Process model are complete.

Reviewing the Forecast

The forecast in cell C26 represents total cycle time. To view its formula in the Microsoft Excel formula bar, select cell C26. The formula is =SUM(C14:C24) — the sum of the assumptions for Steps 1 through 6.

To view and edit the forecast definition:




Choose Define, then Define Forecast or click the toolbar button.

The Define Forecast dialog opens.

Figure 165, following, shows that the forecast has a name, a units entry, and a USL, the upper specification limit of 96 hours. These were entered when the example model was first defined.

Figure 165 Cycle Time Forecast with LSL and Target Entered



As established in the Define phase of this project, 96 hours is the loan processing target as well as the USL.

Enter 96 in the Target edit box and click OK.

Running Simulations

At this point, the model's assumptions and forecast are complete.

To run the simulation and generate a forecast chart, click the Start [Simulation] button



Analyzing Simulations

When you run the simulation for this model, the Cycle Time forecast chart opens (Figure 166, following). Because the forecast is defined with an upper specification limit and a target, the chart is displayed in Split View with the forecast chart beside the process capability metrics. The message with the metrics indicates that the distribution is non-normal (the normality test failed), so capability metrics are calculated from the forecast values.




Figure 166 The Cycle Time Forecast with Process Capability Metrics


Notice that marker lines are displayed for the USL and target. Because the USL is defined, the maximum certainty grabber is automatically set to that value. Because no LSL is defined, the minimum certainty grabber is automatically set to –Infinity. The Certainty field contains 72.31; about 72% of forecasted values fall below the USL.

Looking at the process capability metrics, the forecast distribution is not normal, so the Z scores are not generally appropriate to use. The p(N/C)-above (defects above the USL) is about .28, which confirms the certainty level shown in the forecast chart; more than one-quarter of the forecasted values fall above the USL. The Cpk-upper and Cpk are about 0.21, where you had hoped to see at least 1.00 (the 3-sigma level). The loan specialists you interviewed were correct — cycle time variation is large.

You generate a sensitivity chart (Figure 167) to determine which of the loan processing steps has the most influence on the Cycle Time forecast variance. To do this, select the forecast chart and choose Forecast, then Open Sensitivity Chart.

Figure 167 shows that the loan processing step with the greatest effect on cycle time variance is Document Verification (Step 3). Obviously, this is an ideal target for improvement.

Figure 167 Sensitivity Chart for the Cycle Time Forecast



Investigating Improvement Possibilities

You wonder how much you would need to reduce variation in Document Verification time before overall quality would improve significantly. You decide to increase the probability of a cycle time of four days or less.


- To do this:
- Click the Reset button to reset the simulation.



- 2 Type a new probability in cell G18. Change it to 90%.
- 3 Adjust cell G19 downward to match. Change it to 10%.
- Select cell C18 and open the Define Assumption dialog.

As shown in Figure 168, the two parts of this assumption are both uniform in height over their range, although the first, lower-value one is nine times taller.

Figure 168 Document Verification Assumption, Adjusted



This assumption was previously displayed in Figure 162. Because the probabilities in the table are linked to workbook cells, when you adjusted the workbook, you also adjusted the assumption probabilities.

- Close the Define assumption dialog and run the simulation.
- 6 Review the forecast chart in Figure 169.

-Infinity

Porecast: Cycle Time Edit View Forecast Preferences Help 1,000 Trials Split View 995 Displayed Cycle Time Normality test failed (p-value = 0.04) Metric Forecast values -60 Mean 88 0.06 Standard Deviati 11 50 Ср 0.05 Cpk-lower Cpk-upper 0.2531 40 Cpk 0.2531 0.04 **Probability** 30 Z-LSL 0.03 0.7530 Z-USL 0.7530 20 7st-total 0.02 -0.7470 p(N/C)-below p(N/C)-above 0.2165 p(N/C)-total 0.2165 PPM-below PPM-above 216.548.59 90 100 110 120 PPM-total 216,548.59 hours LSI USL 96

The Cycle Time Forecast with Process Capability Metrics

Crystal Ball sets the upper certainty level at the USL/Target value (96). If the lower certainty grabber is not located at -Infinity, click the lower grabber and drag it toward the Probability axis as shown in Figure 169. Now, the certainty of cycle time falling at or below the USL is about 78%. You decide to make a drastic change and edit the assumption parameters for Document Verification as follows (change 32 to 24 in the first row; change 90% to 100% and 10% to 0%):

96

Target Z-score shift

Forecast values

Certainty: 78.35

1.50 ▼

Step 3: 29 24 100% 48

Now, there is 100% probability that all document verification times will fall between two and three days.

Figure 170 shows that when you run with this new assumption definition, the certainty of falling at or below the USL rises to about 93.5% and PPM-total drops to 62,548. Also, the distribution is now normal; it was not normal for the last simulation. Z-USL and Zst-total are 1.53, the sigma level for this simulation.




Figure 170 Cycle Time Forecast after Adjusting Document Verification

You check the sensitivity chart for this simulation (Figure 171) and find that the impact of Loan Application is now much greater than Document Verification.




Figure 171 Sensitivity Chart after Adjusting Document Verification

You believe that Document Verification has been pared down about as much as it can be, so you decide to experiment further with making the Loan Application cycle time shorter and more consistent.

- To do this:
- Reset the simulation and change the Step 2 parameters to 8, 16, and 24.
- 2 Run the simulation.

Results are much better. Figure 172 shows that the mean cycle time has dropped to 74 hours. The certainty of falling below the USL/Target value is 100% and the sigma level has risen to 3.09. Cpk-upper and Cpk have risen over 1 which, given the shape and location of the distribution, confirms that the sigma level is greater than 3.





Figure 172 Forecast Chart after Loan Application Adjustment

When you check the sensitivity chart (Figure 173), Loan Closing and Loan Disbursement are now the main contributors to cycle time variation. You learn that these processes might be difficult to improve further, so you decide to implement improvements to Document Verification and Loan Application cycle times and then move into the Control phase, discussed in "Control Phase" on page 337.

Figure 173 Sensitivity Chart after Loan Application Adjustment



Packaging Pump Design — Tutorial 4

Design for Six Sigma (DFSS) and other design improvement methodologies are intended to improve product quality at the start of product development.

DFSS has five phases: Define, Measure, Analyze, Design, Validate (or Verify). The ultimate mission of DFSS is to meet general Six Sigma goals of satisfying customer needs while minimizing variation and to accomplish this as early in the development cycle as possible.

The following sections show how DFSS techniques can be used with Crystal Ball to simulate various design options for a fluid pump in a food production setting:

- "Overall Approach" on page 348
- "Using the DFSS Liquid Pump model" on page 351
- "Using OptQuest to Optimize Quality and Cost" on page 359

Overall Approach

Suppose you are a pump manufacturer. You want to design a pump for a liquid packaging system that draws processed fluids from a vat into jars at a consistent rate. You want to meet the specified flow rate target and limits while maximizing performance and minimizing costs.

Using the phases of DFSS, you follow this overall design approach:

- Define Obtain an appropriate flow rate target and limits from the customer and recommend a type of pump based on customer considerations
- Measure Consider the elements of the Known Flow Rate equation to determine the pump design features that affect flow rate and obtain information about tolerances, costs, and so on
- Analyze Use Monte Carlo analysis and sensitivity analysis to study how different design elements affect pump performance and affect production quality

- Design Use optimization techniques with defined tradeoffs between flow rate performance and total cost to determine the design parameters that best meet a defined quality level while reducing cost
- Validate Based on the optimized design parameters, create physical prototypes and test to confirm that this design configuration yields the best-performing, most cost-effective pump

In the Define and Measure phases, you obtain and use information to build a Crystal Ball model. Then, you use Crystal Ball to simulate real-world variations that occur in the actual manufacturing of the product. Finally, if you have Crystal Ball Decision Optimizer, you can use sensitivity analysis and further optimization to arrive at the best balance between quality and cost.

The following sections give details of each phase:

- "Define Phase" on page 349
- "Measure Phase" on page 350
- "Analyze Phase" on page 350
- "Design Phase" on page 350
- "Validate Phase" on page 351

The next main section, "Using the DFSS Liquid Pump model" on page 351, works through these phases using a Crystal Ball example model.

Define Phase


Marketing and production considerations determine the acceptable range of flow rates based on how many units should be produced daily to meet market demand, what is useful production time per day based on plant setup and productivity, and how many pumps are placed in each plant.

The food processing company translates marketing requirements into these engineering specifications and provides them to you:

- Lower specification limit (LSL) = 47.26 ml/sec
- Upper specification limit (USL) = 53.92 ml/sec
- Target = 50.595 ml/sec

Based on these requirements, you propose a displacement pump with a reciprocating piston. Figure 174, following, shows a basic schematic of the system.

Figure 174 The jar-filling system with displacement pump



Measure Phase

You notice that the variables in the Known Flow Rate equation are piston radius, piston stroke length, backflow rate, and motor speed. You observe that the in-house machine shop controls the piston radius and stroke length while the backflow valve and motor are purchased components.

the machine shop indicates it is capable of 3-sigma quality levels within a tolerance of ± 1 millimeter (mm) for the piston radius and stroke length, the purchasing department asks you to use an inexpensive backflow valve and motor.

Analyze Phase

During the Analyze phase, you create a Microsoft Excel workbook that captures the Known Flow Rate equation, characteristics of the pump components used in the equation (piston radius, piston stroke length, motor speed, and backflow rate), flow rate target and limits, machining cost coefficients, plus nine motor and backflow valve cost options.

Then, you define Crystal Ball assumption and forecast cells and simulate the design options.

Design Phase

If you have OptQuest, you can also define decision variables for:

- A range of possible piston radii and associated tolerances (related to standard deviation through capability)
- A range of possible piston stroke lengths and associated tolerances,
- The different possible backflow valve and motor options

After this is done, you can set up quality requirements for flow rate and optimize for the leastcost configuration that meets those quality requirements.

Validate Phase

This phase involves the physical creation of a prototype based on simulation forecasts followed by physical testing to validate the design. Because so much preliminary work was done with spreadsheet modeling instead of building additional prototypes, a final design can be developed quickly and at minimal cost.

Using the DFSS Liquid Pump model

Now, see how to construct a model that effectively simulates the design process with dual goals of maximizing quality and minimizing cost. Then, you run the model and analyze the output using Crystal Ball's process capability features.

This tutorial contains the following review and lesson sections:

- "Starting Crystal Ball" on page 351
- "Activating the Process Capability Features" on page 351
- "Setting the Sampling Seed Value" on page 352
- "Opening the Example Model" on page 352
- "Reviewing the Parts of the Model" on page 353
- "Running Simulations" on page 356
- "Analyzing Simulations" on page 356
- "Adjusting Models" on page 358

If you have Crystal Ball Decision Optimizer, see "Using OptQuest to Optimize Quality and Cost" on page 359 for a useful optimization example.

Starting Crystal Ball

If Crystal Ball is not already started, start it as usual ("Starting Crystal Ball" on page 308).

Activating the Process Capability Features

- To confirm that the process capability features have been activated.
- 1 Choose Run, then Run Preferences.



- 2 Click the Statistics tab in the Run Preferences dialog.
- Check Calculate Capability Metrics.
- Click the Options button to confirm the settings.

For this example, use the defaults: use short-term metrics, calculate metrics from the normal distribution with a significance level of 0.05, and, if the normality test fails, calculate metrics from the forecast values.

Click OK in both dialogs to accept the current settings and close them.

Setting the Sampling Seed Value

Set the sampling seed value to 999 in the Run Preferences dialog, Sampling tab, and use Monte Carlo simulation.


Opening the Example Model

Open DFSS Fluid Pump.xls from the Crystal Ball Examples folder. To choose this example from a list, choose Help, then Crystal Ball, then Examples Guide in Microsoft Excel or choose Start, then All Programs, then Oracle Crystal Ball, then Examples in the Windows taskbar.

(In Microsoft Excel 2007 or later, choose Resources, then Examples Guide.)

The workbook opens, as shown in Figure 175.

Figure 175 Fluid Pump Model



This spreadsheet models the fluid pump design problem to be solved.

Reviewing the Parts of the Model

The liquid pump spreadsheet models a design effort with two basic goals: meeting the flow rate target as closely and consistently as possible and minimizing cost. The following sections describe six main parts of the model, indicated by numbers in Figure 175:

- "Known Flow Rate Formula" on page 354
- "Flow Rate Variable Assumptions" on page 354
- "Flow Rate and Cost Forecasts" on page 354
- "Flow Rate Target and Limits" on page 355
- "Cost Coefficients" on page 355

"Motor and Backflow Valve Cost Options" on page 356

Known Flow Rate Formula



The Known Flow Rate formula specifies the relevant variables discussed earlier: piston radius, piston stroke length, backflow rate, and motor speed.

Flow Rate Variable Assumptions



The flow rate variable assumptions are Crystal Ball assumptions that allow each of the flow rate variables, listed under the Known Flow Rate formula (1), to be included in the Monte Carlo simulation.

- To see how these are defined:
- Select cell K18 and click the Define Assumption button.



Optional: Look at K19, K20, and K21.

Each assumption is defined as a normal distribution with a mean equal to the Nominal Initial Value and a standard deviation equal to the StDev Initial Value. The Mean and Standard Deviation parameters are not entered directly. Instead, they are entered as cell references to the Initial Value cells.

Note: In this model, whenever you see the word "nominal," it means approximate, designed, or theoretical, as opposed to the actual measured value.

Flow Rate and Cost Forecasts

The flow rate and cost forecasts are Crystal Ball forecasts, the output of the Monte Carlo simulation.



The Flow Rate Forecast is based on the Known Flow Rate formula (1).

- To view it:
- Select cell K23.

The formula in that cell is the Known Flow Rate formula (a function of the flow rate variables) expressed in Microsoft Excel format.

Click the Define Forecast button.



Notice that the values in cells E24, E25, and E23 are displayed in the Define Forecast dialog as the process capability lower specification limit (LSL), upper specification limit (USL), and target value, respectively.

They are defined as cell references. To demonstrate this, click in one of the text boxes. The related cell reference is displayed. (Note: If you don't see these text boxes in the Define Forecast dialog, process capability features have not been activated in the Run Preferences dialog.)

The Total Cost Forecast is the sum of the Component Cost Function column values in the upper table. To view it:

Select cell K25.

Notice its formula is simply the sum of the component costs in cells J18 through J21. No target or limits have been established for this forecast.

Flow Rate Target and Limits



The flow rate target and upper and lower specification limits (4) are values defined by the food processing plant in the Define phase of the DFSS process. The previous section shows how they can be used to set the Target, LSL, and USL for the Flow Rate Forecast in cell K23.

Cost Coefficients



The cost coefficients table contains values used to calculate the component cost functions, summed to yield the Total Cost Forecast (3).

- To see how these are used to calculate the component cost functions:
- 1 Select cell J18.

This formula uses the Base Cost Coefficient of the piston, the initial value of the piston radius, the "Tolerance" Cost Coefficient of the piston, and the initial value of the standard deviation. The cost is directly proportional to the radius squared due to raw material usage and inversely proportional to the standard deviation squared (a common method to relate tolerance to cost).

2 Optional: View the contents of cells J19, J20, and J21.

These formulas are also similarly derived from cost coefficients or nominal costs relevant to each particular flow rate variable.

Motor and Backflow Valve Cost Options



The motor and backflow valve cost options are two sets of nine options each provided by the motor and backflow valve vendors. Different performance specifications have different costs associated with them.

Running Simulations

At this point you have created assumptions and forecasts.

To run the simulation and generate forecast charts, click the Start [Simulation] button



Analyzing Simulations

When you run the simulation for this model, two forecast charts are displayed.

The Total Cost Forecast (Figure 176, following) includes no assumption cells. It is displayed as a single value instead of a range of predicted values. The forecasted total cost is \$26.73 for the pump (shown here with the Mean marker line displayed).



Figure 176 The Total Cost Forecast Chart

Because you entered a target and specification limits, the Flow Rate Forecast chart opens in Split View with the forecast chart first and then the process capability metrics, shown in Figure 177.

Forecast: FLOW RATE FORECAST (ml / sec) Edit View Forecast Preferences Help 1,000 Trials 999 Displayed FLOW RATE FORECAST (ml/sec) Normality test passed (p-value = 0.710) 49.3412 Standard Deviation 1.2737 52 0.05 0.8723 48 Cpk-lowe 0.5442 44 Cpk-upper 1 20 0.5442 40 Cpk 0.04 0.6217 Com 36 Probability Z-LSL 1.63 32 Z-USL 3.60 0.03 28 Zst-total 1.63 24 Q 0.1311 p(N/C)-below 0.0513 0.02 20 p(N/C)-above 1.583e-04 16 p(N/C)-total 0.0514 12 PPM-below 51,280,16 0.01 SL = 47.2617 PPM-above 158.30 PPM-total 51 438 45 LSL 47.2617 0.00 USL 53.9283 48.1572 49.2274 53.5080 50.2975 51.3677 52.4379 50.5950 Target ml/sec Z-score shift Certainty: 95.34 47.2617 Infinity

Figure 177 The Flow Rate Forecast with Process Capability Metrics

Notice that marker lines are displayed for the LSL, Target, and USL values. (You might see different marker lines; the Mean and Standard Deviation marker lines have been removed to simplify the screenshot.) Because the LSL and USL are defined, the certainty grabbers are automatically set to those values. According to the value in the Certainty field for this simulation, 95.34% of forecasted values fall between 47.2617 and 53.9283, the specification limits. The mean is about one standard deviation lower than the target.

The distribution curve is normal. Looking at the process capability metrics, the Z-LSL shows the LSL is 1.63 sigmas below the mean. You were hoping to see a value of at least 3. The Cp capability index is lower than 1, so the short-term potential sigma level is less than 3.

You generate a sensitivity chart (Figure 178) to determine which of the flow rate variables has the most influence on the Flow Rate Forecast.




Figure 178 Sensitivity Chart for the Flow Rate Forecast

Figure 167 shows that the flow rate variable with the greatest effect on the Flow Rate Forecast is the piston radius. Reducing the standard deviation of the piston radius should affect the overall flow rate variance and improve the process capability metrics. You wonder how reducing the piston standard deviation by 50% would affect the flow rate forecast quality. You decide to adjust the model and find out.

Adjusting Models

- To see the effect of reducing the standard deviation of the piston radius:
- Select cell H18.
- Reduce the value by half, to 0.16667.
- 3 Enter this same value in the Upper Limit cell, I18.

Now, the table of flow rate variables is displayed as shown in Figure 179.


Figure 179 Flow Rate Variable Initial Values and Limits

		Nominal			StDev			Component	
	Component	Lower	Initial	Upper	Lower	Initial	Upper	Cost	Sampled
	Option #	Limit	Value	Limit	Limit	Value	Limit	Function	Value
Piston Radius (mm)		25.0000	30.0000	50.0000	0.10000	0.16667	0.33333	\$11.70	30.0000
Stroke Length (mm)		25.0000	35.0000	100.0000	0.10000	0.33333	0.33333	\$18.22	36.0000
Motor Speed (rpm)	3	30.0	30.00000	40.0	0.10000	0.24000	0.24000	\$3.25	30.0000
Backflow (ml)	9	0.10	0.50000	0.50	0.00250	0.01500	0.01000	\$0.30	0.5000

Reset and run the simulation again.

The Flow Rate Forecast chart is displayed, similar to Figure 180.

Figure 180 Flow Rate Forecast Chart After Model Adjustment



In Figure 180, quality appears much better than that shown in Figure 166. The mean is less than 1.5 standard deviations lower than the target, and the standard deviation is smaller. USL is almost 6 sigmas away from the mean (Z-USL = 5.68) and LSL is over 2 sigmas away (Z-LSL = 2.49). The certainty of the forecasted flow rate falling within LSL and USL is now over 99%.

The Total Cost Forecast chart has also changed, as shown in Figure 181. The quality improvements result in higher costs. Where the total cost per pump was \$26.73, now it is \$33.47, a 25% increase. You and the team must determine whether to sacrifice quality to lower the cost or whether to test modifications to the other three flow rate variable components as well.



Figure 181 The Total Cost Forecast Chart after Model Adjustment

Using OptQuest to Optimize Quality and Cost

If you have Oracle Crystal Ball Decision Optimizer, Fusion Edition with OptQuest, an optimization tool, you can use it to maximize quality and minimize cost. To do this, you must consider the relationship between design parameters and cost. And, you must consider the relationship between motor and backflow valve options and cost.

The Total Cost Forecast is defined as the sum of the four component costs, examined in "Flow Rate Variable Assumptions" on page 354. The formulas used to calculate these values are based on the mean and standard deviation of the piston radius, the mean and standard deviation of the piston stroke length, nine different motor options, and nine different backflow value options.

To apply OptQuest to this model, follow the steps in these sections:

- "Defining Decision Variables" on page 359
- "Identifying Optimization Goals" on page 360
- "Running OptQuest" on page 361

Defining Decision Variables

To optimize these variables with OptQuest, a set of decision variables are defined in cells C20, C21, E18, E19, H18, and H19 (Figure 182).

Figure 182 Table of Flow Rate with Cells Defined as Decision Variables

	Component Option #	Lower Limit	Nominal Initial Value	Upper Limit	Lower Limit	StDev Initial Value	Upper Limit	Component Cost Function	Sampled Value
Piston Radius (mm)		25.0000	30.0000	50.0000	0.10000	0.33333	0.33333	\$8.33	30.0000
Stroke Length (mm)		25.0000	35.0000	100.0000	0.10000	0.33333	0.33333	\$14.85	36.0000
Motor Speed (rpm)	3	30.0	30.00000	40.0	0.10000	0.24000	0.24000	\$3.25	30.0000
Backflow (ml)	9	0.10	0.50000	0.50	0.00250	0.01500	0.01000	\$0.30	0.5000
Cells C20 C21 Cells F18 F19						Calls H			

- To view the contents of each decision variable, reset Crystal Ball and:
- Select a decision variable cell.
- **Choose Define, then Define Decision.**



In the Define Decision Variable dialog, there is a name, an upper and lower bound, a type setting, and a step size for discrete variables.

Note: If you changed the Piston Radius standard deviation Initial Value and Upper Limit for the last part of the DFSS Pump tutorial, be sure to change them both back to their original values as shown in Figure 182. Also, be sure that these values are included in the decision variable for the Piston Radius standard deviation (Cell H18). Check it as described in steps 1 and 2.

The upper and lower bound for piston machining variables are the same as those in the Lower Limit and Upper Limit tables of Figure 182. For the motor speed and backflow valve variables, the upper and lower bounds are the lowest and highest options (1 and 9, respectively).

In choosing the piston machining values shown in Figure 182, you could identify raw stock availability to define a range of piston radii and lengths for investigation. You could also identify a feasible range of tolerances for machining the radius and length.

Identifying Optimization Goals


If you have OptQuest, you can actually perform the optimization. First, identify optimization goals. For this project, they are:

- Minimize cost.
- Reduce variation of flow rate to at least 3 sigma levels. That is, Z-total, expressed in this example as Zst, should be equal to or greater than 3.

Running OptQuest


- To optimize these goals in OptQuest, reset the simulation in Crystal Ball and:
- Open the Crystal Ball Run Preferences dialog and set Number Of Trials To Run to 1000. For this example, Sampling was set to Latin Hypercube with an Initial Seed Value of 999 (set on the Sampling tab of the Run Preferences dialog).
- Choose Run, then OptQuest. 2
- When the OptQuest wizard opens, click Objectives in the navigation pane to display the Objectives panel if it is not already onscreen.
- Confirm that the Objectives panel opens, as shown in Figure 183.

Figure 183 Forecast Optimization Defined in OptQuest



Click Next until you reach the Options panel (Figure 184, following).

Figure 184 OptQuest Options panel



Choose Run For 1000 Simulations, then click Run.

OptQuest starts looking for the best feasible solution. Ultimately a solution will be found that is similar to that shown in Figure 185.

OptQuest Results Edit View Analyze Preferences Help 685 Simulations Best Solution View Performance Chart FO S60.00 - 00.002 ··Infeasible solutions \$40.00 ◆ Last best solution \$30.00 62 125 188 251 314 377 440 503 566 629 Simulations Best Solution: Simulation # 185 Objectives Minimize the Mean of TOTAL COST FORECAST (\$) \$33.66 Requirements The Zst-total of FLOW RATE FORECAST (ml / sec) must be greater th... 3.14 Constraints Left Side Right Side Decision Variables Backflow Valve Option 3.00 Motor Option 2.00 OPT Piston Radius 33.3779 OPT Piston Radius StDev 0.24904 OPT Stroke Length 28.6045 OPT Stroke Length StDev 0.21453

Figure 185 OptQuest Solution for the Fluid Pump Problem

the optimized values should be similar to these:

Option	Optimized Value
Total cost	\$33.66
Zst (sigma level)	3.14
Backflow valve option	3
Motor option	2
Piston radius (mm)	33.38
Piston radius standard deviation	0.249
Stroke length (cm)	28.60
Stroke length standard deviation	0.214



Using the Process Capability Features

In This Appendix

Introduction	365
About Crystal Ball's Process Capability Features	365
About Quality Improvement Methodologies	366
How Crystal Ball Supports Quality	368
Working Through Process Capability Tutorials	370
Preparing to Use Process Capability Features	370
Analyzing Process Capability Results	372
Capability Metrics List	378

Introduction

This appendix describes Crystal Ball's process capability features, which support quality improvement methodologies such as Six Sigma, DFSS (Design for Six Sigma), and Lean principles.

The following sections introduce the process capability features and describe how to use them:

- "About Crystal Ball's Process Capability Features" on page 365
- "Preparing to Use Process Capability Features" on page 370
- "Analyzing Process Capability Results" on page 372
- "Capability Metrics List" on page 378

About Crystal Ball's Process Capability Features

Use the Statistics tab of the Run Preferences dialog to activate the process capability features. After they are activated, you can:

- Add upper and lower specification limits and target values to forecast definitions ("Setting Specification Limits and Targets" on page 372).
- Display the Capability Metrics view in the forecast window, in addition to the basic Frequency, Cumulative Frequency, Reverse Cumulative Frequency, Statistics, and Percentiles views ("Viewing Capability Metrics" on page 372).

- Use Crystal Ball's Split View feature to display the Capability Metrics view beside one of the forecast chart views ("Viewing Forecast Charts and Capability Metrics Together" on page 373). This is the default.
- View marker lines for specification limits and the target value on forecast charts ("Viewing LSL, USL, and Target Marker Lines" on page 375).
- Extract capability metrics to a worksheet using Analyze, then Extract Data or the Auto Extract forecast preferences ("Extracting Capability Metrics" on page 375).
- Include capability metrics in reports ("Including Capability Metrics in Reports" on page
- Include several capability metrics in OptQuest, if you have a license for it, and specify them as objectives or requirements.
- Access capability metrics in the Crystal Ball Developer Kit using these calls: CB.GetFore, CB.DefineForeND, CB.SetView, CB.GetForeStat, CB.RunPrefsND, CB.GetRunPrefs, CB.CreateRptND, and CB.ExtractDataND (see the Crystal Ball Developer Kit User's Guide for more information)

Several of these calls have equivalent spreadsheet functions and can be used in worksheet cells instead of VBA macros.

The following sections summarize process capability initiatives for quality improvement and describe how Crystal Ball's process capability features support those initiatives:

- "About Quality Improvement Methodologies" on page 366
- "How Crystal Ball Supports Quality" on page 368
- "Working Through Process Capability Tutorials" on page 370

About Quality Improvement Methodologies

Starting about 1975, competitive pressures on manufacturers encouraged them to discover new quality methodologies. Leading companies developed a variety of techniques to save costs and time by improving processes, designing more effectively, and increasing the value of product modifications.

Some of the most successful methodologies have been Six Sigma, Design for Six Sigma (DFSS), Lean principles, and similar tools for quality improvement. The following sections summarize several of these:

- "Methodologies for Improving Processes Six Sigma" on page 367
- "Methodologies for Improving Design DFSS" on page 367
- "Methodologies for Adding Value Lean Principles" on page 367

Methodologies for Improving Processes — Six Sigma

Broken processes reduce quality in a variety of organizations, from banking to manufacturing to healthcare, and more. In the United States, Six Sigma is probably the best known process improvement methodology, although other techniques have similar goals and effectiveness in the field.

Six Sigma is a methodology for improving processes to achieve any or all of the following goals:

- Reducing the defect rate to 3.4 per million (outside the range of 6 standard deviations above and below the mean)
- Creating improvements of 70% or better
- Using the Define-Measure-Analyze-Improve-Control methodology of information gathering, analysis, and action
- Producing rapid improvement of an organization's financial "bottom-line."

Six Sigma methodology is data-driven and focuses on the relationship between system inputs and outputs, both of which can vary. The inputs, or independent variables, are symbolized by X; the outputs, or dependent variables, are symbolized by Y. The goal of a Six Sigma program is to yield maximum quality and minimum variation in Y by controlling variation in all the X variables, the factors that influence it.

Six Sigma teams talk about the "hidden factory" of workers behind the scenes who rework defects and perform other extra tasks to make up for broken processes. Six Sigma programs achieve their goals by discovering the hidden factory in their organization and releasing its workers for more productive tasks.

Methodologies for Improving Design — DFSS

Design for Six Sigma (DFSS) developed following the success of Six Sigma and other process improvement methodologies. If a similar multi-phase improvement technique could be applied to design, costs and time could be saved by eliminating waste and unnecessary prototyping. Successful designs could move to production quicker, reducing costs and increasing an organization's competitive edge. DFSS proved to be a successful tool for improving design in a variety of settings. The five phases of DFSS are Define, Measure, Analyze, Design, and Verify (DMADV).

The ultimate mission of DFSS is to address general Six Sigma goals of meeting customer needs while minimizing variation and to accomplish this as early in the development cycle as possible.

Methodologies for Adding Value — Lean Principles

The Lean approach to quality focuses on eliminating waste, improving efficiency, and increasing value. Lean techniques seek to increase value by improving the flow of value stream activities and reducing waste. Other Lean principles include just-in-time delivery driven by customer pull — not push — and a constant, incremental search for perfection (sometimes called *kaizen*).

Lean practitioners use value stream mapping to identify process steps with the greatest impact on the overall product. Then, Lean projects are created to improve those processes. The "hidden factory" is a major target for Lean programs, as practitioners seek to eliminate rework and scrap. Lean principles are often used with statistical methodologies such as Six Sigma.

How Crystal Ball Supports Quality

Crystal Ball is most helpful in phases of quality improvement programs where uncertainty and variation have the greatest impact. In the Define phase, for example, cost savings forecasting, project cost estimation, and resource scheduling are all subject to uncertainty. These types of analyses are ideal targets for simulation.

In the Analyze and Improve phases, the focus is variation. Simulation can help reduce the variation around tolerances or process steps.

In any phase, defining model inputs (X's) as a distribution of values yields a more realistic range of results. Then, the results (Y's), defined as Crystal Ball forecasts, can be analyzed as demonstrated in the following chapters of this guide.

By default, the forecast certainty range is automatically set to the upper and lower specification limits (USL and LSL), if at least one of them is entered for the forecast. Then, you can tell at a glance the probability of output falling within them, 99.49% in the example in Figure 186. Split View shows capability metrics beside the forecast chart.




Figure 186 Piston Displacement Forecast Chart

Capability Metrics view indicates whether the forecast distribution is normal, to guide the interpretation of the capability metrics. For example, Z metrics are always displayed, although they are generally used only for normal distributions. The Z-total metric indicates the sigma

level of quality. In this example, the level is 2.37 sigmas. For this forecast, 8,824 units per million will be outside the specification limits. The Cp and Cpk values are less than one, which suggests that the quality level is lower than 3 sigmas. To cycle among other views, such as Statistics and Percentiles, press Ctrl+Spacebar.

Note: By default, Crystal Ball applies a 1.5 sigma Z-score shift value when calculating capability metrics. If the organization uses a different value or does not use a Z-score shift value, be sure to adjust this setting in the Capability Options panel.

Crystal Ball's sensitivity charts are also helpful to quality practitioners. They show which inputs, or assumptions, have the greatest impact on the selected forecast. For example, the following sensitivity chart (Figure 187) shows that Connecting Rod Length has the greatest effect on the Piston Displacement forecast shown in the previous figure.




Figure 187 Piston Displacement Sensitivity Chart

Crystal Ball's data extraction and reporting capabilities are also helpful in quality analysis. For example, you can automatically extract capability metrics to a prominent location in the model for immediate access to current sigma levels and other important measures.

If you have OptQuest, it — along with the Crystal Ball Developer Kit and OptQuest Developer Kit — can add further power to the arsenal of quality improvement tools.

When you define certain cells as Crystal Ball decision variables, you can then optimize them with OptQuest and copy the new, optimal values back into the Crystal Ball model. For example, "Using OptQuest to Optimize Quality and Cost" on page 359, shows how you can use OptQuest to adjust several controllable inputs simultaneously to yield a set of optimal results.

Visual Basic for Applications (VBA) developers can use the Crystal Ball Developer Kit to support quality projects in a variety of creative ways. This model uses the Developer Kit to create a colorcoded dashboard that highlights which process steps are lowest in quality and which are most important to the process output.

For other ways to apply Crystal Ball to the own quality improvement projects, review the examples in this guide and examine the other models supplied with Crystal Ball and posted online. For more information, see:

http://www.oracle.com/crystalball

Working Through Process Capability Tutorials

As you learn about the Crystal Ball process capability features, consider working through the following tutorials:

- "Improving Process Quality Tutorial 3" on page 334
- "Packaging Pump Design Tutorial 4" on page 348

They show how the process capability features can be used within quality improvement methodologies to simplify quality initiatives in a variety of settings.

Preparing to Use Process Capability Features

Before you can use the process capability features, you must activate them, set options, and enter limits and target values for at least one forecast. The following sections describe these procedures:

- "Activating the Process Capability Features" on page 370
- "Setting Capability Calculation Options" on page 370
- "Setting Specification Limits and Targets" on page 372

Activating the Process Capability Features

- To activate Crystal Ball's process capability features:
- Choose Run, then Run Preferences to display the Run Preferences dialog.
- Click the Statistics tab.
- 3 **Check Calculate Capability Metrics.**
- Click the Options button to set the capability calculation options, described in the next section.
- Click OK when setting changes are complete.

Setting Capability Calculation Options

After you activate the process capability features, you can set a variety of options to further customize these features for the situation.

- To set the process capability options:
- Display the Statistics tab of the Run Preferences dialog.
- 2 Click the Options button as described in step 5 of the previous section.
 - The Capability Options panel opens.
- Indicate whether metrics should use short-term or long-term formulas, depending on the time-span of the model.
- Optionally, specify a Z-score shift value to be used in long-term formulas.
 - You can specify a value from 0 to 100.
 - The default is 1.5. If don't want to use a Z-score shift value, enter 0 in the box or delete the current value and leave it empty.
- 5 Indicate whether metrics should be calculated from a fitted distribution or directly from the forecast values ("Calculation Method" on page 371).
- When settings are complete, click OK.

Calculation Method

By default, Crystal Ball tries to fit a normal distribution to the forecast values. You can enter a significance level to specify the threshold below which the assumption of normality is rejected. The default level of 0.05 translates into a 95% confidence that a rejection of normality will be correct. Other significance levels typically used are 0.01, 0.025, and 0.1, which translate into 99%, 97.5%, and 90% confidences, respectively.

If normality is rejected, Crystal Ball will then either calculate the metrics directly from the forecast values (the default) or, if you choose, perform a best fit to select the most appropriate continuous probability distribution from which to calculate the metrics.

The normality test and non-normal best fit (if normality is rejected) use the goodness-of-fit test and distribution selection that is set in the Forecast Window tab of the Forecast Preferences dialog (opened by choosing Preferences, then Forecast in the forecast window).

Before you choose to calculate from the best fitting distribution if the distribution is not normal, consider that:

- You are not guaranteed of attaining a good fit to the forecast values, and
- The fitting process might take a long time depending on how many simulation trials you are running.

Note: In unusual circumstances, it is possible for the normality test to fail and the best fitting distribution still be a normal distribution, or for the normality test to pass and the best fitting distribution be non-normal.

Alternatively, you can choose the second main setting, Calculate Metrics From Forecast Values, to bypass the normality test and always calculate the metrics directly from the forecast data.

Setting Specification Limits and Targets

The capability metrics only display if you specify either an upper or lower specification limit (or both) for the forecast. You can also specify an optional target.

- To specify these limits:
- Either define a new forecast or select an existing forecast and choose Define, then Define Forecast.
 - The Define Forecast dialog opens. With process capability features activated, it includes Name, Units, LSL, USL, and Target fields.
- Enter specification limits and target values for this forecast into the appropriate fields.
 - LSL = lower specification limit, USL = upper specification limit, and Target = target value for this forecast. If you prefer, you can enter cell references by typing or browsing.
 - All of these fields are optional, but Crystal Ball only calculates capability metrics if a value is entered for one or both of the specification limits.
- To set forecast preferences at the same time, click the More button, 🦭, beside the Name field.
- When all settings are complete, click OK.

For information about the relationship of the LSL and USL to the certainty range, see "Viewing LSL, USL, and Target Marker Lines" on page 375

Analyzing Process Capability Results

When capability metrics have been activated and appropriate information entered, run a Crystal Ball simulation as usual. Then, you can perform the following capability analyses:


- "Viewing Capability Metrics" on page 372
- "Extracting Capability Metrics" on page 375
- "Including Capability Metrics in Reports" on page 377

Viewing Capability Metrics

After you have defined a forecast with at least one specification limit (and, optionally, a target), you can run a simulation and display capability metrics for the forecast.

- To display capability metrics:
- Define a forecast with LSL, USL, and Target values as described in "Setting Specification Limits and Targets" on page 372.
- Run the simulation and display the chart for that forecast.
- In the forecast window, choose View, then Capability Metrics.
 - A table of metrics opens, similar to Figure 188.

Figure 188 Capability Metrics View



For a description of each statistic, see "Capability Metrics List" on page 378.

The following sections describe other ways you can analyze capability metrics visually:

- "Viewing Forecast Charts and Capability Metrics Together" on page 373
- "Viewing LSL, USL, and Target Marker Lines" on page 375


Viewing Forecast Charts and Capability Metrics Together

It can be helpful to view a forecast chart and its capability metrics side by side. This is now the default view when capability metrics are activated.

- To do this manually:
- Follow the instructions in the previous section to display capability metrics.
- 2 In the forecast window, choose View to open the View menu.
- 3 Choose Split View at the bottom of the menu.
- Open the View menu again and choose Frequency.

The Frequency chart and capability metrics both appear, similar to Figure 189.


Figure 189 Frequency Chart and Capability Metrics in Split View



Optionally, continue opening the View menu and choosing charts or data.

Figure 190 shows a Frequency chart, a Cumulative Frequency chart, plus Statistics and Capability Metrics tables.

Figure 190 Frequency and Cumulative Frequency Charts with Capability Metrics and Statistics in Split View



You can click in any of the Split View panes and use the chart shortcut keys to modify them without using the View or Preferences menus. For a list, see Table 24, "Shortcut Keys for Chart Preferences, All Supported Versions of Microsoft Excel," on page 401.

You can also resize the chart window and drag the horizontal and vertical pane splitters to resize each part of the Split View window.

Viewing LSL, USL, and Target Marker Lines

By default, after you add specification limits and a target to a forecast, markers for these values are displayed on the forecast chart. For an example, see Figure 190.

Notice that the certainty range on the forecast chart changes to match the LSL and USL values. The certainty indicates the probability of falling within those specification limits. To show the certainty of different values, type them into the Minimum and Maximum fields or click the certainty grabbers and drag them to a new position.

- To add specification limit and target marker lines manually or to remove them:
- Choose Preferences, then Chart Preferences in a forecast chart window, or double-click the chart.
- 2 Click the Chart Type tab in the Chart Preferences dialog.
- Scroll down to the bottom of the Marker Lines list.
- Be sure LSL, USL, Target is checked to display the specification limit and target marker lines on that forecast chart.
 - To hide the markers, uncheck LSL, USL, Target.
- When all settings are complete, click OK.

Notice that you can display the mean and other marker lines in addition to the specification limit and target markers.

Extracting Capability Metrics

You can extract capability metrics automatically or manually, as described in the following sections:

- "Extracting Capability Metrics Automatically" on page 375
- "Extracting Capability Metrics Manually" on page 376

Extracting Capability Metrics Automatically

- You can automatically extract capability metrics whenever a simulation runs. To extract metrics automatically:
- 1 In a forecast window, choose Preferences, then Forecast and display the Auto Extract tab of the Forecast Preferences dialog.

- Check Extract Forecast Statistics Automatically... and scroll down to the bottom of the data selection list.
- 3 Check Capability Metrics, enter a Starting Cell, and then click OK.

Now, when you run a simulation, capability metrics are written to the specified area of the worksheet, along with any other data you have requested.

Note: For more information about the Auto Extract feature, see "Auto Extract Tab" on page 70. IMPORTANT! Be sure to choose an open area of the worksheet as the Starting Cell to avoid overwriting the model.

Extracting Capability Metrics Manually

- To extract capability metrics manually after a simulation runs:
- Choose Analyze, then Extract Data to display the Extract Data Preferences dialog.
- 2 Check Capability Metrics at the bottom of the Select Data To Extract list.
- Choose appropriate Forecasts and Assumptions settings and specify locations and other preferences on the Options tab. For more information, see "Extracting Data" on page 165 or click the Help button in the Extract Data Preferences dialog.
- Click OK to extract the data.

Capability metrics are written to the specified location, along with any other data you have requested. See Figure 191 for an example.

Figure 191 Manually Extracted Capability Metrics

	Α	В	C	
1	Capability metrics	FLOW RATE FORECAST (ml / sec)	TOTAL COST FORECAST (\$)	
2	Mean	49.2453	\$26.73	
3	Standard Deviation	1.2122	\$0.00	
4	Ср	0.92		
5	Cpk-lower	0.55		
6	Cpk-upper	1.29		
7	Cpk	0.55		
8	Cpm	0.61		
9	Z-LSL	1.64		
10	Z-USL	3.86		
11	Zst-total	1.64		
12	Zlt	0.14		
13	p(N/C)-below	0.05		
14	p(N/C)-above	0.00		
15	p(N/C)-total	0.05		
16	PPM-below	50,879.49		
17	PPM-above	55.92		_
14 4	▶ ▶ Description (1	Model Data	<u> </u>	

Including Capability Metrics in Reports


- To include capability metrics in full, forecast, or custom reports:
- Choose Analyze, then Create Report to display the Create Report dialog.
- Click a report type: Full, Forecast, or Custom.

If you choose Full or Forecast, the capability metrics are displayed in a block for each forecast following the statistics and percentiles for that forecast. Additional process capability information is displayed in the summary and all selected marker lines are displayed in forecast and overlay charts.

If you choose Custom, the Custom Report dialog opens.

- If it is not already highlighted, highlight Forecasts in the Report Sections list.
 - The Forecasts Details list is displayed. With the process capability features activated, Capability Metrics is checked by default.
- If you don't want to include the capability metrics for some reason, uncheck that setting in the Forecasts Details list. Otherwise, leave it checked and follow the instructions in online help or "Defining Custom Reports" on page 164 to finish defining the custom report.
- When all settings are complete, click OK to generate the report.
 - Capability metrics are displayed with other forecast data, similar to Figure 192.

Figure 192 Frequency Report with Capability Metrics



Capability Metrics List

The following table lists and defines the statistics shown in Capability Metrics view and indicates whether each statistic is displayed for long-term or short-term data. There is a "y" if it is displayed and an "n" if it does not display. For a discussion of the equations used to calculate each of these statistics, see the online Oracle Crystal Ball Statistical Guide.

Z scores are typically reported only for normal data. Crystal Ball always displays Z scores. It is up to the user to determine if the values are appropriate.

Table 11 Capability Statistics Calculated by Crystal Ball

Metric	Long- term	Short- term	Description
Mean	у	у	Mean of the forecast values
Standard deviation	у	у	Standard deviation of the forecast values
Ср	n	у	Short-term capability index indicating what quality level the forecast output is potentially capable of producing. It is defined as the ratio of the specification width to the forecast width. If a Cp is equal to or greater than 1, then a short-term 3-sigma quality level is possible.

Metric	Long- term	Short- term	Description
Рр	у	n	Long-term capability index indicating what quality level the forecast output is potentially capable of producing. It is defined as the ratio of the specification width to the forecast width. If a Cp is equal to or greater than 1, then a long-term 3-sigma quality level is possible.
Cpk-lower	n	у	One-sided short-term capability index; for normally distributed forecasts, the ratio of the difference between the forecast mean and lower specification limit over three times the forecast short-term standard deviation; often used to calculate process capability indices with only a lower specification limit.
Ppk-lower	у	n	One-sided long-term capability index; for normally distributed forecasts, the ratio of the difference between the forecast mean and lower specification limit over three times the forecast long-term standard deviation; often used to calculate process capability indices with only a lower specification limit.
Cpk-upper	n	у	One-sided short-term capability index; for normally distributed forecasts, the ratio of the difference between the forecast mean and upper specification limit over three times the forecast short-term standard deviation; often used to calculate process capability indices with only an upper specification limit.
Ppk-upper	у	n	One-sided long-term capability index; for normally distributed forecasts, the ratio of the difference between the forecast mean and upper specification limit over three times the forecast long-term standard deviation; often used to calculate process capability indices with only an upper specification limit.
Cpk	n	у	Short-term capability index (minimum of calculated Cpk-lower and Cpk-upper) that takes into account the centering of the forecast with respect to the midpoint of the specified limits; a Cpk equal to or greater than 1 indicates a quality level of 3 sigmas or better.
Ppk	у	n	Long-term capability index (minimum of calculated Ppk-lower and Ppk-upper) that takes into account the centering of the forecast with respect to the midpoint of the specified limits; a Ppk equal to or greater than 1 indicates a quality level of 3 sigmas or better.
Cpm	n	у	Short-term Taguchi capability index; similar to Cpk but considers a target value, which may not necessarily be centered between the upper and lower specification limits.
Ppm	у	n	Long-term Taguchi capability index; similar to Ppk but considers a target value, which may not necessarily be centered between the upper and lower specification limits.
Z-LSL	у	у	The number of standard deviations between the forecast mean and the lower specification limit.
Z-USL	у	у	The number of standard deviations between the forecast mean and the upper specification limit.
Zst	у	n	For short-term metrics when only one specification limit is defined, equal to Z-LSL if there is only a lower specification limit or Z-USL if there is only an upper specification limit.
Zst-total	n	у	For short-term metrics when both specification limits are defined, the number of standard deviations between the short-term forecast mean and the lower boundary of combining all defects onto the upper tail of the normal curve. Also equal to Zlt-total plus the Z-score shift value if a long-term index is available.
ZIt	n	у	For long-term metrics when only one specification limit is defined, equal to Z-LSL if there is only a lower specification limit or Z-USL if there is only an upper specification limit.

Metric	Long- term	Short- term	Description
Zlt-total	у	n	For long-term metrics when both specification limits are defined, the number of standard deviations between the long-term forecast mean and the lower boundary of combining all defects onto the upper tail of the normal curve. Also equal to Zst-total minus the Z-score shift value if a short-term index is available.
p(N/C)-below	у	у	Probability of a defect below the lower specification limit; DPUBELOW
p(N/C)-above	у	у	Probability of a defect above the upper specification limit; DPUABOVE
p(N/C)-total	у	у	Probability of a defect outside the lower and upper specification limits; DPUTOTAL
PPM-below	у	у	Defects below the lower specification limit, per million units
PPM-above	у	у	Defects above the upper specification limit, per million units
PPM-total	у	у	Defects outside both specification limits, per million units
LSL	у	у	Lower specification limit, the lowest acceptable value of a forecast involved in process capability, or quality, analysis.
USL	у	у	Upper specification limit, the highest acceptable value of a forecast involved in process capability analysis.
Target	у	у	The ideal target value of a forecast involved in process capability analysis.
Z-score shift	у	у	An optional shift value to use when calculating long-term capability metrics. The default, set in the Capability Options panel, is 1.5.



Maximizing the Use of Crystal Ball

In This Appendix

Introduction	381
Simulation Accuracy	381
Simulation Speed	383
Sample Size	384
Correlated Assumptions	384
Crystal Ball and Multi-processor Computers	385

Introduction

This appendix contains information that you can use to improve the overall performance of Crystal Ball. These improvements occur in terms of the accuracy of the model or speed of the results.

Simulation Accuracy

The accuracy of the simulation is primarily governed by two factors:

- The number of trials, or length, of the simulation
- The sampling method

Generally speaking, the more trials you run in a simulation, the greater the accuracy of the statistics and percentiles information. This greater accuracy comes at the expense of lengthier simulation times and higher memory usages (see later sections on simulation speed and memory usage). Also, for a given number of trials, the accuracy of the statistics and percentiles greatly depends on the shape and nature of the forecast distribution.

If you are not sure how many trials to run for a specific level of accuracy, you can use the precision control feature of Crystal Ball to automatically determine the appropriate number of trials to run. For a detailed picture of a simulation's statistical accuracy, you can run the Bootstrap tool to generate a forecast chart for each statistic or percentile of interest.

The sampling method is the other primary factor governing simulation accuracy. During a simulation, Monte Carlo sampling generates natural, "what-if" type scenarios while Latin Hypercube's sampling is constrained, but more accurate. See "Sampling Method" on page 383 for further discussion on sampling methods.

Precision Control

The precision control feature sets how precise you want the forecast statistics to be. Crystal Ball Crystal Ball runs the simulation until the selected forecast statistics reach the required precision as determined by calculating confidence intervals.

See the Oracle Crystal Ball Statistical Guide for more information about how Crystal Ball calculates confidence intervals.

Generally speaking, as more trials are calculated, the confidence intervals narrow and statistics become more precise. The precision control feature in Crystal Ball uses this characteristic of confidence intervals to determine when a specified precision of a statistic has been reached. It compares the specified precision to the confidence interval. When the calculated confidence interval drops to less than the specified precision, the simulation stops.

For each forecast, you can specify precision in either absolute terms in units of the forecast, or in relative terms as percentages. These settings are made on the Precision tab of the expanded Define Forecast dialog or the Forecast Preferences dialog. Each method, absolute or relative, has its own benefits and drawbacks.

Specifying precision in absolute terms offers greater control of the simulation when the shape and scale of the forecast distribution is roughly known. For example, for a Gross Profit forecast (from the Vision Research model) that ranges from \$25.5 to \$64.0 million dollars, you can require the precision of the mean to be within plus or minus \$100,000 or some other convenient measure of accuracy. However, with the same forecast range, an absolute accuracy of \$1000 might require an unreasonably large number of trials to reach. So, the drawback of using absolute precision is that it might require experimentation to determine reasonable accuracy values.

Specifying precision in relative terms offers greater control of the simulation when the shape and scale of the forecast distribution are largely unknown and you are interested in the accuracy only as it relates to the overall distribution itself. In the previous Gross Profit example, you might not know or care if the distribution ranges from \$25,500 to \$64,000 dollars or from \$25.5 to \$64.0 million dollars. You might require only that the simulation's estimate of the mean fall within plus or minus 5% of itself.

You might encounter the drawback of using relative precision when the forecast statistic is close to zero. For example, suppose a forecast's distribution straddles the break-even point of zero. A relative precision of 5% of the mean, or roughly \$0.5 million, results in a very small confidence interval (relative to the full range width of \$49.1 million) that might take an unexpectedly large number of trials to satisfy.

Finally, Crystal Ball combines the individual forecast precision options with the confidence level value found in the Trials tab of the Run Preferences dialog to calculate confidence intervals. Generally, it is a good idea to leave this value at 95% or 90% so that you can have a high degree of confidence that the precision requirements have been met. However, if you have a large number of forecasts defined with precision control set, you can adjust the confidence level up or down to globally change the accuracy of all forecasts together.

Sampling Method

Selecting Monte Carlo or Latin Hypercube sampling affects how random numbers are generated for individual assumptions.

In almost all cases, Latin Hypercube produces more accurate forecast statistics (especially the mean) given the same number of trials as Monte Carlo, because it is a more consistent sampling method. If you are primarily interested in the accuracy of the statistics, you should select Latin Hypercube as the sampling method in the Sampling tab of the Run Preferences dialog.

If you are primarily interested in evaluating how the spreadsheet model behaves under various "what-if" scenarios, you should select Monte Carlo as the sampling method. Monte Carlo produces assumptions with the most randomness and simulates real-life situations the best.

Simulation Speed

Monte Carlo simulations can be very time-consuming. You can change a number of factors that affect the speed of simulations. The factors are listed below in order of importance:

1. Change the Speed preferences.

In the Speed tab of the Run Preferences dialog are options that can substantially increase the speed of the simulation. Of course, if the Extreme Speed setting is available, you can experience dramatic cuts in run times.

If the model is incompatible with Extreme speed or the license does not include it, you can, in increasing order of helpfulness:

- Redraw forecasts less often (available in Extreme speed).
- Update worksheets only every few seconds or so.
- Suppress chart windows during simulations (available in Extreme speed).
- Minimize Microsoft Excel workbooks during simulations.
- Minimize Microsoft Excel workbooks and suppress chart windows (the fastest combination of settings).
- Use the precision control feature.

The precision control feature can be used to either increase the precision of the simulations or increase the speed of the simulations. If you set the confidence level to a high number, the simulations will be more precise, but might run significantly longer. However, if you do not need as precise a result, you can set the confidence level to a lower number and the simulation speed will increase.

Using this feature to speed up the model will require you to experiment with different confidence levels.

3. Reduce the size of the model by reducing the number of assumptions, forecasts, and correlations.

Large models require more time per trial. For example, a model that takes 3 or 4 seconds per recalculation cycle will take up to an hour to simulate 1,000 trials.

Greater numbers of assumptions and forecasts slow the simulation, especially if the assumptions and forecasts are scattered across many spreadsheets in the model. Start by examining the structure and nature of the model to locate possible efficiencies. You can also use the sensitivity feature or the Tornado Chart tool to determine which assumptions contribute the least amount of uncertainty to the most important forecasts. Freeze or eliminate the least important assumptions from the simulation.

Correlated assumptions can also consume a significant amount of processing time; the time grows geometrically as the number of correlated assumptions increases.

4. Reduce the use of other applications.

Quitting other applications and closing or minimizing windows can be helpful in reducing overhead and increasing simulation speed.

5. Increase the system's RAM.

The amount of RAM in the computer has a large effect on the speed of simulations. Modern operating systems give applications such as spreadsheets the appearance of additional RAM through the use of virtual memory.

Virtual memory enables you to run a greater number of applications than would otherwise be possible, but slows down overall processing speed because the system is frequently accessing the hard drive. If you hear the hard disk being used during a simulation, there might not be enough RAM to hold all parts of the simulation. Buying more RAM or turning off virtual memory (if possible) are solutions to this problem.

Sample Size

The sample size setting is located on the Sampling tab of the Run Preferences dialog. Sample size, which is initially set to 500, affects Latin hypercube sampling. It divides each assumption's distribution into a number of intervals of equal probability. The sample size governs the number of intervals for each distribution. Crystal Ball Crystal Ball generates an assumption value for each interval according to the interval's probability distribution.

While any sample size greater than 100 should produce sufficiently acceptable results, you can set this number higher to maximize accuracy. There is no absolute limit to sample size, although samples greater than 100,000 work best with at least 1 GB RAM and may take a long time to run. The increased accuracy resulting from the use of larger samples, however, requires additional memory. If memory becomes an issue, reduce the sample size and consider adding more RAM.

Correlated Assumptions

There is a practical limit of about 250 fully correlated assumptions (assumptions that are correlated to every other assumption) and up to about 1,000 serially correlated assumptions (assumptions that are correlated to one or two other assumptions) for each workbook. If you have a large number of correlated assumptions:

For coefficients that are close to zero and are nearly independent, remove the correlation.

For coefficients that are close to one, replace one of the assumptions with a formula in the spreadsheet.

If Crystal Ball detects inconsistently correlated assumptions when a simulation is running, it first determines whether small adjustments to the correlation coefficients are possible. This process might take a long time, depending on the number of correlated assumptions. Crystal Ball displays the message "Examining the Correlation Coefficients". If you get this message, you should probably stop and redefine the correlations.

If small adjustments to the correlation coefficients are possible, a dialog opens and prompts you to cancel the simulation or continue with adjusted coefficients.

Select one of the following responses:

- Click Adjust Coefficients This Time Only to continue the simulation with the adjusted coefficients.
 - Since adjusted coefficients are not saved permanently, a dialog appears again if you stop and restart the simulation.
- To continue the simulation, replacing the original correlation coefficients with the adjusted ones, click Adjust Coefficients Permanently.
 - Correlation coefficients that were specified using a cell reference in place of an actual value are replaced with a permanent value (the cell references are removed).
- To terminate the simulation, click Cancel.

If small adjustments to the correlation coefficients are not possible, the simulation stops with an error message prompting you to reexamine the coefficients. To solve the problem, you can usually make large coefficients smaller or change the spreadsheet model to use formulas to calculate assumptions with large correlation coefficients. Creating a report containing just the assumptions might make it easier to spot problems.

Crystal Ball and Multi-processor Computers

Subtopics

- Crystal Ball and Multiple Processors
- Crystal Ball and Multi-threading

In general, Crystal Ball users should see a linear increase in simulation speed with an increase in processor speed. Using a computer with multiple processors and multiple threads can further enhance performance as described in the following sections.

Crystal Ball and Multiple Processors

In computers with two or more processors, users should see a slight increase in simulation speed because chart updating occurs on a separate thread. Typically, Extreme Speed is faster than Normal Speed, even with multiple processors. In Normal Speed, when a seed value is set and

multiple processors are used with a model that has embedded Crystal Ball distribution functions (CB.Normal, and so on) in the Microsoft Excel spreadsheet, results may not be exactly the same between simulations. However the results are valid. This is because Microsoft Excel does not guarantee execution order in this case. This can be corrected by running in Extreme Speed or Normal Speed with a single thread.

Crystal Ball and Multi-threading

If you are using a computer with multiple processors—for example, dual cores—you can use a new "multi-threaded" option in Microsoft Excel 2007 or later to split up spreadsheet model recalculations into separate tasks. These tasks can then be run independently on each processor to speed up the overall recalculation time. Since time for one calculation is reduced, the time to run an entire simulation is also reduced. For two processors, the increase in simulation speed can be roughly anywhere from 10% to 50%, depending on the model. Crystal Ball uses Microsoft Excel's multi-threading setting by default when performing simulation calculations in Microsoft Excel 2007 on multi-core or multi-processor computers.

- To activate multi-threading in Microsoft Excel 2007 or later:
- Click the Office button.
- Select Microsoft Excel Options, then Advanced.
- In the Advanced Options dialog, scroll to the Formulas group and then check "Enable multi-threaded calculation."
- Click **OK** to accept the setting and close the dialog.

To use multi-threading efficiently, you should be working with a spreadsheet model that:

- Is large (that is, it takes more than 0.5 sec for each recalculation).
- Can easily be divided into separate tasks by Microsoft Excel 2007 or later (for example, it might have separate chains or groups of formulas that do not depend on each other).
- Is running at Normal Speed, since Extreme Speed does not increase benefits with this feature.

Depending on model size, it is possible that performance can be improved on multi-core or multi-processor computers by manually disabling the use of multi-threading. In general, smaller models run more slowly with multi-threading and larger models run faster. Changing this setting on single-core or single-processor computers has no impact. To disable multi-threading in Microsoft Excel 2007 before running Crystal Ball models, consult the Microsoft Excel documentation.

Accessibility

In This Appendix

Documentation Accessibility	387
Access to Oracle Support for Hearing-Impaired Customers	387
Enabling Accessibility for Crystal Ball	388
Activating and Using Accessibility Mode	388
Using the Tab and Arrow Keys in the Crystal Ball User Interface	390
Accessing Charts without a Mouse	390
Tabbing Through the Crystal Ball Control Panel	391
Using Microsoft Windows Large Fonts in Crystal Ball	391
Controlling Screen Flicker when Resizing Charts	392
Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2003 or Earlier	392
Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2007 or Later	402

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other marketleading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at http://www.oracle.com/accessibility/.

Access to Oracle Support for Hearing-Impaired Customers

Oracle customers have access to electronic support through My Oracle Support or by calling Oracle Support at 1.800.223.1711. Hearing-impaired customers in the U.S. who wish to speak to an Oracle Support representative may use a telecommunications relay service (TRS). Information about the TRS is available at http://www.fcc.gov/cgb/consumerfacts/trs.html, and a list of telephone numbers is available at http://www.fcc.gov/cgb/dro/trsphonebk.html. International hearing-impaired customers should use the TRS at +1.605.224.1837. An Oracle

Support engineer will respond to technical issues according to the standard service request process.

Enabling Accessibility for Crystal Ball

You do not need to enable a minimum level of accessibility specifically for Crystal Ball; dialog controls are always accessible through tabbing and shortcut keys. Also, results data can be extracted to Microsoft Excel spreadsheets and pasted into Microsoft PowerPoint slides, which are accessible through Microsoft Office. For information about Microsoft Excel or Microsoft PowerPoint accessibility, refer to Microsoft Office product documentation.

For maximum accessibility, particularly involving features for users with visual impairments, Crystal Ball can be switched into Accessibility mode. For details, see "Activating and Using Accessibility Mode" on page 388.

Other topics concerning accessibility:

- "Documentation Accessibility" on page 387
- "Access to Oracle Support for Hearing-Impaired Customers" on page 387
- "Activating and Using Accessibility Mode" on page 388
- "Using the Tab and Arrow Keys in the Crystal Ball User Interface" on page 390
- "Accessing Charts without a Mouse" on page 390
- "Tabbing Through the Crystal Ball Control Panel" on page 391
- "Using Microsoft Windows Large Fonts in Crystal Ball" on page 391
- "Controlling Screen Flicker when Resizing Charts" on page 392
- "Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2003 or Earlier" on page
- "Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2007 or Later" on page 402

Activating and Using Accessibility Mode

Subtopics

- Activating Accessibility Mode
- A Note About Accessible Default Cell Preferences

Accessibility is the extent to which products, systems, and so on can be used by people with a wide range of disabilities. Specifically, accessibility is the ability of products to support the needs of users with visual impairments—including blindness, partial sight (low vision), and color blindness—and motor impairments that prevent them from using a mouse or a standard keyboard.

Crystal Ball supports accessibility at all times through keyboard shortcuts and other means mentioned in "Enabling Accessibility for Crystal Ball" on page 388. Crystal Ball users who require additional support can activate Accessibility mode with the following features:

- When defining new forecast charts, the certainty grabber and certainty level marker lines are displayed by default.
- Series or other distinctions in forecast, overlay, trend, and sensitivity charts have different patterns, not just different colors. Filtered data in scatter charts has a pattern.
- Accessible text is provided in blank charts to indicate whether data is available or why it is not.
- The Compare Run Modes chart now plots one series with a pattern.
- Microsoft Excel charts are included in reports by default and values of the chart series are written out to the report spreadsheet.
- The cell preference Add comment to cell is active and comments are added to new assumption, decision variable, and forecast cells defined in new workbooks. For additional information, see "A Note About Accessible Default Cell Preferences" on page 389.
- The default size of assumption charts in reports is changed from 50% to 100%.
- To ensure consistent progress readings, the simulation progress bar is hidden in the Crystal Ball control panel when an OptQuest optimization is running.
- Constants to support the new Accessibility mode preference are added to the Crystal Ball Developer Kit for calls that set or return run preference settings.

Activating Accessibility Mode

- To activate Accessibility mode:
- 1 Select Enable accessibility options in the Crystal Ball Welcome screen or the Options tab of the Run Preferences dialog.
- Select Use Crystal Ball in the Welcome screen or click OK in the Run Preferences dialog.
- To deactivate Accessibility mode, clear the Enable accessibility options selection in either the Welcome screen or the Run Preferences dialog.

A Note About Accessible Default Cell Preferences

The default Cell Preferences settings in Crystal Ball do not meet accessibility standards for differentiation by means other than color. Both pattern and color are used in Accessibility mode. When Accessibility mode is activated, new workbooks have accessible defaults as described in the bulleted list presented previously in the section.

- To apply accessible defaults to existing workbooks:
- Activate Accessibility mode.

- 2 Open one or more existing workbooks.
- 3 Select File, and then New to create a new workbook.
- With the new workbook selected, select **Define**, and then **Cell Preferences**. (In Microsoft Excel 2007 or later, select **Define**, and then **Cell Prefs**.)
 - The Cell Preferences dialog opens with accessibility default settings selected.
- 5 Click Apply To.
- 6 In the Apply To dialog, select All tabs, and then All open and new workbooks.
- 7 Click OK.
 - The accessibility defaults are applied to all open workbooks.
- 8 Save the workbooks to keep the new defaults.

Using the Tab and Arrow Keys in the Crystal Ball User Interface

The main Crystal Ball menubar and menu commands are accessed with shortcut keys. After a menu is open, the Tab key or Down Arrow key highlights commands in a circular sequence (tabbing from the last item moves focus to the first item). Using Shift+Tab or the Up Arrow highlights commands in the opposite direction.

Default tab order in the Crystal Ball dialogs flows from left to right, top to bottom. Tab can be used to access the first item of a subwindow in a dialog, but then arrow keys are needed to move to additional items. For example, in the Distribution Gallery, Alt+b displays six basic distributions and places the focus on one of them. Using one of the arrow keys moves focus to an adjacent distribution. Tab moves focus to the OK button at the bottom of the dialog.

For additional information about navigation in Crystal Ball, see "Accessing Charts without a Mouse" on page 390.

Accessing Charts without a Mouse

For general information about navigating through Crystal Ball without a mouse, see "Using the Tab and Arrow Keys in the Crystal Ball User Interface" on page 390. The information given there about the Tab key and tab order applies to controls in Crystal Ball charts. However, additional rules apply for accessing data tables, such as statistics and percentiles, for navigating through multiple views in the same window—for example, in Split View, and for navigating through multiple charts.

Tabbing Through Data Tables

If you use the Tab key to select a data view in a chart, a cell is selected. Each time you press Tab, the next cell in that row is selected. At the end of the row, the first cell in the next row is selected. At the end of the table, the next Tab advances to an additional view or the next control.

Note: When you use Tab to access a data table, the first cell in the table is not always selected. If you leave the table after only a few Tab key uses, press Shift+Tab several times to return to the table and test whether additional data is available.

Tabbing Through Multiple Views

If a chart displays multiple views, the default tab order changes. Initially, the first label or control at the top is selected and then the second until all labels and controls above the chart have been selected. The next time you use the Tab key, the first chart view is selected (if present), then the next chart view, and so on until all chart views have been selected. Then, the first data table is selected (if present). After navigation through all data tables, the first label or control below the chart views is selected (if present) and navigation proceeds as usual.

Except for scatter charts, you can use Ctrl+(view number) to select a chart view without tabbing. For example, suppose a chart in Split View contains a column with two chart views and a second column with two data tables. Considering tab order, the first chart view is 1, the second chart view is 2, the first data table is 3, and the second data table is 4. Then, Ctrl+2 selects the second chart view and Ctrl+3 selects the first data table. Once a data table is selected, Tab moves through rows of data as usual.

Tabbing Through Multiple Charts

If several charts are open but none are in Split View, you can use Ctrl+(chart number) to cycle through the charts. Ctrl+1 selects the first chart, Ctrl+2 selects the second, and so on.

Note: Once a chart is selected, you can use the View menu commands (Alt+command shortcut key) or the shortcut keys listed in Table 5, "Shortcut keys for Chart Preferences," on page 116 to cycle through available chart views.

Tabbing Through the Crystal Ball Control Panel

You can use the Tab key to access labels and controls in the Crystal Ball Control Panel. When you select the More button and press Enter, an expanded view with simulation statistics becomes accessible. They show how fast the simulation ran and how many assumptions, decision variables, and forecasts were included in it.

In expanded view, press Ctrl+s to access the table of simulation statistics. Then, you can use Tab and Shift+Tab to move through the labels and values presented there.

Using Microsoft Windows Large Fonts in Crystal Ball

If the Large Fonts setting is used in Microsoft Windows, screen resolution must be set higher than 1024x768 pixels.

Controlling Screen Flicker when Resizing Charts

Excessive screen flickering can result when Crystal Ball users resize assumption and forecast charts. To reduce flickering below the limits required by accessibility standards, open Windows Control Panel and make the following setting: Select **Display**, then **Appearance**, then **Effects**, and then clear the checkbox for **Show window contents while dragging**.

Note: You need to shut down and restart Crystal Ball and Microsoft Excel to ensure that this setting is activated.

Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2003 or Earlier

Subtopics

- Crystal Ball Toolbar
- Crystal Ball Menus
- Shortcut key Combinations in Microsoft Excel 2003 or Earlier
- Distribution Gallery Shortcut Keys
- Chart Preference Shortcut Keys


This section describes the Crystal Ball toolbar and associated menu commands for Crystal Ball used with Microsoft Excel 2003 or earlier supported versions.

Note: The information in this appendix is for those using Crystal Ball with Microsoft Excel 2003 or earlier. If you are using Crystal Ball with Microsoft Excel 2007 or later, see "Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2007 or Later" on page 402 instead.

Crystal Ball Toolbar

As you point to each of the Crystal Ball toolbar buttons, a tool tip identifies it, as shown in Figure 193.

Figure 193 Crystal Ball Toolbar



Crystal Ball Menus

When Crystal Ball is open in Microsoft Excel, the following special menus are added to the standard Microsoft Excel menus:

- Define menu lets you define and select assumptions, decision variables, and forecasts; freezes assumptions, decision variables, and forecasts so they are not included in simulations; also lets you copy, paste and clear assumptions, decision variables, and forecasts as well as set cell preferences.
- Run menu starts, stops, continues, and resets simulations; runs simulations a step at a time; if available, starts OptQuest or Predictor; opens the Crystal Ball tools: Batch Fit, Correlation Matrix, Tornado Chart, Bootstrap, Decision Table, Scenario Analysis, 2D Simulation; saves and restores simulation results; and sets run preferences.
- Analyze menu creates and opens charts and reports; cascades and closes windows; and extracts data.
- Help, then Crystal Ball menu Opens online help for Crystal Ball; displays the following items: the About Box for Crystal Ball, Oracle Crystal Ball User's Guides and other documentation in Adobe Acrobat format, Crystal Ball examples, the Crystal Ball Welcome Screen, lets you access Crystal Ball licensing features, and more.

For more information on each command, see the tables later in this appendix or view Help from within Crystal Ball using the Help menu or icon.

Distribution Gallery Menus

The Distribution Gallery has the following menus:

- Edit menu lets you copy, paste, modify, and delete, and print distributions.
- View menu determines how distributions are displayed in the Distribution Gallery and shows or hides the Description pane.
- Categories menu lets you create, modify, delete, reorder, and share distribution categories.
- Help menu displays help for the Distribution Gallery or the active distribution.

Shortcut key Combinations in Microsoft Excel 2003 or Earlier

Use the following Alt+key combinations in Microsoft Excel 2003 or earlier supported versions of Microsoft Excel to execute Crystal Ball's Define, Run, Analyze, and Help menu commands and chart windows commands without using the mouse. For shortcut keys in Microsoft Excel 2007 or later, see "Shortcut Key Combinations in Microsoft Excel 2007 or later" on page 404.

Note: Oracle tries to avoid keyboard shortcut conflicts; if a key combination is the same as one on a Microsoft Excel menu, you can use the down arrow key to access the command you want. Commands are listed in the same order in which they are displayed in menus.

Table 12 Crystal Ball Define Menu Commands, Microsoft Excel 2003 or Earlier

Button	Command	Command Action	Keystrokes
	Define Assumption	Opens the Define Assumption dialog to define a Crystal Ball assumption cell	Alt+n, a
Ŷ	Define Decision	Opens the Define Decision Variable dialog to define a Crystal Ball decision variable cell	Alt+n, d
ofh	Define Forecast	Opens the Define Forecast dialog to define a Crystal Ball forecast cell	Alt+n, f
	Select All Assumptions	Selects all assumption cells on the active worksheet	Alt+n, m
②	Select All Decisions	Selects all decision variable cells on the active worksheet	Alt+n, i
Ab	Select All Forecasts	Selects all forecast cells on the active worksheet	Alt+n, r
n/a	Select	Selects specified Crystal Ball data cells	Alt+n, s
n/a	Freeze	Excludes specified Crystal Ball data cells from a simulation	Alt+n, z
90	Copy Data	Copies Crystal Ball cell definitions to specified cells	Alt+n, c
(C)	Paste Data	Pastes copied Crystal Ball cell definitions into specified cells	Alt+n, p
	Clear Data	Removes Crystal Ball cell definitions from specified cells	Alt+n, e
n/a	Cell Preferences	Opens the Cell Preferences dialog	Alt+n, I

Table 13 Crystal Ball Run Menu Commands, Microsoft Excel 2003 or Earlier

Button	Command	Command Action	Keystrokes
	Start Simulation	Begins running a Crystal Ball simulation	Alt+r, s
	Continue Simulation	Resumes a stopped simulation	Alt+r, c
44	Reset Simulation	Clears generated simulation data (Crystal Ball data cells remain defined)	Alt+r, r
	Stop Simulation	Stops a simulation	Alt+r, s
	Single Step	Steps through a simulation one trial at a time	Alt+r, i

Button	Command	Command Action	Keystrokes
4	OptQuest	Starts OptQuest for results optimization (if available)	Alt+r, o
p And	Predictor	Starts Predictor for time series analysis	Alt+r, p
n/a	More Tools	Offers the Crystal Ball general and integration tools for additional analysis and charting	Alt+r, t
n/a	Batch Fit (general)	Automatically fits selected continuous probability distributions to multiple data series	Alt+r, t, f
n/a	Bootstrap (general)	Addresses the reliability and accuracy of forecast statistics	Alt+r, t, b
n/a	Correlation Matrix (general)	Rapidly defines and automates correlations of assumptions	Alt+r, t, c
n/a	Data Analysis (general)	Loads data from spreadsheets directly into Crystal Ball forecast charts	Alt+r, t, y
n/a	Decision Table (general)	Evaluates the effects of alternate decisions in a simulation model	Alt+r, t, d
n/a	Scenario Analysis (general)	Displays what inputs created particular outputs	Alt+r, t, s
n/a	Tornado Chart (general)	Individually analyzes the impact of each model variable on a target outcome	Alt+r, t, t
n/a	2D Simulation (general)	Independently addresses uncertainty and variability using a two-dimensional simulation	Alt+r, t, 2
n/a	Enterprise Performance Management (integration)	If available, enables Crystal Ballto run simulations on Enterprise Performance Management data using Smart View	Alt+r, t, m
n/a	Strategic Finance (integration)	If available, opens the Strategic Finance Setup wizard for setting up a workbook that allows Crystal Ball to run with Strategic Finance data	Alt+r, t, e
n/a	Compare Run Modes	Compares simulation run time in Normal and Extreme speed	Alt+r, t, r
n/a	Save Results	Saves current simulation results for later display and analysis	Alt+r, v
n/a	Restore Results	Restores saved results files for display and analysis	Alt+r, e
	Run Preferences	Opens the Run Preferences dialog for setting trials, sampling, and other run preferences	Alt+r, u

 Table 14
 Crystal Ball Analyze Menu Commands, Microsoft Excel 2003 or Earlier

Button	Command	Command Action	Keystrokes
_	Assumption Charts	Opens the Assumption Charts dialog so you can display or manage assumption charts	Alt+a, a
dh	Forecast Charts	Opens the Forecast Charts dialog so you can display or manage forecast charts	Alt+a, f
	Overlay Charts	Opens the Overlay Charts dialog so you can display existing or create new overlay charts	Alt+a, o

Button	Command	Command Action	Keystrokes
	Trend Charts	Opens the Trend Charts dialog so you can display existing or create new trend charts	Alt+a, t
F	Sensitivity Charts	Opens the Sensitivity Charts dialog so you can display existing or create new sensitivity charts	Alt+a, s
****	Scatter Charts	Opens the Scatter Charts dialog so you can display existing or create new scatter charts	Alt+a, s
n/a	OptQuest Charts	If OptQuest is available, displays OptQuest charts following an optimization	Alt+a, q
n/a	Predictor Charts	If Predictor is available, displays Predictor charts following a time series forecast	Alt+a, p
n/a	Open Selected Cells	Opens associated charts for selected cells	Alt+a, n
n/a	Cascade	Arranges open charts for easier viewing	Alt+a, c
n/a	Close All	Closes all open charts	Alt+a, I
	Create Report	Opens the Create Report dialog so you can create a report with data and charts for the simulation and results	Alt+a, r
	Extract Data	Copies selected simulation data to a worksheet for additional analysis	Alt+a, d

Table 15 Crystal Ball Help Menu Commands, Microsoft Excel 2003 or Earlier

Button	Command	Command Action	Keystrokes
n/a	Technical Support	Displays the main Crystal Ball Web page with a link to Crystal Ball technical support	Alt+h, b, t
?	Crystal Ball Help	Displays online help for Crystal Ball	Alt+h, b, h
n/a	Crystal Ball Documentation	Displays a list of available online documentation for Crystal Ball	Alt+h, b, c
n/a	Oracle EPM Documentation	Displays the OTN index for Oracle Enterprise Performance Management documentation, including Crystal Ball	Alt+h, b, d
n/a	Examples Guide	Displays a list of available example models so you can load the choice into Crystal Ball	Alt+h, b, e
n/a	Welcome Screen	Displays a Welcome screen that lets you automatically activate the process capability features for quality programs such as Six Sigma or set percentile preferences frequently used in the oil and gas industry	Alt+h, b, w
n/a	Licensing	Displays the Activate a License dialog so you can enter a serial number and activate a Crystal Ball license	Alt+h, b, l
n/a	About Crystal Ball	Offers version and other information about the current Crystal Ball release including the username	Alt+h, b, b

Table 16 Crystal Ball General Chart Menu Commands, All Supported Versions of Microsoft Excel

Command	Command Action	Keystrokes
Bring Crystal Ball to front	A Microsoft Windows shortcut for switching between applications	Alt+Tab, then Oracle Crystal Ball
Edit, then Copy Chart	Copies the active Crystal Ball chart to the clipboard	Alt+e, c
Edit, then Page Setup	Defines margins, paper size, orientation, and other layout information for printing	Alt+e, u
Edit, then Print Preview	Displays the active chart as it will print with the current Page Setup settings	Alt+e, v
Edit, then Print	Prints the current chart to the specified printer or file	Alt+e, p
View, then Frequency	If available, sets the current chart to show the number or frequency of values occurring in a given interval (bin)	Alt+v, f
View, then Cumulative [Frequency]	If available, sets the current chart to show the number or proportion (percentage) of values less than or equal to a given amount	Alt+v, c
View, then Reverse Cumulative [Frequency]	If available, sets the current chart to show the number or proportion (percentage) of values greater than or equal to a given amount	Alt+v, r
View, then Statistics	If available, sets the current chart to show a set of descriptive statistics for charted values	Alt+v, s
View, then Percentiles	If available, sets the current chart to show percentile information, in 10% increments	Alt+v, p
View, then Goodness Of Fit	If available, sets the current chart to show goodness-of-fit statistics for all fitted distributions	Alt+v, g
View, then Capability Metrics	If available, sets the current forecast chart to show a set of process capability metrics based on the specification limit(s) set for that forecast	Alt+v, a
View, then Split View	If available, sets the current chart to show data, such as statistics, beside the associated chart	Alt+v, I
Preferences, then Chart Preferences	Open the Chart Preferences dialog for changing chart appearance settings	Alt+p, c

Table 17 Assumption-specific Chart Menu Commands, All Supported Versions of Microsoft Excel

Command	Command Action	Keystrokes
Edit, then Add To Gallery	Adds the active assumption to the Distribution Gallery in the specified category	Alt+e, a
Parameters, then Show Cell References	Shows cell references instead of values when cell references are entered into parameter fields	Alt+m, c
Preferences, then Assumption	Opens the Assumption Preferences dialog for choosing a view and controlling chart display	Alt+p, a
Help, then Assumption Help	Displays online help about assumptions	Alt+h, a

 Table 18
 Forecast-specific Chart Menu Commands, All Supported Versions of Microsoft Excel

Command	Command Action	Keystrokes
Forecast, then Open Sensitivity Chart	Creates or displays a sensitivity chart for the active forecast	Alt+f, s
Forecast, then Fit Probability Distribution	Uses the default or currently selected distributions and ranking method to fit a distribution to the active forecast chart	Alt+f, f
Preferences, then Forecast	Opens the Forecast Preferences dialog for controlling chart display, distribution fitting specifications, and more	Alt+p, f
Help, then Forecast Help	Displays online help about the forecast chart	Alt+h, f

 Table 19
 Overlay-specific Chart Menu Commands, All Supported Versions of Microsoft Excel

Command	Command Action	Keystrokes
Overlay, then Choose Forecasts	Displays the Choose Forecasts dialog to add or remove forecasts from the active overlay chart	Alt+o, c
Overlay, then Remove All	Removes all forecasts from the active overlay chart	Alt+o, r
Overlay, then Fit Probability Distribution	Uses the default or currently selected distributions and ranking method to fit a distribution to all forecasts on the active overlay chart	Alt+o, f
Preferences, then Overlay	Opens the Overlay Preferences dialog for controlling chart display and distribution fitting specifications	Alt+p, o
Help, then Overlay Help	Displays online help about the overlay chart	Alt+h, o

Table 20 Sensitivity-specific Chart Menu Commands, All Supported Versions of Microsoft Excel

Command	Command Action	Keystrokes
View, then Contribution To Variance Chart	Displays the sensitivity chart in terms of percentage of variance contributed by each assumption	Alt+v, c
View, then Rank Correlation Chart	Displays the sensitivity chart in terms of rank correlation of the assumptions in their effect on the forecast	Alt+v, r
View, then Sensitivity Data	Displays contribution to variance and rank correlation sensitivity measures for each assumption	Alt+v, s
Sensitivity, then Choose Target Forecast	Opens the Choose Forecast dialog so you can choose a forecast for sensitivity analysis	Alt+s, t
Sensitivity, then All Assumptions	Uses all assumptions in creating the sensitivity chart	Alt+s, a
Sensitivity, then Choose Assumptions	Displays the Choose Assumptions dialog so you can choose one or more assumptions for the sensitivity chart	Alt+s, c
Sensitivity, then Open Scatter Chart	Opens a scatter chart using the current target forecast and assumptions	Alt+s, s
Preferences, then Sensitivity	Opens the Sensitivity Preferences dialog for choosing a view, controlling chart display, and more	Alt+p, s
Help, then Sensitivity Help	Displays online help about the sensitivity chart	Alt+h, s

Table 21 Trend-specific Chart Menu Commands, All Supported Versions of Microsoft Excel

Command	Command Action	Keystrokes
View, then Centered On Median	Displays certainty bands centered on the median of each forecast value	Alt+v, m
View, then Cumulative	Displays certainty bands anchored at the lower end of the forecast range; shows the certainty that the forecast values will be at or below a given value	Alt+v, c
View, then Reverse Cumulative	Displays certainty bands anchored at the higher end of the forecast range; shows the certainty that the forecast values will be at or above a given value	Alt+v, r
Trend, then Choose Forecasts	Opens the Choose Forecasts dialog so you can choose forecasts to use in the trend chart	Alt+t, f
Trend, then Remove All	Removes all forecasts from the active trend chart	Alt+t, r
Preferences, then Trend	Opens the Trend Preferences dialog for choosing a view and controlling chart display	Alt+p, t
Help, then Trend Help	Displays online help about the trend chart	Alt+h, t

Table 22 Scatter-specific Chart Menu Commands, All Supported Versions of Microsoft Excel

Command	Command Action	Keystrokes
View, then Scatter View (1xN)	Plots selected assumptions and/or forecasts against a target assumption or forecast	Alt+v, s
View, then Matrix View (NxN)	Plots selected assumptions and/or forecasts against each other	Alt+v, m
Scatter, then Choose Data	Opens the Choose Data dialog so you can choose assumptions and forecasts to use in the scatter chart	Alt+s, d
Scatter, then Remove All	Removes all assumptions and forecasts from the active scatter chart	Alt+s, r
Preferences, then Scatter	Opens the Scatter Preferences dialog for choosing a view and controlling chart display	Alt+p, s
Help, then Scatter Help	Displays online help about the scatter chart	Alt+h, t

Table 23 Shortcut Keys for Controls without Labels, All Supported Versions of Microsoft Excel

Shortcut Key	Control	Description
Ctrl+Plus key	Equivalent to pressing the More button in the Define Assumption or Define Forecast dialog	Displays all settings in the dialog, basic and advanced (or preferences)
Ctrl+Minus key (hyphen or dash)	Equivalent to pressing the Less button in the Define Assumption or Define Forecast dialog	Hides advanced settings (or preferences) while displaying the basic settings
Alt+[Minimum truncation grabber (text box), Define Assumption dialog	Truncates the distribution at the low end of the value range and sets the value in the Minimum text box
Alt+]	Maximum truncation grabber (text box), Define Assumption dialog	Truncates the distribution at the high end of the value range and sets the value in the Maximum text box
Ctrl+Up arrow, Ctrl+Down arrow	Series list box in charts; other similar list boxes	Moves up and down through the list box so a series can be selected from the list

Shortcut Key	Control	Description
Ctrl+w, Ctrl+Shift+w	Warning footnotes in the Predictor Data Screening and Results windows	Enables screen readers to read warning footnotes if present
Ctrl+z	Zoom control in Predictor charts	Activates and inactivates the zoom control for enlarging and reducing Predictor chart display

Distribution Gallery Shortcut Keys

Note: In Microsoft Excel 2007 or later, if you click the upper half of the Define Assumption icon, or if the assumption has already been defined, the Distribution Gallery opens. If you click the lower half of the icon, a list of All, Basic, or Favorite distributions opens, depending on the active category in the Distribution Gallery.

You can use the keyboard to select distributions in the Distribution Gallery.

- To do this:
- 1 Choose **Define**, then **Define Assumption** or click the **Define Assumption** button to open the **Distribution**Gallery.
- 2 Type the first letter of the name of the target distribution, for example **n** for **normal**.
 - If more than one distribution starts with that letter, keep pressing to cycle through the distributions.
- 3 When the target distribution is highlighted, type Alt+o (for **OK**) to open its Define Assumption dialog. Or, you can press Alt+f to fit a distribution to it or Alt+c to cancel the selection and close the dialog.
- To select the description pane of the Distribution Gallery, press Ctrl+d.
- To apply a shortcut key to a custom category:
- 1 Open the Crystal Ball Distribution Gallery.
- 2 Edit the category name:

For a new category, select Categories, then New. In the Name field, enter an ampersand (&) before a letter in the category name. That letter becomes the shortcut key. Click OK to save.

For an existing category, select Categories, then Properties, add the ampersand to the existing name, and click OK to save.

Each category should use a unique shortcut key. The shortcut keys for the default categories are: Alt+b for the Basic category, Alt+a for All, and Alt+i for Favorites.

Chart Preference Shortcut Keys

Double-clicking in a chart area opens the Chart Preferences dialog.

Table 24 lists key combinations that can be used to cycle through settings available in the Chart Preferences dialog. Most of these commands work on the primary distribution — the theoretical probability distribution for assumptions, and the generated values for forecasts and overlay charts.

Table 24 Shortcut Keys for Chart Preferences, All Supported Versions of Microsoft Excel

Shortcut key	Command Equivalent	Description
Ctrl+d	View menu; Preferences, then <i>chartname</i> , then <i>chartname</i> Window, then View	Cycles through chart views — Frequency, Cumulative Frequency, Reverse Cumulative Frequency (for assumption and forecast charts)
Ctrl+b; Ctrl+g	Preferences, then Chart, then General, then Density	Cycles through bins or group interval values to adjust the number of data bins used to create the chart
Ctrl+l	Preferences, then Chart, then General, then Gridlines	Cycles through gridline settings: None, Horizontal, Vertical, Both
Ctrl+t	Preferences, then Chart, then Chart Type, then Type	Cycles through chart types: Area, Line, Column; for sensitivity charts: Bar (direction), Bar (magnitude), Pie (in Contribution To Variance view)
Ctrl+w	Preferences, then Chart, then General, then 3D Chart	Cycles between two-dimensional and three-dimensional chart display
Ctrl+m	Preferences, then Chart, then Chart Type, then Marker Lines	Cycles through central tendency marker lines: None, Mean, Median, Mode (except for sensitivity and trend charts)
Ctrl+n	Preferences, then Chart, then General, then Legend	Toggles the legend display on and off
Ctrl+p	Preferences, then Chart, then Chart Type, then Marker Lines, then Percentiles	Cycles through percentile markers: None, 10%, 20%,90%
Ctrl+Spacebar	View menu; Preferences, then chartname, then chartname Window, then View	Cycles through window views when Microsoft Excel is not in Edit mode: Chart, Statistics, Percentiles, Goodness Of Fit (if distribution fitting is selected — except for trend charts), and Capability Metrics (if process capability features are activated)

Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2007 or Later

Subtopics

- Introduction
- Crystal Ball Ribbon in Microsoft Excel 2007 or Later
- Shortcut Key Combinations in Microsoft Excel 2007 or later
- Distribution Gallery Shortcut Keys
- Chart Preference Shortcut Keys
- Compatibility and File Conversion Issues

Introduction

This section describes the Crystal Ball ribbon for those who are using Crystal Ball with Microsoft Excel 2007 or later. If you are using an earlier version of Microsoft Excel, see "Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2003 or Earlier" on page 392 instead.


See "Compatibility and File Conversion Issues" on page 409 for important information about using Crystal Ball models created in earlier versions of Microsoft Excel with Microsoft Excel 2007 or later or the Microsoft Excel 2007 or later Compatibility Pack.

Crystal Ball Ribbon in Microsoft Excel 2007 or Later

"Crystal Ball Commands and Shortcut Keys in Microsoft Excel 2003 or Earlier" on page 392 describes the Crystal Ball toolbar and menus in Microsoft Excel 2003 and earlier versions of Microsoft Excel. While Crystal Ball's functionality is the same in Microsoft Excel 2007 or later, the command structure or interface is different.

In earlier versions of Microsoft Excel, including Microsoft Excel 2003, Crystal Ball adds three menus to the Microsoft Excel menubar. Many menu commands are also accessible through icons on the Crystal Ball toolbar. In Microsoft Excel 2007 or later, there is a single Crystal Ball ribbon that contains all the commands. The ribbon is a combination menubar and toolbar. It contains both menus and tool icons. Figure 194 shows the Crystal Ball ribbon in Microsoft Excel 2007 or later.

Figure 194 Crystal Ball Ribbon in Microsoft Excel 2007 or later



The Crystal Ball ribbon contains five groups, described in the following sections:

• "Define Commands" on page 403

- "Run Commands" on page 403
- "Analyze Commands" on page 403
- "Tools Commands" on page 403
- "Help Commands" on page 404

Define Commands

The Define commands specify settings for the three types of Crystal Ball data cells: assumptions, decision variables, and forecasts.

You can also:

- Set cell preferences
- Select Crystal Ball data cells
- Copy, paste, and clear Crystal Ball data
- Freeze data cells to exclude them from Crystal Ball simulations

For commands and shortcut keys, see Table 25 on page 404.

Run Commands

You can use the basic Run commands to start, stop, continue, reset, and single-step through Crystal Ball simulations. You can use other Run commands to:

- Save or restore Crystal Ball simulation results
- Set run preferences, which control the number of trials, the sampling method, and other simulation options

For commands and shortcut keys, see Table 26 on page 405.

Analyze Commands

You can use the Analyze commands to:

- Create and view Crystal Ball charts
- Create reports
- Extract data for external use

For commands and shortcut keys, see Table 27 on page 405.

Tools Commands

You can use the Tools commands to access the Crystal Ball tools, Predictor, and OptQuest with an appropriate license. For commands and shortcut keys, see Table 28 on page 407.

Help Commands

The Help commands display online help, online documents, example models, the Crystal Ball About Box, and more. For commands and shortcut keys, see Table 29 on page 408.

Shortcut Key Combinations in Microsoft Excel 2007 or later

The following tables list Alt+key combinations in Microsoft Excel 2007 or later so you can use the Define, Run, Analyze, and Help commands in the Crystal Ball ribbon without using the mouse:

- Table 25, "Crystal Ball Define Commands, Microsoft Excel 2007 or Later," on page 404
- Table 26, "Crystal Ball Run Commands, Microsoft Excel 2007 or Later," on page 405
- Table 27, "Crystal Ball Analyze Commands, Microsoft Excel 2007 or Later," on page 405
- Table 28, "Crystal Ball Tools Commands, Microsoft Excel 2007 or Later," on page 407
- Table 29, "Crystal Ball Help Commands, Microsoft Excel 2007 or Later," on page 408

For shortcut keys in Microsoft Excel 2003, see "Shortcut key Combinations in Microsoft Excel 2003 or Earlier" on page 393.

Note: Oracle tries to avoid keyboard shortcut conflicts; if a key combination is the same as one on a Microsoft Excel menu, you can use the down arrow key to access the command you want.

Table 25 Crystal Ball Define Commands, Microsoft Excel 2007 or Later

Button	Command	Command Action	Keystrokes
	Define Assumption	Opens the Define Assumption dialog to define a Crystal Ball assumption cell	Alt+c, a
Y	Define Decision	Opens the Define Decision Variable dialog to define a Crystal Ball decision variable cell	Alt+c, d
	Define Forecast	Opens the Define Forecast dialog to define a Crystal Ball forecast cell	Alt+c, f
5	Сору	Copies Crystal Ball cell definitions to specified cells	Alt+c, c
(C)	Paste	Pastes copied Crystal Ball cell definitions into specified cells	Alt+c, p
9	Clear	Removes Crystal Ball cell definitions from specified cells	Alt+c, e
4	Select	Offers the Select menu, to select assumption, decision variable, and forecast cells	Alt+c, I

Button	Command	Command Action	Keystrokes
	Select All Assumptions	Selects all assumption cells on the active worksheet	Alt+c, I, a
>	Select All Decisions	Selects all decision variable cells on the active worksheet	Alt+c, I, d
A h	Select All Forecasts	Selects all forecast cells on the active worksheet	Alt+c, I, f
n/a	Select Some	Selects specified Crystal Ball data cells	Alt+c, I, s
	Freeze	Excludes specified Crystal Ball data cells from a simulation	Alt+c, g
	Cell Prefs	Opens the Cell Preferences dialog to change the appearance and other properties of Crystal Ball data cells	Alt+c, y

Table 26 Crystal Ball Run Commands, Microsoft Excel 2007 or Later

Button	Command	Command Action	Keystrokes
	Start	Begins running a Crystal Ball simulation and resumes a stopped simulation.	Alt+c, s
	Stop	Stops a simulation	Alt+c, o
44	Reset	Clears generated simulation data (Crystal Ball data cells remain defined)	Alt+c, r
	Step	Steps through a simulation one trial at a time	Alt+c, i
n/a	Trials	Sets the number of trials for a Crystal Ball simulation; edits the Trials setting in the Run Preferences dialog	Alt+c, j
n/a	Save Results	Saves current simulation results for later display and analysis	Alt+c, w, v
n/a	Restore Results	Restores saved results files for display and analysis	Alt+c, w, e
Y	Run Preferences	Opens the Run Preferences dialog for setting trials, sampling, and other run preferences	Alt+c, u

 Table 27
 Crystal Ball Analyze Commands, Microsoft Excel 2007 or Later

Button	Command	Command Action	Keystrokes
	View Charts	Offers a gallery of charts and chart window management commands; use the Down arrow (\emptyset) to choose a specific chart	Alt+c, v
	Assumption Charts	Opens the Assumption Charts dialog so you can display or manage assumption charts	Alt+c, v

Button	Command	Command Action	Keystrokes
a b	Forecast Charts	Opens the Forecast Charts dialog so you can display or manage forecast charts	Alt+c, v
-	Sensitivity Charts	Opens the Sensitivity Charts dialog so you can display existing or create new sensitivity charts	Alt+c, v
	Overlay Charts	Opens the Overlay Charts dialog so you can display existing or create new overlay charts	Alt+c, v
	Trend Charts	Opens the Trend Charts dialog so you can display existing or create new trend charts	Alt+c, v
	Scatter Charts	Opens the Scatter Charts dialog so you can display or manage scatter charts	Alt+c, v
	OptQuest Charts	If OptQuest is available, opens OptQuest charts following an optimization	Alt+c, v
profession and the second	Predictor Charts	Opens Predictor charts following a time series forecast	Alt+c, v
n/a	Open Selected Cells	Opens associated charts for selected cells	Alt+c, v, n
n/a	Cascade	Arranges open charts for easier viewing	Alt+c, v, c
n/a	Close All	Closes all open charts	Alt+c, v, I
	Create Report	Lets you create a report of simulation results by choosing from a gallery of predefined and custom report types	Alt+c, m
	Assumptions Report	Creates a report of only assumption results	Alt+c, m, a
(Decision Variables Report	Creates a report of only decision variable results	Alt+c, m, d
•	Forecasts Report	Creates a report of only forecast results	Alt+c, m, f
	Full Report	Creates a report of all simulation results	Alt+c, m, r
	Index Report	Creates a report of only Crystal Ball data cell names and locations	Alt+c, m, i
4	OptQuest Report	If OptQuest is available, creates a report of selected OptQuest results following an optimization	Alt+c, m, o

Button	Command	Command Action	Keystrokes
	Predictor Report	Creates a report of selected Predictor results following a time series forecast	Alt+c, m
	Custom Report	Creates a report with only selected simulation results	Alt+c, m, c
n/a	Report Preferences	Sets various report options	Alt+c, m, p
	Extract Data	Copies selected simulation data to a worksheet for additional analysis	Alt+c, x

 Table 28
 Crystal Ball Tools Commands, Microsoft Excel 2007 or Later

Button	Command	Command Action	Keystrokes
(OptQuest	Starts OptQuest for results optimization	Alt+c, q
man and a second	Predictor	Starts Predictor for time series analysis	Alt+c, t, p
X	More Tools	Offers the Crystal Ball general tools and integration tools for additional analysis and charting	Alt+c, t
4	Batch Fit (general)	Automatically fits selected continuous probability distributions to multiple data series	Alt+c, t, f
dith	Bootstrap (general)	Addresses the reliability and accuracy of forecast statistics	Alt+c, t, b
	Correlation Matrix (general)	Rapidly defines and automates correlations of assumptions	Alt+c, t, c
	Data Analysis (general)	Imports data directly from spreadsheets into Crystal Ball forecasts for further analysis	Alt+c, t, a
Y	Decision Table (general)	Evaluates the effects of alternate decisions in a simulation model	Alt+c, t, d
all	Scenario Analysis (general)	Displays what inputs created particular outputs	Alt+c, t, s
	Tornado Chart (general)	Individually analyzes the impact of each model variable on a target outcome	Alt+c, t, t
	2D Simulation (general)	Independently addresses uncertainty and variability using a two-dimensional simulation	Alt+c, t, 2

Button	Command	Command Action	Keystrokes
	Enterprise Performance Management (integration)	If available, enables Crystal Ballto run simulations on Enterprise Performance Management data using Smart View	Alt+c, t, m
	Strategic Finance Setup (integration)	If available, opens the Strategic Finance Setup wizard for setting up a workbook that allows Crystal Ball to run with Strategic Finance data	Alt+c, t, e
	Compare Run Modes	Compares simulation run time in Normal and Extreme speed (if available)	Alt+c, t, r

Table 29 Crystal Ball Help Commands, Microsoft Excel 2007 or Later

Button	Command	Command Action	Keystrokes
?	Help	Displays online help for Crystal Ball	Alt+c, h
6	Resources	Displays menus for Crystal Ball technical support, Oracle OTN for downloads, Crystal Ball and EPM documentation, examples, the Welcome screen, and the licensing dialog	Alt+c, n
n/a	Technical Support	Opens the Crystal Ball Web page with a link to technical support	Alt+c, n, t
n/a	Crystal Ball Documentation	Displays a list of available online documentation for Crystal Ball	Alt+c, n, c
n/a	Oracle EPM Documentation	Displays the OTN index for Oracle Enterprise Performance Management documentation, including Crystal Ball	Alt+c, n, d
n/a	Examples Guide	Displays a list of available example models so you can load the choice into Crystal Ball	Alt+c, n, e
n/a	Welcome Screen	Displays a Welcome Screen that lets you automatically activate the process capability features for quality programs such as Six Sigma or set percentile preferences frequently used in the Oil and Gas industry.	Alt+c, n, w
n/a	Licensing	Displays the Activate a License dialog for entering a Crystal Ball serial number and activating a license	Alt+c, n, l
0	About [Crystal Ball]	Offers version and other information about the current Crystal Ball release including the username	Alt+c, b

The following information concerns menus and commands in the Distribution Gallery and chart windows. This information is the same for all supported versions of Windows:

- Table 16, "Crystal Ball General Chart Menu Commands, All Supported Versions of Microsoft Excel," on page 397
- Table 17, "Assumption-specific Chart Menu Commands, All Supported Versions of Microsoft Excel," on page 397
- Table 18, "Forecast-specific Chart Menu Commands, All Supported Versions of Microsoft Excel," on page 398
- Table 19, "Overlay-specific Chart Menu Commands, All Supported Versions of Microsoft Excel," on page 398

- Table 20, "Sensitivity-specific Chart Menu Commands, All Supported Versions of Microsoft Excel," on page 398
- Table 21, "Trend-specific Chart Menu Commands, All Supported Versions of Microsoft Excel," on page 399
- Table 22, "Scatter-specific Chart Menu Commands, All Supported Versions of Microsoft Excel," on page 399

Distribution Gallery Shortcut Keys

These are described in "Distribution Gallery Shortcut Keys" on page 400.

Chart Preference Shortcut Keys

These are described in "Chart Preference Shortcut Keys" on page 401.

Compatibility and File Conversion Issues

Microsoft Excel 2007 or later workbooks can be saved in several file formats that are significantly different from those for previous versions of Microsoft Excel. This version of Crystal Ball has been carefully designed to preserve Crystal Ball data in existing workbooks as long as you follow a few simple rules when opening and saving files created in previous versions of Microsoft Excel.

In general:

- 1. Be sure Crystal Ball is loaded before opening a workbook with extension .xls that contains Crystal Ball data or before saving a model to any format in Microsoft Excel 2007 or later.
- 2. Always save files in .xls format to share with users of Crystal Ball who are using Microsoft Office 2003, XP, or 2000.

For detailed information, see the migration appendix in the Oracle Crystal Ball Installation and Licensing Guide.

Files from earlier versions of Crystal Ball — including the Crystal Ball Example files — are converted automatically when you save them in the current version of Crystal Ball on Microsoft Excel 2007 or later or earlier versions of Microsoft Excel. Files from earlier versions are opened in compatibility mode. [Compatibility Mode] is displayed after the workbook name in the title bar.

Note: When you open the Example files in Windows Vista, they are read-only. To edit or customize them, save them under another name before you run a simulation.

G

Bibliography

In This Appendix

Bootstrap	411
Monte Carlo Simulation	411
Probability Theory and Statistics	412
Random Variate Generation Methods	412
Specific Distributions	413
Tornado Charts and Sensitivity Analysis	413
Two-Dimensional Simulation.	414
Uncertainty Analysis	414
Sequential Sampling with SIPs	414

Bootstrap

Efron, Bradley, and Robert J. Tibshirani. "An Introduction to the Bootstrap," Monographs on Statistics and Applied Probability, vol. 57. New York: Chapman and Hall, 1993.

Mooney, C. Z., and R.D. Duval. "Bootstrapping: A Nonparametric Approach to Statistical Inference," Sage University Paper Series on Quantitative Applications in the Social Sciences, series no. 07-095. Newbury Park, CA: Sage, 1993.

Monte Carlo Simulation

Burmaster, D.E., and P.D. Anderson, "Principles of Good Practice for the Use of Monte Carlo Techniques in Human Health and Ecological Risk Assessments," Risk Analysis 14(4): 477-481 (1994).

Hammersley, J.M., and D.C. Handscomb. Monte Carlo Methods. New York: Chapman and Hall, 1965.

Kalos, Malvin H., and Paula A. Whitlock. Monte Carlo Methods. New York: John Wiley & Sons, 1986.

Morgan, Byron J.T. Elements of Simulation. Portland, ME: Chapman and Hall, 1984.

Rubinstein, Reuven Y. Simulation and the Monte Carlo Method. New York: John Wiley & Sons, 1981.

Thommes, M.C. Proper Spreadsheet Design. Boston: Boyd and Fraser Publishing Co., 1992.

Thompson, K.M., D.E. Burmaster, and E.A.C. Crouch. "Monte Carlo Techniques for Quantitative Uncertainty Analysis in Public Health Risk Assessments," Risk Analysis 12(1): 53-63 (1992).

Yakowitz, Sidney J. Computational Probability and Simulation. Reading, MA: Addison-Wesley Publishing Co., 1977.

Probability Theory and Statistics

Alder, Henry L., and Edward B. Roessler. Introduction to Probability and Statistics, 6th Ed. San Francisco: W.H. Freeman & Co, 1977.

Biswas, Suddhendu. Topics in Statistical Methodology. New York: John Wiley & Sons, Inc., 1991.

Fraser, D.A.S. Probability and Statistics: Theory and Applications. North Scituate, MA: Duxbury Press, 1976.

Johnson, Norman L., and Samuel Kotz. Distributions in Statistics: Continuous Multivariate Distributions. New York: John Wiley & Sons, Inc., 1972.

Kapur, K.C., and L.R. Lamberson. Reliability in Engineering Design. New York: John Wiley & Sons, 1977.

Larson, Harold J. Introduction to Probability Theory and Statistical Inference, 2nd Ed. New York: John Wiley & Sons, 1974.

Logothetis, N., and V. Rothschild. Probability Distributions. New York: John Wiley & Sons, Inc., 1986.

Spurr, William, and Charles P. Bonini. Statistical Analysis for Business Decisions. Homewood, IL: Richard D. Irwin, Inc., 1973.

Wadsworth, George P., and Joseph G. Bryan. Introduction to Probability and Random Variables. New York: McGraw-Hill Book Co., 1960.

Winkler, Robert L., and William L. Hays. Statistics: Probability, Inference, and Decision. New York: Holt, Rinehart, and Winston, 1975.

Random Variate Generation Methods

Devroye, Luc. Non-Uniform Random Variate Generation. New York: Springer-Verlag, 1986.

Iman, Ronald L., and W.J. Conover. "A Distribution-Free Approach to Inducing Rank Correlation Among Input Variables," Communications in Statistics B11(3):1982.

Iman, Ronald L., and J.M. Davenport. "Rank Correlation Plots for Use With Correlated Input Variables," Communications in Statistics B11(3): (1982).

Iman, Ronald L., and J.M. Davenport. Latin Hypercube Sampling (A Program User's Guide). Technical Report SAND79-1473, Sandia National Laboratories, 1980.

Kennedy, William J., Jr., and James E. Gentle. Statistical Computing. New York: Marcel Dekker, Inc., 1980.

Knuth, Donald E. The Art of Computer Programming, Vols I-III. Reading, MA: Addison-Wesley Publishing Co., 1969.

Newman, Thomas G., and Patrick L. Odell. The Generation of Random Variates. New York: Hafner Publishing Co., 1971.

Press, William H., et al. Numerical Recipes in C, 2nd Ed. Cambridge, England: Cambridge University Press, 1993.

Odeh, R.E., and J.O. Evans. "Percentage Points of the Normal Distribution," Applied Statistics. London: Royal Statistical Society, 1974.

Sedgewich, Robert. Algorithms. Reading, MA: Addison Wesley Publishing Co., 1983.

Specific Distributions

Extreme Value Distribution

Castillo, Enrique. Extreme Value Theory in Engineering. London: Academic Press, 1988.

Lognormal Distribution

Aitchison, J., and J.A. Brown. The Lognormal Distribution. New York: Cambridge University Press, 1973.

Weibull Distribution

King, James R. Probability Charts for Decision Making, Rev. Ed. New York: Industrial Press, Inc., 1981.

Henley, Ernest J., and Hiromitsu Kumamoto. Reliability Engineering and Risk Assessment. Englewood Cliffs, NJ: Prentice Hall, Inc., 1981.

Tornado Charts and Sensitivity Analysis

Clemen, Robert T. Making Hard Decisions: An Introduction to Decision Analysis, 2nd Ed. Belmont, CA: Duxbury Press, 1997.

Morgan, M. Granger, and Max Henrion; with a chapter by Mitchell Small. Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis. New York: Cambridge University Press, 1990.

Two-Dimensional Simulation

Burmaster, David E., and Andrew M. Wilson. "An Introduction to Second-Order Random Variables in Human Health Risk Assessments," Human and Ecological Risk Assessment 2(4): 892-919 (1996).

Frey, H. Christopher, and David S. Rhodes. "Characterizing, Simulating, and Analyzing Variability and Uncertainty: An Illustration of Methods Using an Air Toxics Emissions Example," Human and Ecological Risk Assessment 2(4): 762-797 (1996).

Hoffman, F.O., and J.S. Hammonds. "Propagation of Uncertainty in Risk Assessments: The Need to Distinguish Between Uncertainty Due to Lack of Knowledge and Uncertainty Due to Variability" Risk Analysis 14(5): 707-712 (1994).

Rai, S. N., D. Drewski, and S. Bartlett. "A General Framework for the Analysis of Uncertainty and Variability in Risk Assessment," Human and Ecological Risk Assessment 2(4): 972-989 (1996).

Uncertainty Analysis

Morgan, M. Granger, and Max Henrion; with a chapter by Mitchell Small. Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis. New York: Cambridge University Press, 1990.

Hoffman, F.O., and J.S. Hammonds. "Propagation of Uncertainty in Risk Assessments: The Need to Distinguish Between Uncertainty Due to Lack of Knowledge and Uncertainty Due to Variability" Risk Analysis 14(5): 707-712 (1994).

International Atomic Energy Agency, "Evaluating the Reliability of Predictions Made Using Environmental Transfer Models," IAEA Safety Series No. 100:1-106. STI/PUB/835. Vienna, Austria, 1989.

Kelton, W. David, and Averill M. Law. Simulation Modeling & Analysis, 3rd Ed. New York: McGraw-Hill, Inc. 2000.

Sequential Sampling with SIPs

The SIPs format described in "Sequential Sampling with Custom Distributions" on page 291 conforms to the Stochastic Library approach set forth in the following articles:

Savage, Sam, Stefan Scholtes, and Daniel Zweidler. "Probability Management, Part 1," OR/MS Today 33(1): 21-28 (February 2006).

Savage, Sam, Stefan Scholtes, and Daniel Zweidler. "Probability Management, Part 2," OR/MS Today 33(2): 60-66 (April 2006).

These articles are available for download at:

http://www.probabilitymanagement.org

Glossary

algorithm A rule that specifies how to solve a particular problem.

assumption An estimated value or input to a spreadsheet model.

assumption cell A value cell in a spreadsheet model that has been defined as a probability distribution.

base case The value in a Crystal Ball assumption, decision variable, or forecast cell at the start of a simulation.

CDF Cumulative distribution function that represents the probability that a variable will fall at or below a given value.

certainty bands In a trend chart, a graphic depiction of a particular certainty range for each forecast.

certainty level The percentage of values in the certainty range compared to the number of values in the entire range.

certainty range The linear distance for the set of values between the certainty grabbers on the forecast chart.

coefficient of variability also coefficient of variance or coefficient of variation. A measure of relative variation that relates the standard deviation to the mean. Results can be represented in percentages for comparison purposes.

continuous probability distribution A probability distribution that describes a set of uninterrupted values over a range. In contrast to the discrete distribution, the continuous distribution assumes there is an infinite number of possible values.

correlation In Crystal Ball, a dependency that exists between *assumption cells*.

correlation coefficient A number between -1 and 1 that specifies mathematically the degree of positive or negative correlation between *assumption cells*. A *correlation* of 1 indicates a perfect positive correlation, minus 1 indicates a perfect negative correlation, and 0 indicates there is no correlation.

cumulative frequency distribution A chart that shows the number or proportion (or percentage) of values less than or equal to a given amount.

decision variable A Crystal Ball variable in the model that you can control.

decision variable cell Cells that contain the values or variables that are within the control to change. The decision variable cells must contain simple numeric values, not formulas or text.

deterministic model Another name for a *spreadsheet model* which yields single-valued results.

discrete probability distribution A probability distribution that describes distinct values, usually integers, with no intermediate values. In contrast, the continuous distribution assumes there is an infinite number of possible values.

display range The linear distance for the set of values displayed on the forecast chart.

dominant A relationship between distributions in which one distribution's values for all percentile levels are higher than another's. *See also subordinate*.

entire range The linear distance from the minimum *forecast value* to the maximum forecast value.

forecast A statistical summary of the assumptions in a spreadsheet model, output graphically or numerically.

forecast cell Cells that contain formulas that refer to one or more assumption and decision variable cells and combine the values in the assumption, decision, and other cells to calculate a result.

forecast definition The forecast name and parameters assigned to a cell in a Crystal Ball dialog.

forecast filtering A process by which Crystal Ball discards forecast values outside or inside a specified range.

forecast formula A formula that has been defined as a forecast cell.

forecast value also trial A value calculated by the forecast formula during an iteration. These values are kept in a list for each forecast, and are summarized graphically in the forecast chart and numerically in the descriptive statistics.

formula cell A cell that contains a mathematical formula.

frequency also frequency count The number of times a value recurs in a group interval.

frequency distribution A chart that graphically summarizes a list of values by subdividing them into groups and displaying their frequency counts.

goodness-of-fit A set of mathematical tests performed to find the best fit between a standard probability distribution and a data set.

grabber also certainty grabber or truncation grabber A control that lets you use the mouse to change values and settings.

group interval A subrange of a distribution that allows similar values to be grouped together and given a frequency count.

iteration also trial A three-step process in which Crystal Ball generates random numbers for assumption cells, recalculates the spreadsheet model or models, and displays the results in a forecast chart.

kurtosis The measure of the degree of peakedness of a curve. The higher the kurtosis, the closer the points of the curve lie to the mode of the curve. A normal distribution curve has a kurtosis of 3.

Latin Hypercube sampling In Crystal Ball, a sampling method that divides an assumption's probability distribution into intervals of equal probability. The number of intervals corresponds to the Minimum Sample Size option available in the Run Preferences dialog. A random number is then generated for each interval.

Compared with conventional Monte Carlo sampling, Latin Hypercube sampling is more precise because the entire range of the distribution is sampled in a more even, consistent manner. The increased accuracy of this method comes at the expense of added memory requirements to hold the full Latin Hypercube sample for each assumption. (See "Setting Sampling Preferences" on page 79.)

mean The familiar arithmetic average of a set of numerical observations: the sum of the observations divided by the number of observations.

mean standard error The Standard Deviation of the distribution of possible sample means. This statistic gives one indication of how accurate the simulation is.

median The value midway (in terms of order) between the smallest possible value and the largest possible value.

mode That value which, if it exists, occurs most often in a data set.

model sensitivity The overall effect that a change in an assumption cell produces in a forecast cell. This effect is solely determined by the formulas in the spreadsheet model.

Monte Carlo simulation A system which uses random numbers to measure the effects of uncertainty in a spreadsheet model.

outlying values Values generated during a simulation on the extreme end of a distribution that are excluded from the display range.

PDF Probability density function that represents the probability that an infinitely small variable interval will fall at a given value.

probabilistic model A system whose output is a distribution of possible values. In Crystal Ball, this system includes a spreadsheet model (containing mathematical relationships), probability distributions, and a mechanism for determining the combined effect of the probability distributions on the model's output (Monte Carlo simulation).

probability (Classical Theory) The likelihood of an event.

probability distribution also distribution A set of all possible events and their associated probabilities.

random number A mathematically selected value which is generated (by a formula or selected from a table) to conform to a probability distribution.

random number generator A method implemented in a computer program that is capable of producing a series of independent, random numbers.

range The difference between the largest and smallest values in a data set.

rank correlation also Spearman's rank correlation $\,A\,$ method whereby assumption values are replaced with their ranking from lowest value to highest value using the integers 1 to $\,N\,$ prior to computing the correlation coefficient. This method allows the distribution types to be ignored when correlating assumptions.

relative probability also relative frequency A value, not necessarily between 0 and 1, that indicates probability when used in a proportion.

reverse cumulative frequency distribution A chart that shows the number or proportion (or percentage) of values greater than or equal to a given amount.

risk The uncertainty or variability in the outcome of some event or decision.

seed value The first number in a sequence of random numbers. A given seed value produces the same sequence of random numbers every time you run a simulation.

sensitivity The amount of uncertainty in a forecast cell that is a result of both the uncertainty (probability distribution) and model sensitivity of an assumption cell.

sensitivity analysis The computation of a forecast cell's sensitivity with respect to the assumption cells.

skewed An asymmetrical distribution.

skewed, negatively A distribution in which most of the values occur at the upper end of the range.

skewed, positively A distribution in which most of the values occur at the lower end of the range.

skewness The amount a curve differs from a normal, symmetrical distribution. The greater the degree of *skewness*, the more points of the curve lie to either side of the peak of the curve. A normal distribution curve, having no *skewness*, is symmetrical. Skewness is computed by finding the third moment about the mean and dividing by the cube of the standard deviation.

spreadsheet model Any spreadsheet that represents an actual or hypothetical system or set of relationships.

standard deviation The square root of the variance for a distribution. A measurement of the variability of a distribution, i.e., the dispersion of values around the mean. (See formulas in "Standard deviation" in the "Statistical Definitions" chapter of the online *Oracle Crystal Ball Statistical Guide*.)

subordinate A relationship between distributions in which one distribution's values for all percentile levels are lower than another's. *See also dominant*.

trial also iteration A three-step process in which Oracle Crystal Ball, Fusion Edition generates random numbers for assumption cells, recalculates the spreadsheet model or models, and displays the results in a forecast chart.

trial as used to describe a parameter in certain probability distributions The number of times a given experiment is repeated.

value cell A cell that contains a simple numeric value.

variable A quantity that can assume any one of a set of values and is usually referenced by a formula.

variance The square of the standard deviation; i.e., the average of the squares of the deviations of a number of observations from their mean value.

Variance can also be defined as a measure of the dispersion, or spread, of a set of values about a mean. When values are close to the mean, the variance is small. When values are widely scattered about the mean, the variance is large. (See formulas in "Variance" in the "Statistical Definitions" chapter of the online *Oracle Crystal Ball Statistical Guide*.)

virtual memory Memory which uses the hard drive space to store information after you run out of random access memory. Virtual memory supplements the random access memory.

workbook a Microsoft Excel file composed of at least one worksheet.

worksheet a Microsoft Excel file in which you work and store the data. A worksheet is part of a workbook.

Index

Numbers	assumption charts, 28
2D Simulation tool, 32	assumption parameters
2D simulation tool, 214	alternate, 45
	assumptions
Λ.	correlated, 49
A absolute call references 44, 45	correlating, 49
absolute cell references, 44, 45	defining, 40, 316
accessibility activating, 388	defining from assumptions, 114
cell preferences, 389	entering, 40
charts, 390	freezing, 82
Control Panel, 391	selecting and reviewing, 72
documentation, 387	assumptions, fitting to distributions, 340
large fonts, 391	Auto Extract preferences, 70
screen flicker, 392	
support, 387	В
Tab and Arrow keys, 390	Batch Fit
Accessibility mode	example, 177
activating, 389	tool, 173
activating accessibility, 388	Batch Fit tool, 31
activating process capability features, 337, 351	bibliography, 413
alternate parameters, 45, 235	Bootstrap
analysis, simulation, 342, 356	example, 199
Analyze menu, 393	tool, 195
Analyze phase	Bootstrap tool, 31
liquid pump example, 350	
loan processing example, 336	C
analyzing simulation results, 95	capability metrics, 90, 365
analyzing time-series data, 33	liquid pump example, 356, 358
Anderson-Darling test, 47, 113	loan processing example, 344
Apply To, 124	categories
applying settings, 124	creating distribution, 59
assumption	deleting distribution, 60
preferences, 52	rearranging distribution, 60
assumption cells	sharing distribution, 61, 62, 63
defined, 39	viewing and editing distribution properties, 60
liquid pump example, 354	categories of distributions, 59
loan processing example, 339	Category pane of the Distribution Gallery, 54
	0 / 1

	11 . 11 . 1 . 20
caveats for tornado and spider charts, 195	comparing distributions, 286
Cell menu, 393	compatibility, file, 409
cell preferences, 74	consulting, 17
accessibility defaults, 389	contact information, 17
cell references, 44	continuing a simulation, 83
absolute, 45	continuous
dynamic and static, 44	distributions, 232
relative, 44	Control Panel
cells	Crystal Ball, 85
freezing, 82	showing after reset, 85
certainty level, 311, 330	Control Panel accessibility, 391
certainty levels	Control Panel, Crystal Ball, 313
forecast chart, 98	Control Panel, hiding, 85
trend charts, 141	Control phase
chart preferences, 115	loan processing example, 337
shortcut keys, 401	conversion, file, 409
chart preferences shortcut keys, 116	copying charts, 125
charts, 29	copying distributions, 56
	correlated assumption limit, 384
assumption, 28 closing, 127	÷
6	correlating assumptions, 49
copying, 125	Correlation Chart, 51
customizing, 115, 332	Correlation Matrix
deleting, 127	example, 184
displaying, 125	tool, 181
forecast, 24	Correlation Matrix tool, 32
interpreting, 329	correlations
managing, 124	accuracy, 381
OptQuest, 29	defined, 182
overlay, 25	specifying, 49
overlay with 2D Simulation, 220	creating distribution categories, 59
pasting, 126	creating distributions, 56
preferences, 115	Crystal Ball
printing, 126	Monte Carlo simulation, 23
scatter, 27, 153	quality support, 368
sensitivity, 26, 142	starting, 337, 351
spider, 189	system requirements, 14
tornado, 187	Crystal Ball Control Panel, 85, 313
trend, 26	Crystal Ball menus, 393
trend with 2D Simulation, 220	Crystal Ball toolbar buttons, 392, 402
charts and accessibility, 390	Crystal Ball tools, 31
chi-square distribution, 247	custom distribution, 270
chi-square test, 47, 113	important notes, 284
closing charts, 127	sequential sampling, 291
commands	custom distributions, 321
menu, 392, 393	customizing
ribbon, 402	charts, 332
Compare Run Modes tool, 33, 300	customizing charts, 115

D	DFSS liquid pump model, using, 351
data	dialogs
editing, 70	Select, 74
entering custom, 279	discrete
extracting, 30	distributions, 232, 233
grouped, 276	ranges, 276
Data Analysis	display range preferences, 102
example, 225	displaying charts, 125
Data Analysis tool, 32	distribution
Decision Table	preferences, 52
example, 206	distribution categories, 59
tool, 203	creating, 59
Decision Table tool, 32	deleting, 60
decision variable cells, 359	rearranging, 60
defined, 39	sharing, 61, 62, 63
decision variables	distribution category properties
defining, 65	editing and viewing, 60
optimizing, 33	distribution fitting, 46
default settings, 124	overlay charts, 134, 135
Define Forecast dialog, 67	overview, 46
Define phase	Distribution Gallery
liquid pump example, 349	Category pane, 54
loan processing example, 335	Description pane, 55
defining	displaying, 53
assumption cells, 40	Distribution pane, 54
assumptions, 316	managing categories, 59
forecasts, 326	managing distributions, 55
defining assumptions, 40	menubar and buttons, 54
defining assumptions from forecasts, 114	using, 53
deleting charts, 127	window, 53
deleting distribution categories, 60	Distribution Gallery menus, 393
deleting distributions, 58	Distribution Gallery, Basic, 317
Description pane, 55	Distribution pane, 54
Design for Six Sigma, 367	distributions, 235
defined, 348	binomial, 320
phases, 348	continuous, 232
tutorial, 348	copying, 56
Design phase	creating, 56
liquid pump example, 350	custom, 321
DFSS, 367	deleting, 58
defined, 348	discrete, 233
DMADV, 348	discrete or continuous, 232
phases, 348	fitting to a forecast, 113
tutorial, 348	managing, 55
DFSS example	modifying, 57
liquid pump example, 348	normal, 324
DFSS Fluid Pump.xls, 352	parameters, 235, 288

pasting, 56	distribution types, 104
triangular, 319	format, 103
truncating, 285, 324	groups, 112
uniform, 317	Forecast Preference dialog, 67
distributions as Microsoft Excel functions, 290	forecast preferences, 67
DMADV, 348	Forecast Preferences dialog
DMAIC, 335	Auto Extract tab, 70
documentation accessibility, 387	Filter tab, 69
dynamic cell referencing, 44	Forecast Window tab, 68
,	Precision tab, 68
-	Forecast Window preferences, 68
E and the state of	forecasts
editing distribution category properties, 60	additional features, 112
emailing distribution categories, 62	defining, 66
entering custom data, 279	defining assumptions from, 114
Enterprise Performance Management, 33	fitting distributions, 46, 113
Enterprise Performance Management integration	percentiles view, 107
tool, 172	selecting and reviewing, 72
Enterprise Performance Management tool, 33	statistics view, 107
erlang distribution, 247	understanding, 97
examples	forecasts, defining, 326
Batch Fit, 177	formulas in parameters, 45
Bootstrap, 199	freezing assumptions, 82
Correlation Matrix, 184	freezing cells, 82
Data Analysis, 225	functions in Crystal Ball spreadsheets, 90
Decision Table, 206	Futura Apartments model, 309
Scenario Analysis, 211	,
Tornado Chart, 192	
Two-dimensional Simulation, 217	G
expanded Define Forecast dialog, 67	global macros, 94
extract data, 165	grouped data, 276
extracting data, 30	groups, 112
	guidelines
F	analyzing simulation results, 95
file compatibility and conversion, 409	
Filter preferences, 69	Н
fitting distributions	help, 16
locking parameters, 49	Help menu, 393
overview, 46	hidden factory, 368
fitting distributions to assumptions, 340	how this manual is organized, 14
fitting distributions to data, 46	110 (1 4110 1111111111 10 01 911111111 11 11
fitting distributions to forecasts, 113	
forecast cells, 341	I
defined, 39	Improve phase
liquid pump model, 354	loan processing, 344
forecast charts, 24	loan processing example, 336
analyzing, 95	integration, 33, 172
	integration tools, 33, 172

interpreting charts, 329	managing distributions
	Distribution Gallery, 55
K	marker lines, limits and target, 342, 356
	Measure phase
chartcut 302 303 402	liquid pump example, 350
shortcut, 392, 393, 402	loan processing example, 336
Kolmogorov-Smirnov test, 47, 114	menu commands, 392, 393
	menus
L	Analyze, 393
large fonts, 391	Cell, 393
Latin hypercube sampling	Crystal Ball, 393
accuracy, 381	Distribution Gallery, 393
Lean principles, 367	Help, 393
liquid pump example	Run, 393
adjusting, 358	Microsoft Excel charts
Analyze, 350	pasting to, 125
capability metrics, 358	Microsoft Excel functions and distributions, 290
Define, 349	model sensitivity, 142
Design, 350	modifying distributions, 57
illustrated, 352	monotonic relationships, 188
limits and target, 355	Monte Carlo simulation, 312
Measure, 350	benefits, 23
parts, 353	defined, 22
Validate, 351	in Crystal Ball, 23
viewing, 352	
Loan Process.xls, 338	N
loan processing example, 335	N
Analyze phase, 336	normal distributions, 324
Control phase, 337	
Define phase, 335	0
Improve phase, 336, 344	online help, 16
Measure phase, 336	optimization goals, 360
metrics, 344	optimizing decision variables, 33
using, 337	Options preferences, 81
loan processing model, 338	OptQuest, 33
locking parameters, 49	running, 361
lognormal distribution	using, 358
parameter sets, 254	OptQuest charts, 29
LSL, 355	outliers, 97
	outlying values, 97
B.4	overlay chart
M	2D Simulation, 220
macros	overlay charts, 25
global, 94	customizing, 134
user-defined, 91	distribution fitting, 134, 135
managing categories	overview, 130
Distribution Gallery, 59	Overlay Example, 132
managing charts, 124	,

P	and Crystal Ball, 368
p(N/C)-above, 342	·
parameters	D
alternate, 45, 235	R
formulas, 45	range names in parameters, 45
locking, 49	rearranging distribution categories, 60
range names, 45	references, 413
Parameters menu, 235, 323	absolute cell, 45
parameters, distribution, 288	cell, 44
pasting charts, 125, 126	dynamic and static, 44
pasting distributions, 56	relative cell, 44
percentiles	relative cell references, 44
view, 107	relative cell referencing, 44
PPM-above, 342	relative probability, 272
PPM-total, 342	reports, 29
precision control	overview, 161
defining with forecast, 66	requirements, 14
displayed percentiles, 107	resetting a simulation, 83
displayed statistics, 107	restoring results, 88
increasing accuracy, 382	restoring simulation results, 89
Precision preferences, 68	results, saving and restoring, 88
predicting trends, 33	ribbon commands, 402
Predictor, 29, 33	Run menu, 393
Predictor charts, 29	run preferences, 78
preferences for charts, 115	running a simulation, 83, 342, 356
Preferences menu	running OptQuest, 361
distributions, 52	running simulations, 310, 328
printing charts, 126	
probability, 230	S
probability distributions, 235	sample size, 384
chi-square, 247	Sampling preferences, 79
comparing, 286	saving and restoring simulation results, 88
custom, 270	saving results, 88
erlang, 247	saving simulation results, 88
selecting, 233	scatter charts, 27
process capability features, 33, 365	creating, 154
activating, 337, 351	customizing, 156
profit, determining, 311	overview, 153
properties	Scenario Analysis
distribution category, 60	example, 211
publishing distribution categories, 61	running, 211
pump example, 348	selecting a target forecast, 210
pump example, 5 to	starting, 210
	tool, 209
Q	using, 210
quality features, 33, 365	Scenario Analysis tool, 32
quality improvement, loan processing example, 344	screen capture notes, 16
quality methodologies	sercen capture notes, 10

screen flicker, 392	liquid pump example, 355
second-order variables, 222	spider chart, 189
Select dialog, 74	spreadsheet functions in Crystal Ball, 90
selecting assumptions, forecasts, and other data types,	starting Crystal Ball, 337, 351
130	starting simulations, 83
selecting some data, 74	static cell referencing, 44
sending distribution categories, 62	Statistics preferences, 82
sensitivity analysis	statistics view, 107
accuracy, 381	stopping a simulation, 83
calculation, 145	Strategic Finance integration tool, 172
sensitivity chart, 342, 357	Strategic Finance tool, 33
sensitivity charts, 26	subscribing to distribution categories, 61
creating, 143	support
customizing, 147	with hearing impairments, 387
overview, 142	support resources, 17
sequential sampling, 291	system requirements, 14
setting defaults, 124	
setting forecast preferences, 67	Т
setting preferences	tables of 2D Simulation results, 219
Options, 81	target, 355
Sampling, 79	t
Statistics, 82	t, 342
Trials, 79	technical support, 17
sharing distribution categories, 61, 62, 63	time-series analysis, 33
shortcut keys, 392, 393, 402	tool for comparing run modes, 300
shortcut keys for charts, 116, 401	toolbar buttons, 392, 402
simulation	tools, 33, 172
analyzing, 342, 356	2D Simulation, 32, 214
running, 342, 356	Batch Fit, 31, 173
simulation results	Bootstrap, 31, 195
restoring, 89	Compare Run Modes, 33
saving, 88	Correlation Matrix, 32, 181
using, 89	Data Analysis, 32
simulation results, saving and restoring, 88	Decision Table, 32, 203
simulations	Scenario Analysis, 32, 209
accuracy, 381	Strategic Finance, 33
running, 310, 328	Tornado Chart, 32, 187
speed, 383	Tornado Chart
starting, 83	description, 187
Single Step, 84	example, 192
SIPs, 291	tool, 187
Six Sigma, 367	uses, 187
DMAIC, 335	Tornado Chart tool, 32
Six Sigma features, 33, 365	training, 17
SIZE	trend analysis, 33
sample, 384	trend chart with 2D Simulation, 220
specification limits	trend charts, 26

```
certainty levels, 141
  creating, 137
  customizing, 137
  overview, 136
  value axis, 141
Trials preferences, 79
triangular distributions, 319
truncating distributions, 285, 324
Tutorial 1, loan processing example, 335
Tutorial 2, pump design with DFSS, 348
Two-dimensional Simulation
  example, 217
U
uncertainty, 214
uniform distributions, 317
unpublishing distribution categories, 63
user-defined macros, 91
using OptQuest, 358
using simulation results, 89
USL, 355
  L
     L, 342
V
Validate phase
  liquid pump example, 351
value axis settings, 141
variability
  vs. uncertainty, 214
variables
  defining decision, 65
  second-order, 222
viewing distribution category properties, 60
virtual memory, 384
Vision Research model, 316
W
Welcome screen, 35
who should use this program, 14
Z
Z scores, 342
```