

ORACLE® CRYSTAL BALL DECISION OPTIMIZER,
FUSION EDITION

RELEASE 11.1.1.3.00

OPTQUEST USER'S GUIDE

ORACLE
ENTERPRISE PERFORMANCE
MANAGEMENT SYSTEM

Crystal Ball Decision Optimizer OptQuest User's Guide, 11.1.1.3.00

Copyright © 1988, 2009, Oracle and/or its affiliates. All rights reserved.

Authors: EPM Information Development Team

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States 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, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Chapter 1. Welcome	9
Introduction	9
How This Manual is Organized	9
Screen Capture Notes	10
Getting Help	10
Additional Resources	11
Chapter 2. Overview	13
Introduction	13
What OptQuest Does	13
How OptQuest Works	14
About Optimization Models	15
Optimization Objectives	16
Forecast Statistics	16
Minimizing or Maximizing	16
Requirements	17
Decision Variables	17
Constraints	18
Model and Solution Feasibility	19
Efficient Frontier Analysis	19
Efficient Portfolios	20
OptQuest and Process Capability	21
Chapter 3. Setting Up and Optimizing a Model	23
Introduction	23
Overview	23
For Users of OptQuest Versions Earlier Than 11.1.1.x	24
Developing a Crystal Ball Optimization Model	24
Developing the Worksheet	24
Defining Assumptions, Decision Variables, and Forecasts	25
Setting Crystal Ball Run Preferences	25
Starting OptQuest	26

Selecting the Forecast Objective	26
Selecting Decision Variables to Optimize	27
Specifying Constraints	28
Specifying Constraints in Simple Entry Mode	29
Specifying Constraints in Advanced Entry Mode	29
Constraints Editor and Related Buttons	31
Constraint Rules and Syntax	32
Constraints and Cell References in Advanced Entry Mode	33
Constraint Types	34
Setting Options	34
Advanced Options	35
Running Optimizations	35
OptQuest Control Panel Buttons and Commands	36
OptQuest Results Window	36
Interpreting the Results	40
Viewing a Solution Analysis	40
Running a Longer Simulation of the Results	42
Printing OptQuest Results	42
Viewing Charts in Crystal Ball	42
Creating OptQuest Reports	42
Extracting OptQuest Data	44
Saving optimization models and settings	45
Closing OptQuest	45
Setting Up Efficient Frontier Analysis in OptQuest	46
Efficient Frontier Variable Bound Example	46
Transferring Settings from .opt Files	46
Learning More About OptQuest	48
Chapter 4. OptQuest Tutorials	49
Introduction	49
Tutorial 1 — Futura Apartments Model	49
Running OptQuest	50
Tutorial 2 — Portfolio Allocation Model	54
Problem Description	55
Using OptQuest	55
Chapter 5. Examples Using OptQuest	67
Overview	67
Product Mix	69
Product Mix Problem Statement	69

Product Mix Spreadsheet Model	70
Product Mix OptQuest Solution	70
Hotel Design and Pricing Problem	72
Hotel Design Problem Statement	72
Hotel Design Spreadsheet Model	73
Hotel Design OptQuest Solution	73
Budget-constrained Project Selection	75
Project Selection Problem Statement	75
Project Selection Spreadsheet Model	76
Project Selection OptQuest Solution	77
Groundwater Cleanup	79
Groundwater Cleanup Problem Statement	79
Groundwater Cleanup Spreadsheet Model	79
Groundwater Cleanup OptQuest Solution	81
Oil Field Development	83
Oil Field Development Problem Statement	83
Oil Field Development Spreadsheet Model	84
Oil Field Development OptQuest Solution	85
Portfolio Revisited	86
Portfolio Revisited Problem Statement	87
Portfolio Revisited Method 1: Efficient Frontier Optimization	87
Portfolio Revisited Method 2: Multi-objective Optimization	89
Tolerance Analysis	92
Tolerance Analysis Problem Statement	93
Tolerance Analysis Spreadsheet Model	93
Tolerance Analysis OptQuest Solution	95
Inventory System Optimization	96
Inventory System Problem Statement	96
Inventory System Spreadsheet Model	98
Inventory System OptQuest Solution	99
Drill Bit Replacement Policy	102
Drill Bit Replacement Problem Statement	102
Drill Bit Replacement Spreadsheet Model	103
Drill Bit Replacement OptQuest Solution	104
Gasoline Supply Chain	105
Gasoline Supply Chain Statement of Problem	106
Gasoline Supply Chain Spreadsheet Model	107
Gasoline Supply Chain OptQuest Solution	108

Appendix A. Optimization Tips and Notes	111
Introduction	111
Model Types	111
Optimization Models Without Uncertainty	111
Optimization Models With Uncertainty	112
Discrete, Continuous, or Mixed Models	113
Linear or Nonlinear Models	114
Factors That Affect Optimization Performance	114
Simulation Accuracy	115
Number of Decision Variables	116
Base Case Values	116
Bounds and Constraints	117
Requirements	117
Complexity of the Objective	118
Simulation Speed	118
Precision Control	118
Sensitivity Analysis Using a Tornado Chart	119
Maintaining Multiple Optimization Settings for a Model	120
Other OptQuest Notes	121
Automatic Resets of Optimizations	121
Constraint Formula Limitations	121
Minor Limit Violations With Continuous Forecasts	122
Solutions Still Ranked Even With No Feasible Solution	122
Referenced Assumption and Forecast Cells	122
Decision Variables and Ranges With the Same Name	122
Linear Constraints Can Be Evaluated As Nonlinear	122
 Appendix B. Accessibility	 123
Introduction	123
Accessibility Notes	123
Accessibility of Code Examples in Documentation	124
Accessibility of Links to External Web Sites in Documentation	124
Enabling Accessibility for Crystal Ball	124
Using the Tab and Arrow Keys in the Crystal Ball Decision Optimizer User Interface	124
TTY Access to Oracle Support Services	124
OptQuest Wizard Keyboard Command Equivalents	124
OptQuest Results Window Menus	125
OptQuest Control Panel Keyboard Shortcuts	126

Appendix C. References and Bibliography	129
Introduction	129
References	129
Spreadsheet Design	130
Optimization Topics	130
Metaheuristics	130
Stochastic (Probabilistic) Optimization Theory	130
Multiobjective Optimization	130
Financial Applications	130
Quality and Six Sigma Applications	131
Petrochemical Engineering Applications	131
Inventory System Applications	131
Glossary	133
Index	137

1

Welcome

In This Chapter

Introduction	9
How This Manual is Organized	9
Screen Capture Notes	10
Getting Help	10
Additional Resources	11

Introduction

Welcome to OptQuest, an optimization option available in Oracle Crystal Ball Decision Optimizer, Fusion Edition.

OptQuest enhances Crystal Ball by automatically searching for and finding optimal solutions to simulation models. Simulation models by themselves can only give you a range of possible outcomes for any situation. They do not tell you how to control the situation to achieve the best outcome

Using advanced optimization techniques, OptQuest finds the right combination of variables to produce accurate results. Suppose you use simulation models to answer questions such as “What are likely sales for next month?” Now, you can find the price points that maximize monthly sales. Suppose you ask, “What will production rates be for this new oil field?” Now, you can also determine the number of wells to drill to maximize net present value. Suppose you wonder, “Which stock portfolio should I pick?” With OptQuest, you can choose the one that yields the greatest profit with limited risk.

Like Crystal Ball, OptQuest is easy to learn and easy to use. With its wizard-based design, you can start optimizing your own models in under an hour. All you need to know is how to use a Crystal Ball spreadsheet model. From there, this manual guides you step by step, explaining OptQuest terms, procedures, and results.

How This Manual is Organized

Besides this Welcome chapter, the *OptQuest User Manual* includes the following additional chapters and appendices:

- [Chapter 2, “Overview”](#)

This chapter contains a description of optimization models and their components.

- [Chapter 3, “Setting Up and Optimizing a Model”](#)

This chapter provides step-by-step instructions for setting up and running an optimization in OptQuest.
- [Chapter 4, “OptQuest Tutorials”](#)

This chapter contains two tutorials designed to give you a quick overview of OptQuest’s features and to show you how to use the program. Read this chapter if you need a basic understanding of OptQuest.
- [Chapter 5, “Examples Using OptQuest”](#)

This chapter contains a variety of examples to show the types of problems that OptQuest can solve.
- [Appendix A, “Optimization Tips and Notes”](#)

This appendix describes different factors that enhance the performance of OptQuest’s features.
- [Appendix B, “Accessibility”](#)

The appendix provides a summary of OptQuest’s menus and a list of the commands you can execute directly from the keyboard.
- [Appendix C, “References and Bibliography”](#)

The appendix lists references describing OptQuest’s methodology, theory of operation, and comparisons to other optimization software packages. This appendix is designed for the advanced user.
- [Glossary](#)

This section is a compilation of terms specific to OptQuest as well as statistical terms used in this manual.

Screen Capture Notes

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

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

Getting Help

As you work in OptQuest, you can display online help in a variety of ways:

- Click the Help button in a dialog, .
- Press F1 in a dialog.

Note: In Excel 2007, click Help at the right end of the Crystal Ball ribbon. Note that if you press F1 in Excel 2007, Excel help is displayed 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.

Additional Resources

Oracle offers technical support, training, and additional resources to increase the effectiveness with which you can use Crystal Ball products.

For more information about all of these resources, see the Crystal Ball Web site at:

<http://www.oracle.com/crystalball>

2

Overview

In This Chapter

Introduction.....	13
What OptQuest Does	13
How OptQuest Works	14
About Optimization Models	15
Optimization Objectives	16
Decision Variables	17
Constraints.....	18
Model and Solution Feasibility	19
Efficient Frontier Analysis.....	19
OptQuest and Process Capability	21

Introduction

This chapter describes the three major elements of an optimization model: the objective, decision variables, and optional constraints. It also describes other elements required for models with uncertainty, such as forecast statistics and requirements, and ends with discussions of feasibility, Efficient Frontier analysis, and using optimization with Crystal Ball’s process capability features.

What OptQuest Does

Most simulation models have variables that you can control, such as how much to charge for rent or how much to invest. In Crystal Ball, these controlled variables are called decision variables. Finding the optimal values for decision variables can make the difference between reaching an important goal and missing that goal.

Obtaining optimal values generally requires that you search in an iterative or ad hoc fashion. A more rigorous method systematically enumerates all possible alternatives. This process can be very tedious and time consuming even for small models, and it is often not clear how to adjust the values from one simulation to the next.

OptQuest overcomes the limitations of both the ad hoc and the enumerative methods by intelligently searching for optimal solutions to your simulation models. You describe an optimization problem in OptQuest and then let it search for the values of decision variables that maximize or minimize a predefined objective. In almost all cases, OptQuest will efficiently find

an optimal or near-optimal solution among large sets of possible alternatives, even when exploring only a small fraction of them.

The easiest way to understand what OptQuest does is to apply it to a simple example. “[Tutorial 1 — Futura Apartments Model](#)” on page 49 demonstrates basic OptQuest operation.

How OptQuest Works

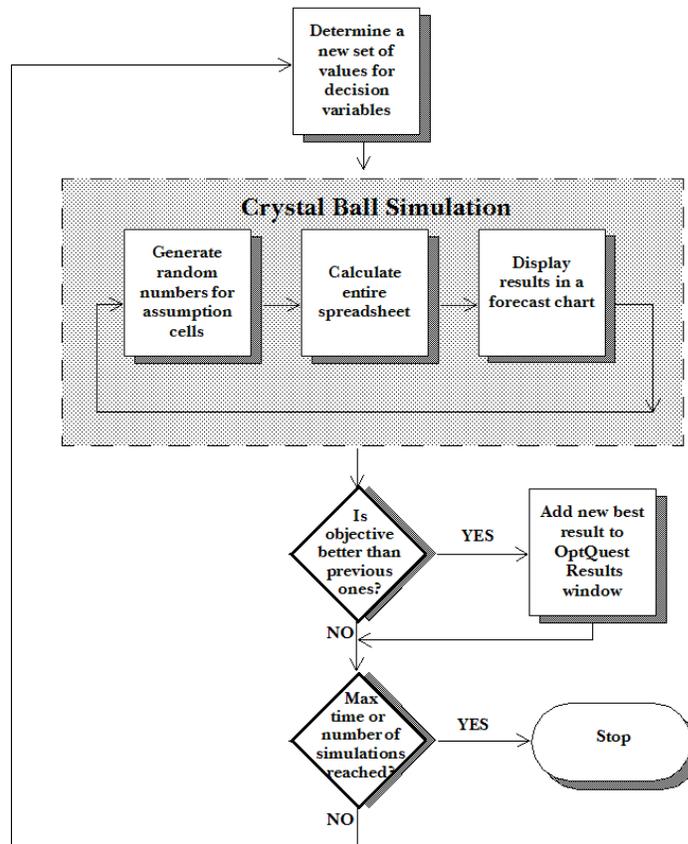
Traditional search methods work well when finding local solutions around a given starting point with model data that are precisely known. These methods fail, however, when searching for global solutions to real world problems that contain significant amounts of uncertainty. Recent developments in optimization have produced efficient search methods capable of finding optimal solutions to complex problems involving elements of uncertainty.

OptQuest incorporates metaheuristics to guide its search algorithm toward better solutions. This approach uses a form of adaptive memory to remember which solutions worked well before and recombines them into new, better solutions. Since this technique doesn’t use the hill-climbing approach of ordinary solvers, it does not get trapped in local solutions, and it does not get thrown off course by noisy (uncertain) model data. You can find more information on OptQuest’s search methodology in the references listed in [Appendix C, “References and Bibliography.”](#)

Once you describe an optimization problem (by selecting decision variables and the objective and possibly imposing constraints and requirements), OptQuest invokes Crystal Ball to evaluate the simulation model for different sets of decision variable values. OptQuest evaluates the statistical outputs from the simulation model, analyzes and integrates them with outputs from previous simulation runs, and determines a new set of values to evaluate. This is an iterative process that successively generates new sets of values. Not all of these values improve the objective, but over time this process provides a highly efficient trajectory to the best solutions.

As shown in the following flow chart, the search process continues until OptQuest reaches some termination criteria, either a limit on the amount of time devoted to the search or a maximum number of simulations.

Figure 1 OptQuest flow



About Optimization Models

In today's competitive global economy, people are faced with many difficult decisions. Such decisions might involve thousands or millions of potential alternatives. A model can provide valuable assistance in analyzing decisions and finding good solutions. Models capture the most important features of a problem and present them in a form that is easy to interpret. Models often provide insights that intuition alone cannot.

An OptQuest optimization model has four major elements: an objective, optional requirements, Crystal Ball decision variables, and optional constraints.

- **Optimization Objectives**—Elements that represents the target goal of the optimization, such as maximizing profit or minimizing cost, based on a forecast and related decision variables.
- **Requirements**—Optional restrictions placed on forecast statistics. All requirements must be satisfied before a solution can be considered feasible.
- **Decision Variables**—Variables over which you have control; for example, the amount of product to make, the number of dollars to allocate among different investments, or which projects to select from among a limited set.
- **Constraints**—Optional restrictions placed on decision variable values. For example, a constraint might ensure that the total amount of money allocated among various

investments cannot exceed a specified amount, or at most one project from a certain group can be selected.

For direct experience in setting up a model and running an optimization, see [“Tutorial 2 — Portfolio Allocation Model”](#) on page 54.

Optimization Objectives

Each optimization model has one objective that mathematically represents the model’s goal as a function of the assumption and decision variable cells, as well as other formulas in the model. OptQuest’s job is to find the optimal value of the objective by selecting and improving different values for the decision variables.

When model data are uncertain and can only be described using probability distributions, the objective itself will have some probability distribution for any set of decision variables. You can find this probability distribution by defining the objective as a forecast and using Crystal Ball to simulate the model.

Forecast Statistics

You cannot use an entire forecast distribution as the objective, but must characterize the distribution using a single summary measure for comparing and choosing one distribution over another. So, to use OptQuest, you must select a statistic of one forecast to be the objective. You must also select whether to maximize or minimize the objective, or set it to a target value.

The statistic you choose depends on your goals for the objective. For maximizing or minimizing some quantity, the mean or median are often used as measures of central tendency, with the mean being the more common of the two. For highly skewed distributions, however, the mean might become the less stable (having a higher standard error) of the two, and so the median becomes a better measure of central tendency.

For minimizing overall risk, the standard deviation and the variance of the objective are the two best statistics to use. For maximizing or minimizing the extreme values of the objective, a low or high percentile might be the appropriate statistic. For controlling the shape or range of the objective, the skewness, kurtosis, or certainty statistics might be used. If you are working with Six Sigma or another process quality program, you might want to use process capability metrics in defining the objective. For more information on these statistics, see the Glossary, online help, and the online *Oracle Crystal Ball Statistical Guide*.

Minimizing or Maximizing

Whether you want to maximize or minimize the objective depends on which statistic you select to optimize. For example, if your forecast is profit and you select the mean as the statistic, you would want to maximize the profit mean. However, if you select the standard deviation as the statistic, you might want to minimize it to limit the uncertainty of the forecast.

Requirements

Requirements restrict forecast statistics. These differ from constraints, since constraints restrict decision variables (or relationships among decision variables). Requirements are sometimes called "probabilistic constraints," "chance constraints," "side constraints," or "goals" in other literature.

When you define a requirement, you first select a forecast (either the objective forecast or another forecast). As with the objective, you then select a statistic for that forecast, but instead of maximizing or minimizing it, you give it an upper bound, a lower bound, or both (a range).

If you want to perform an Efficient Frontier analysis, you can define requirements with variable bounds. For more information, see [“Efficient Frontier Analysis” on page 19](#).

Requirement Examples

In the Portfolio Allocation example of [Chapter 4, “OptQuest Tutorials,”](#) the investor wants to impose a condition that limits the standard deviation of the total return. Because the standard deviation is a forecast statistic and not a decision variable, this restriction is a requirement.

The following are some examples of requirements on forecast statistics that you could specify:

```
95th percentile >= 1000
```

```
-1 <= skewness <= 1
```

```
Range 1000 to 2000 >= 50% certainty
```

Decision Variables

Decision variables are variables in your model that you can control, such as how much rent to charge or how much money to invest in a mutual fund. Decision variables aren't required for Crystal Ball models, but are required for OptQuest models. You define decision variables in Crystal Ball using Define, Define Decision or by clicking the Define Decision button in the toolbar or Excel 2007 ribbon.

When you define a decision variable in Crystal Ball, you define its:

- **Bounds**—Defines the upper and lower limits for the variable. OptQuest searches for solutions for the decision variable only within these limits.
- **Type**—Defines whether the variable type is discrete, continuous, binary, category, or custom:
 - **Continuous** — A variable that can be fractional (that is, it is not required to be an integer and can take on any value between its lower and upper bounds; no step size is required and any given range contains an infinite number of possible values.
 - **Discrete** — A variable that can only assume values equal to its lower bound plus a multiple of its step size; a step size is any number greater than zero but less than the variable's range.

- Binary — A decision variable that can be 0 or 1 to represent a yes-no decision, where 0 = no and 1 = yes.
- Category — A decision variable for representing attributes and indexes; can assume any discrete integer between the lower and upper bounds (inclusive), where the order (or direction) of the values does not matter (nominal). The bounds must be integers.
- Custom — A decision variable that can assume any value from a list of specific values (two values or more). You can enter a list of values or a cell reference to a list of values in the spreadsheet. If a cell 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. If you enter values in a list, they should be separated by a valid list separator -- a comma, semicolon, or other value specified in the Windows regional and language settings.

For details, refer to the *Oracle Crystal Ball User's Guide*.

- **Step Size**—Defines the difference between successive values of a discrete decision variable in the defined range. For example, a discrete decision variable with a range of 1 to 5 and a step size of 1 can only take on the values 1, 2, 3, 4, or 5; a discrete decision variable with a range of 0 to 2 with a step size of 0.25 can only take on the values 0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 1.75, and 2.0.

The cell value becomes the base case value, or starting value for the optimization.

Note: If changing the type of a decision variable causes the base case to fall outside the range of values that are valid for that type, a new base case value is selected. The base case changes to the nearest acceptable value for the new type.

In an optimization model, you select which decision variables to optimize from a list of all the defined decision variables. The values of the decision variables you select will change with each simulation until the best value for each decision variable is found within the available time or simulation limit.

Constraints

Constraints are optional settings in an optimization model. They restrict the decision variables by defining relationships among them. For example, if the total amount of money invested in two mutual funds must be \$50,000, you can define this as:

```
mutual fund #1 + mutual fund #2 = 50000
```

OptQuest only considers combinations of values for the two mutual funds whose sum is \$50,000.

Or if your budget restricts your spending on gasoline and fleet service to \$2,500, you can define this as:

```
gasoline + service <= 2500
```

In this case, OptQuest considers only combinations of values for gasoline and service at or below \$2,500.

Not all optimization models need constraints.

Model and Solution Feasibility

A feasible solution is one that satisfies all defined constraints and requirements. A solution is infeasible when no combination of decision variable values can satisfy the entire set of requirements and constraints. Note that a solution (i.e., a single set of values for the decision variables) can be infeasible by failing to satisfy the problem requirements or constraints, but this doesn't imply that the problem or model itself is infeasible.

However, constraints and requirements can be defined in such a way that the entire model is infeasible. For example, suppose that in the Portfolio Allocation problem in Chapter 1, the investor insists on finding an optimal investment portfolio with the following constraints:

```
Income fund + Aggressive growth fund <= 10000
```

```
Income fund + Aggressive growth fund >= 12000
```

Clearly, there is no combination of investments that will make the sum of the income fund and aggressive growth fund no more than \$10,000 and at the same time greater than or equal to \$12,000.

Or, for this same example, suppose the bounds for a decision variable were:

```
$15,000 <= Income fund <= $25,000
```

And a constraint was:

```
Income fund <= 5000
```

This also results in an infeasible problem.

You can make infeasible problems feasible by fixing the inconsistencies of the relationships modeled by the constraints. OptQuest detects optimization models that are constraint-infeasible and reports them to you.

If a model is constraint-feasible, OptQuest will always find a feasible solution and search for the optimal solution (that is, the best solution that satisfies all constraints).

When an optimization model includes requirements, a solution that is constraint-feasible might be infeasible with respect to one or more requirements.

After first satisfying constraint feasibility, OptQuest assumes that the user's next highest priority is to find a solution that is requirement-feasible. Therefore, it concentrates on finding a requirement-feasible solution and then on improving this solution, driven by the objective in the model.

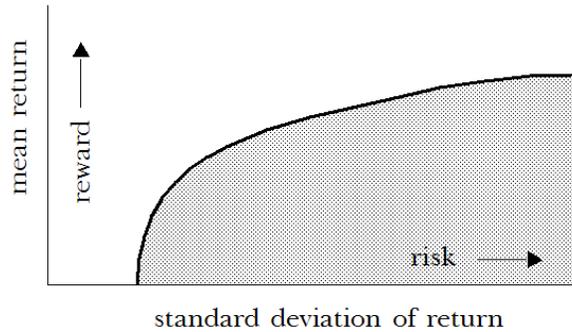
Efficient Frontier Analysis

Efficient Frontier analysis calculates the curve that plots an objective value against changes to a requirement or constraint. A typical use is for comparing portfolio returns against different risk levels so that investors can maximize return and minimize risk. If you want to use this type of analysis, you need to define a range of values for a requirement or constraint bound. For instructions and more information, see [“Setting Up Efficient Frontier Analysis in OptQuest” on page 46](#).

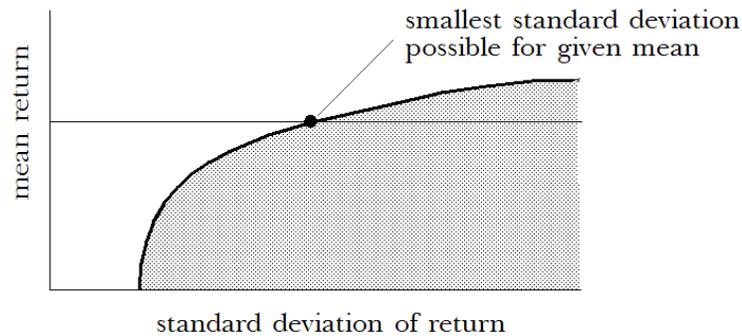
One use for Efficient Frontier analysis is to allocate funds among a portfolio of investments in the most efficient way. The Description page of Portfolio Revisited.xls describes this technique. “Efficient Portfolios” on page 20, following, offers the concepts behind it.

Efficient Portfolios

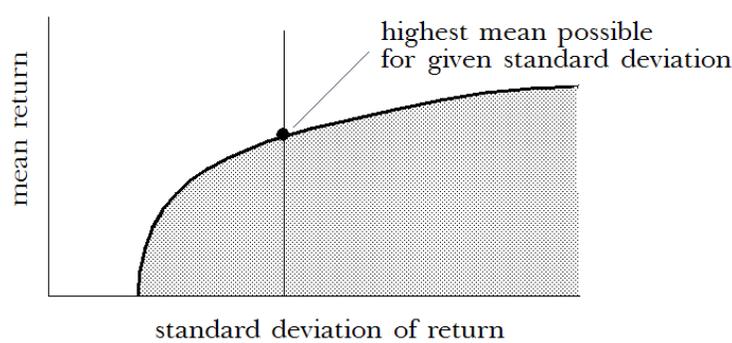
If you were to examine all the possible combinations of investment strategies for the assets described for Portfolio Revisited.xls, you would notice that each portfolio had a specific mean return and standard deviation of return associated with it. Plotting the means on one axis and the standard deviations on another axis, you can create a graph like this:



Points on or under the curve represent possible combinations of investments. Points above the curve are unobtainable combinations given the particular set of assets available. For any given mean return, there is one portfolio that has the smallest standard deviation possible. This portfolio lies on the curve at the point that intersects the mean of return.



Similarly, for any given standard deviation of return, there is one portfolio that has the highest mean return obtainable. This portfolio lies on the curve at the point that intersects the standard deviation of return.



Portfolios that lie directly on the curve are called efficient (see Markowitz, 1991 listed in [“Financial Applications” on page 130](#)), since it is impossible to obtain higher mean returns without generating higher standard deviations, or lower standard deviations without generating lower mean returns. The curve of efficient portfolios is often called the efficient frontier.

Portfolios that lie below the curve are called inefficient, meaning better portfolios exist with either higher returns, lower standard deviations, or both.

The example in [“Tutorial 2 — Portfolio Allocation Model” on page 54](#) uses one technique to search for optimal solutions on the efficient frontier. This method uses the mean and standard deviation of returns as the criteria for balancing risk and reward.

You can also use other criteria for selecting portfolios. Instead of using the mean return, you could select the median or mode as the measure of central tendency. These selection criteria would be called median-standard deviation efficient or mode-standard deviation efficient. Instead of using the standard deviation of return, you could select the variance, range minimum, or low-end percentile as the measure of risk or uncertainty. These selection criteria would be mean-variance efficient, mean-range minimum efficient, or mean-percentile efficient.

The mode is usually only available for discrete-valued forecast distributions where distinct values might occur more than once during the simulation.

OptQuest and Process Capability

You can use OptQuest to support process capability programs such as Six Sigma, Design for Six Sigma (DFSS), Lean principles, and similar quality initiatives. To do this, activate the Crystal Ball process capability features by checking Calculate Capability Metrics on the Statistics tab of the Run Preferences dialog. Once you do that, define a lower specification limit (LSL), upper specification limit (USL), or both for a forecast in the Define Forecast dialog. (You can also define an optional value target.)

Once you have defined at least one of the specification limits, you can optimize capability metrics for that forecast. The process capability metrics appear with other forecast statistics in the OptQuest Objectives panel. When you copy the values back to the model, the optimized values, relevant forecast charts, and the capability metrics table appear with the workbook. See the *Oracle Crystal Ball User's Guide* for more information.

3

Setting Up and Optimizing a Model

In This Chapter

Introduction.....	23
Overview	23
Developing a Crystal Ball Optimization Model	24
Starting OptQuest.....	26
Selecting the Forecast Objective	26
Selecting Decision Variables to Optimize.....	27
Specifying Constraints	28
Setting Options.....	34
Running Optimizations.....	35
Interpreting the Results.....	40
Saving optimization models and settings.....	45
Closing OptQuest.....	45
Setting Up Efficient Frontier Analysis in OptQuest	46
Transferring Settings from .opt Files.....	46
Learning More About OptQuest.....	48

Introduction

This chapter describes how to use OptQuest, step by step. It also gives details about each of the panels and dialogs in OptQuest, including all the fields and options.

Overview

- To set up and optimize a model with OptQuest, follow these steps:
 - 1 Create a Crystal Ball model of the problem.
 - 2 Define the decision variables within Crystal Ball.
 - 3 In OptQuest, select the forecast objective and define any requirements.
 - 4 Select decision variables to optimize.
 - 5 Specify any constraints on the decision variables.
 - 6 Select optimization settings.

7 Run the optimization.

8 Interpret the results.

For Users of OptQuest Versions Earlier Than 11.1.1.x

If you used a version of OptQuest earlier than 11.1.1.x, be aware of some significant changes. As you have discovered, the user interface is redesigned to be easier to use. For added flexibility, there are now five types of decision variables.

Another difference is that .opt files are no longer used to store optimization settings. For more information on saving optimization settings and options, see [“Saving optimization models and settings” on page 45](#). An .opt file viewer is provided to help you transfer settings from .opt files to current model workbooks. For instructions, see [“Transferring Settings from .opt Files” on page 46](#).

Developing a Crystal Ball Optimization Model

Before using OptQuest, you must first develop a useful Crystal Ball model. This involves building a well-tested spreadsheet in Excel, and then defining assumptions and forecast cells using Crystal Ball. You should refine the Crystal Ball model and run several simulations to ensure that the model is working correctly and that the results are what you expect.

Developing the Worksheet

You should build your spreadsheet model using principles of good design, since this makes understanding and modifying it easier.

The spreadsheet should include:

- A descriptive title.
- An input data area separate from the output and any working space. Place all input variables in their own cells where you can later define them as assumptions or decision variables.
- A working space for all complex calculations, formulas, and data tables.
- A separate output section that provides the model results.

Examine the Portfolio Allocation spreadsheet model ([Figure 20](#)) for an example.

Note that all assumptions are in rows 5 through 8. Rows 13 through 16 are reserved for decision variables, created by users during the OptQuest tutorials. Forecast cells reference these input variable cells in their calculations, not values directly. Therefore, you could easily change any values, and the forecast calculations would be automatically updated.

Other tips that improve the usefulness of your spreadsheet are:

- Reference input data only with cell references or range names so that any changes are automatically reflected throughout the worksheet.
- Use formats, such as currency or comma formats, appropriately.

- Divide complex calculations into several cells to minimize the chance for error and enhance understanding.
- Place comments next to formula cells for explanation, if needed.
- Consult a reference such as those listed in [Appendix C](#) for further discussion of good spreadsheet design.

Defining Assumptions, Decision Variables, and Forecasts

Once you build and test the spreadsheet, you can define your assumptions, decision variables, and forecasts. For more information on defining assumptions, decision variables, and forecasts, see the *Oracle Crystal Ball User's Guide*.

Setting Crystal Ball Run Preferences

To set Crystal Ball run preferences, select Run, Run Preferences. For optimization purposes, you should usually use the following Crystal Ball settings:

- Trials tab — Maximum number of trials to run set to 1000.
Central-tendency statistics such as mean, median, and mode usually stabilize sufficiently at 500 to 1000 trials per simulation. Tail-end percentiles and maximum and minimum range values generally require at least 2000 trials.
- Sampling tab — Sampling method set to Latin Hypercube.
Latin Hypercube sampling increases the quality of the solutions, especially the accuracy of the mean statistic.
- Sampling tab — Random Number Generation set to Use Same Sequence Of Random Numbers with an Initial Seed Value of 999.
The initial seed value determines the first number in the sequence of random numbers generated for the assumption cells. Then, you can repeat simulations using the same set of random numbers to accurately compare the simulation results. If you do not set an initial seed value, OptQuest will automatically pick a random seed and use that starting seed for each simulation that is run.
When your Crystal Ball forecast has extreme outliers, run the optimization with several different seed values to test the solution's stability.
- Speed tab — Run in Extreme Speed if possible.
After you define the assumptions, decision variables, and forecasts in Crystal Ball, you can begin the optimization process in OptQuest.

Starting OptQuest

► To start OptQuest:

1 Choose Run, OptQuest.

The OptQuest wizard starts.

2 Set up the optimization by completing each wizard panel. The first step of this process is selecting a forecast objective to optimize.

Note: This version of OptQuest does not use .opt files. If you would like to retrieve settings from existing .opt files for use in this version of OptQuest, see [“Transferring Settings from .opt Files” on page 46.](#)

Selecting the Forecast Objective

When the OptQuest wizard starts, the Objectives panel opens, similar to [Figure 12](#). (The first time you start the wizard, the Welcome screen opens. Click Next to display the Objectives panel.)

In the Objectives panel, you choose a forecast statistic to maximize, minimize, or set to a target value. Optionally, you can define one or more requirements either on the objective forecast or on other forecasts.

[Figure 21](#) shows a default objective including the first forecast found in the model. If you have questions about settings, click the Help button.

Note: Each optimization must include one and only one objective. However, you can define several objectives and exclude those not used in the current optimization. To exclude an objective, check the box in the Exclude column for that objective.

► To define a forecast objective and, optionally, define requirements:

1 Click Add Objective.

A default objective is displayed in the Objectives area.

2 Review the default objective definition. It has the format Operation > Statistic > Forecast.

- a. First, if the model has more than one forecast, does the default objective include the same forecast you want to include in the objective? If not, click the underlined forecast and replace it with your selection. If more than ten forecasts are available, **More Forecasts** is displayed at the bottom of the list. You can select it to display a forecast selection dialog.
- b. Next, do you want to maximize a statistic for that forecast? If you would prefer to minimize the statistic or set it to a target value, click the underlined operation and choose an alternative.
- c. Finally, is the underlined statistic the one you prefer to use. If not, click it and choose a different one. If you have activated Crystal Ball’s process capability features and have defined an LSL or USL, the process capability statistics are available in the list of statistics.

Note: For many problems, the mean (expected value) of the forecast is the most appropriate statistic to optimize, but it need not always be. For example, investors who want to maximize the upside potential of their portfolios might want to use the 90th or 95th percentile as the objective. The results would be solutions that have the highest likelihood of achieving the largest possible returns. Similarly, to minimize the downside potential of the portfolio, they might use the 5th or 10th percentile as the objective to minimize the possibility of large losses. You can use other statistics to realize different objectives. See the Glossary, online help, and the online *Oracle Crystal Ball Statistical Guide* for a description of all available statistics.

3 Optionally, define requirements.

- a. To add a requirement, click Add Requirement. A default requirement is displayed.
- b. First, look at the default statistic. Is it the one you want to use? To review the list of available choices, click the underlined statistic and select a different one if you want. Depending on your choice, the requirement statement could change.
- c. Next, review the forecast. If you want, click the underlined forecast and choose another.
- d. Then, review the requirement operator. The selected statistic can be less than or equal to a selected value, greater than or equal to a selected value, or between two selected values (including the values). Click the underlined limit to choose another. If you choose Between, an additional target value is displayed.
- e. Finally, review and adjust the target value or values. To change a value, click it and then type a new number over it.
- f. You can repeat steps 3a through 3e to add additional requirements. New requirements are duplicates of the last one entered.
- g. To delete a requirement, click it and then click Delete.
- h. If you want to set variable bounds for Efficient Frontier analysis, select a variable and click Efficient Frontier. For details, see [“Efficient Frontier Analysis” on page 19](#).

4 When objective and requirement settings are complete, click Next.

The Decision Variables panel opens.

Note: You can create multiple requirements without using all of them at once. If you check the Exclude box, that requirement is not used in the current OptQuest simulation.

Selecting Decision Variables to Optimize

When you click Next in the Objectives panel, the Decision Variables panel opens, similar to [Figure 22](#). It lists every decision variable, frozen or not, defined in all open Excel workbooks.

The next step of the optimization process is selecting decision variables to optimize. The value of each decision variable changes with each simulation until OptQuest finds values that yield the best objective. For some analyses, you might fix the values of certain decision variables and optimize the rest.

By default, all decision variables in all open workbooks are shown, even those that are frozen in your model. Frozen decision variables have a check in the Freeze column. If you want, you can uncheck them and include them in the optimization. Be aware, though, that if you freeze or unfreeze a decision variable, you are also changing it in your model.

OptQuest uses the limits, base case (start value), and decision variable type you entered when you defined the decision variables.

If you check Show Cell Locations, the following additional columns appear in the Decision Variables panel: Cell Address, Worksheet, and Workbook.

► To confirm and change selections:

- 1 Review the listed variables. Check Freeze for any that you do not want to include in the OptQuest optimization.
- 2 Optionally, change the lower and upper bounds, base case, or decision variable type for any listed decision variable. Highlight the existing value and type over it. This changes the decision variable definition in your worksheet.

Note the following about these settings:

- The tighter the bounds you specify, the fewer values OptQuest must search to find the optimal solution. However, this efficiency comes at the expense of missing the optimal solution if it lies outside the specified bounds.
- By default, OptQuest uses the base case cell values in your Crystal Ball model as the suggested starting solution. If the suggested values lie outside of the specified bounds or do not meet the problem constraints, OptQuest ignores them.

Note: You can sort decision variables in the Decision Variables panel by name, type, freeze status, cell address, worksheet, or workbook. To sort, click the column heading. An arrow is displayed at the right end of the column to show the direction of the sort. The sort column and direction of the decision variables is stored as a global preference and is also used to set the order of the decision variables in the reports and extracted data.

- 3 When your decision variable selections are complete, click Next.

The Constraints panel opens.

Specifying Constraints

In OptQuest, constraints limit the possible solutions to a model in terms of relationships among the decision variables. You can use the Constraints panel to specify linear and nonlinear constraints. For example, in [“Tutorial 2 — Portfolio Allocation Model”](#) on page 54, the total investment was limited to \$100,000. In the Constraints panel, this limit is expressed by the formula:

```
Money Market fund + Income fund + Growth and Income fund + Aggressive  
Growth fund = 100000
```

By default, the Constraints panel opens in Simple Entry mode. In this mode, most of the constraint formula is entered into cells in your spreadsheet. You then complete the constraint formula on the Constraints panel using a simple conditional expression like `Sheet1!A1 <= 100`.

For more information, see the following section, [“Specifying Constraints in Simple Entry Mode” on page 29](#).

If you move to Advanced Entry mode, you can enter constraint formulas directly. See [“Specifying Constraints in Advanced Entry Mode” on page 29](#).

Specifying Constraints in Simple Entry Mode

When you click Next in the Decision Variables panel or click Constraints in the navigation list, the Constraints panel opens, similar to [Figure 23](#).

By default, the Constraints panel opens in Simple Entry mode. If you click Add Constraint, you can reference cells with formulas for the left and right sides of the constraint formula and you can choose an operator. Alternatively, you can enter a value for the right side or left side. For information about allowable constraint formulas, see [“Constraint Rules and Syntax” on page 32](#).

For an example of using Simple Entry mode, see [“Specifying Constraints” on page 59](#).

Specifying Constraints in Advanced Entry Mode

► To use the Constraints panel in Advanced Entry mode:

- 1 **Switch to Advanced Entry mode by checking Advanced Entry in the upper right corner of the Constraints editor.**
- 2 **In the Constraints editor, enter a mathematical formula. You can use the buttons at the bottom of the Constraints panel to help you edit the formula.**

For information on the Constraints editor syntax, see [“Constraint Rules and Syntax” on page 32](#).

You can also enter parts of a constraint formula into spreadsheet cells and then reference those cells, separated by an operator, in a formula. See [“Constraints and Cell References in Advanced Entry Mode” on page 33](#).

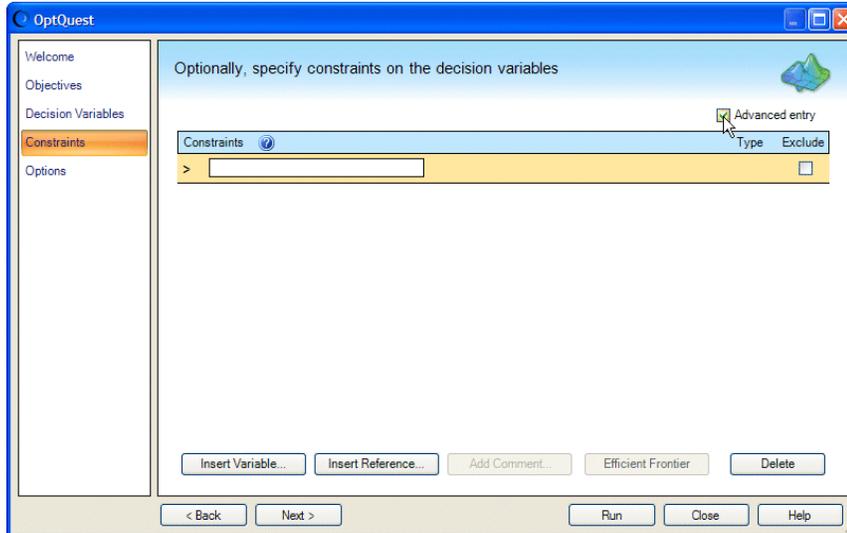
- 3 **Enter any additional constraints on their own lines.**
- 4 **When you are done, click Next to display the Options panel.**

Note: You can copy and paste constraints to duplicate them for further editing.

Advanced Entry Example

To enter Advanced Entry mode, check Advanced Entry in the upper right corner of the Constraints panel as shown in [Figure 2](#). A Constraints edit box opens.

Figure 2 Entering Advanced Entry mode in the Constraints panel



At first the Constraints edit box is blank. A series of buttons near the bottom of the dialog can help you create a formula in it. You can enter a linear or nonlinear formula and you can enter any number of formulas, as long as each constraint formula is on its own line. For details, see [“Constraints Editor and Related Buttons”](#) on page 31.

In this case, supposed you want to create a formula that adds all the decision variable values and specifies that they should be equal to \$100,000, as discussed in [“Tutorial 2 — Portfolio Allocation Model”](#) on page 54.

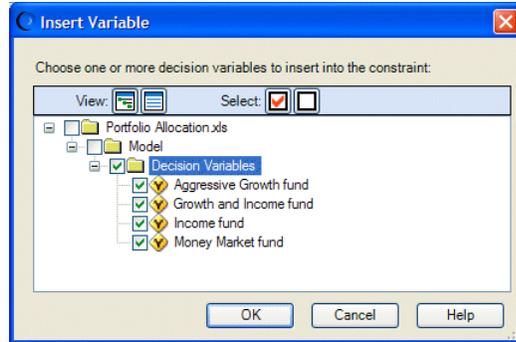
Constraints Editor Example

► To create this formula:

- 1 Click **Insert Variable**.

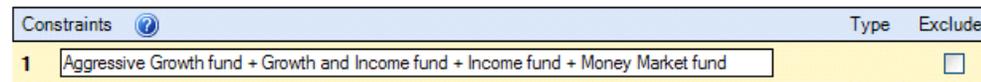
The Insert Variable dialog opens.

Figure 3 Insert Variable dialog, Portfolio Allocation model



- 2 Since you want to include all four decision variables in the constraint formula, put a check in front of each name. To check all four at once, check the box in front of Decision Variables. Then, click OK.

The variables appear in the edit box as a sum:



- 3 After **Money Market fund**, type an equals sign (=).
- 4 Enter the total investment as **\$100,000** (without the dollar sign or comma), so that the final constraint looks like:

Money market fund + Income fund + Growth and income fund + Aggressive growth fund = 100000

Note: Don't use "\$" or a comma in constraints. See [“Constraint Rules and Syntax”](#) on page 32 for other rules about constraint formulas.

- 5 Click **Next** to continue.

The Options panel opens, similar to [“Constraint Rules and Syntax”](#) on page 32.

Constraints Editor and Related Buttons

The upper part of the Constraints panel is the Constraints editor. The lower part of the Constraints panel contains buttons that perform the following tasks in Advanced Entry mode:

Button	Description
Insert Variable	Lists all available decision variables you can insert. If you choose more than one, they are automatically added to the Constraints editor with plus (+) signs between them.
Insert Reference	Displays the Cell Reference dialog, where you can either point to a cell or enter a formula to include in the constraint formula you are creating. For more information, see “Constraints and Cell References in Advanced Entry Mode” on page 33.
Add Comment	Displays the Add Comment dialog, where you can enter a comment that describes the constraint. The comment is displayed in the Constraints panel above the constraint. It also is displayed in the OptQuest Results window to identify the constraint.

Button	Description
Efficient Frontier	Adds a requirement with a variable upper or lower bound for use in Efficient Frontier analysis. For more information, see “Efficient Frontier Analysis” on page 19 . If you have already added a variable requirement on the Objectives panel, a message is displayed and asks if you want to use the new one instead.
Delete	Deletes the currently selected constraint.

To add a variable or a reference to a constraint, place your cursor where you want the variable and then either type the variable name or click the Insert Variable button and select one or more variables in the list. You can define any number of constraints.

Constraint Rules and Syntax

In general, constraint formulas are like standard Excel formulas. Each constraint formula:

- Is constructed of mathematical combinations of constants, selected decision variables, and other elements.
- Must each be on its own line.
- Can be linear or nonlinear. You can multiply a decision variable by a constant (linear), and you can multiply it by another decision variable (nonlinear).
- Cannot have commas, dollar signs, or other non-mathematical symbols.

In Advanced Entry mode, decision variables can be entered directly by name but in Simple Entry mode, they can only be referenced within spreadsheet formulas by cell location or range name.

The mathematical operations allowed in constraint formulas are:

Table 1 Mathematical operations in the Constraints panel OptQuest

Operation	Syntax	Example
Addition	Use + between terms	$\text{var1} + \text{var2} = 30$
Subtraction	Use - between terms	$\text{var1} - \text{var2} = 12$
Multiplication	Use * between terms	$4.2 * \text{var1} \geq 9$
Division	Use / between terms	$4.2 / \text{var1} \geq 18$
Equalities and inequalities	Use =, <=, or >= between left and right sides of the constraint. Notice that < and > are treated as <= and >= for constraints involving continuous decision variables.	$\text{var1} * \text{var2} \leq 5$
Exponents	Use ^ between a term and the exponential power	var1^3

Notice that the above examples are for Advanced Entry mode. In Simple Entry mode, the expression on the left side of the operator would be entered into a spreadsheet cell. The actual

formula in the Constraints panel would include a cell reference, the operator, and either a value or another cell reference. For an example, see [Figure 26](#).

Note: Although these examples always show a formula on the left side of the operator, you can actually have a formula (or a cell reference to a formula in the spreadsheet) on either the left or the right side.

You can also use Excel functions and range names in constraint formulas.

If you are using Advanced Entry mode, calculations occur according to the following precedence. Multiplication and division occur first, and then addition and subtraction. For example, $5 * E6 + 10 * F7 - 26 * G4$ means: Multiply 5 times the value in cell E6, add that product to the product of 10 times the value in cell F7, and then subtract the product of 26 times the value in cell G4 from the result. You can use parentheses to override precedence. If you are using Simple Entry mode, you are creating formulas in Excel and Excel's precedence rules apply.

Constraints and Cell References in Advanced Entry Mode

“[Specifying Constraints in Simple Entry Mode](#)” on [page 29](#) describes how you can create formulas in spreadsheet cells and then reference them when creating constraints. You can also use cell references in Advanced Entry mode to simplify constraint formulas.

- To do this in Advanced Entry mode:
 - 1 Enter a formula for the left side of the constraint into a spreadsheet cell. The example in “[Specifying Constraints in Simple Entry Mode](#)” on [page 29](#) has `=SUM(C13:C16)` entered into cell G13.
 - 2 Consider what to use for the right side of the formula. It can be a single value or a formula that resolves to a constant.
 - 3 Decide on the relationship between the left and right side: `=`, `<=`, `>=`.
 - 4 Run OptQuest and display the Constraints panel.
 - 5 With the cursor in a constraint formula edit box, click Insert Reference. Point to the cell with the left side of the formula and click OK.
 - 6 Following the cell reference, type the relationship operator.
 - 7 Click Insert Reference again and point to the cell for the right side of the formula. Click OK again. Alternately, you can type a numeric value instead of using a cell reference

You can add additional constraints or other OptQuest settings and run the optimization when settings are complete.

For best results, avoid putting an entire formula, including operator, in a cell and then referencing that cell in a constraint formula that tests whether the formula is true or false. For example, suppose cell G6 contains `=SUM(B2:E2) >= 10`. You should avoid defining a constraint as `G6 = TRUE`. This method does not provide OptQuest with the information it needs to improve the solution.

Instead, you should break up the left-hand and right-hand parts of the equation and make sure the conditional operator (=, >=, <=) is entered in the constraints panel. In this example, cell G6 could contain =SUM(B2:E2) and the constraint could be written G6 >= 10.

Constraint Types

Constraints can be linear, nonlinear, or constant (in special situations):

- Linear constraints are more efficient in generating feasible solutions to try. They are evaluated by OptQuest before a solution is generated.
- Nonlinear constraints are evaluated by Excel before a simulation is run. They may be slower to evaluate if they contain many Excel functions or refer to many formulas in the spreadsheet. They are less efficient at generating feasible solutions.
- Constant constraints are generally an error unless a user-defined macro or the Crystal Ball Auto Extract feature is used to set values in a referenced spreadsheet cell. For more about user-defined macros and constant constraints, see information about the OptQuest Developer Kit in the *Oracle Crystal Ball Developer's Guide*.

When you create a constraint, its type is displayed after the formula.

Setting Options

When you click Next in the Constraints panel or click Options in the navigation list, the Options panel opens, similar to [Figure 17](#).

You can use the Options panel to set OptQuest options including optimization length (time or number of simulations), Crystal Ball simulation preferences, optimization type (with or without simulation, window display, automatic decision variable value settings, and more.

Note: If you saved settings in a version of OptQuest earlier than 11.1.1, you will need to set new options in this version of OptQuest.

► To change the settings:

- 1 Choose the settings you want, typing any new numeric values.

Settings are as follows:

Table 2 OptQuest Options panel settings

Option	Description
Optimization Control	Settings that control how long the optimization runs. Choose "Run for ___ simulations" or "Run for ___ minutes" and enter the desired value. The defaults are 1000 simulations and 5 minutes. You can also click the Run Preferences button to change settings in the Crystal Ball Run Preferences dialog.

Option	Description
Type Of Optimization	Choose "With simulation (stochastic)" to run a simulation on the assumption variables or choose "Without simulation (deterministic)" to use the base case (cell value) for the assumption cells.
While Running	Settings that control chart window display. You can choose "Show chart windows as defined" for maximum information or "Show only target forecast window" for fastest performance. "Update only for new best solutions" is checked by default to enhance performance and will only show results related to the best solution. Uncheck this setting to see the forecast results for each solution.
Decision Variable Cells	Choose "Leave set to original values" to keep the original base case values in decision variable cells, the default. At the end of an optimization, you can copy any solution OptQuest tried (including the best solution) to these cells if you want. Choose "Automatically set to best solution" to update decision variable cells in the workbook to the best solution found at the end of the optimization.
Advanced Options	Click this button to display the Advanced Options dialog, where you can choose to stop a simulation early if the desired confidence level or number of non-improving solutions is met. For details, see "Advanced Options" on page 35 .

2 When options settings and all other required OptQuest settings are complete, click Run.

Advanced Options

The OptQuest advanced options control whether the optimization stops automatically under certain conditions.

- The first setting, Enable Low-Confidence Testing, stops the active optimization if the confidence interval around the forecast objective indicates that the current solution is inferior to the current best solution. This only works if the statistic used for the forecast objective is the mean, standard deviation, or a percentile.

This setting uses the Confidence Level setting on the Trials tab of the Crystal Ball Run Preferences dialog to determine the confidence interval.
- The second setting, Automatically Stop After ___ Non-improving Solutions, stops the active simulation if the specified number of solutions are calculated without generating a new best solution. The default setting is unchecked (off) with a value of 500.

Running Optimizations

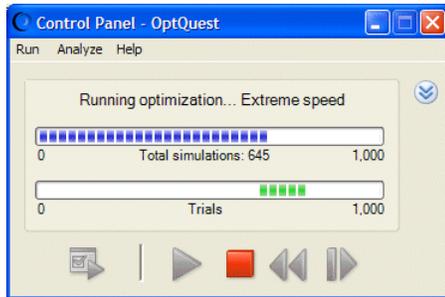
To run an optimization, click Run at the bottom of any OptQuest wizard panel. Once the optimization starts, you can use buttons in the Control Panel to stop, pause, continue, or restart at any time.

You cannot work in Crystal Ball or Excel or make changes in OptQuest when running an optimization, but you can work in other programs. Do not close Excel, Crystal Ball, or OptQuest while running an optimization.

OptQuest Control Panel Buttons and Commands

You can use the buttons and commands on the OptQuest Control Panel for starting and stopping an optimization (Figure 4).

Figure 4 OptQuest Control Panel



The Control Panel menus are the same as the Crystal Ball Run and Analyze menus. The Help menu describes the Control Panel. The following buttons are available:

Action	Button	Description
Run Preferences		Opens a dialog for controlling optimizations.
Start or Continue		Starts a new optimization or continues a paused optimization.
Pause or Stop		Pauses or stops the current optimization.
Reset		Resets the current optimization and closes all results.

The progress bars help you keep track of individual simulations and the optimization as a whole. If simulations are running faster than one per second, you will see a "marquee" style progress bar. If an optimization is set to run for a maximum amount of time, the upper progress bar shows elapsed time instead of number of stimulations. A notification message is displayed if the optimization stops early because a set confidence level is reached or there has been no solution improvement for a set number of simulations.

If you click the More button, , a panel opens below the control buttons with additional information about the optimization.

OptQuest Results Window

When an optimization is complete, you can view the OptQuest Results window for information about the current optimization. The following sections describe different Results window views:

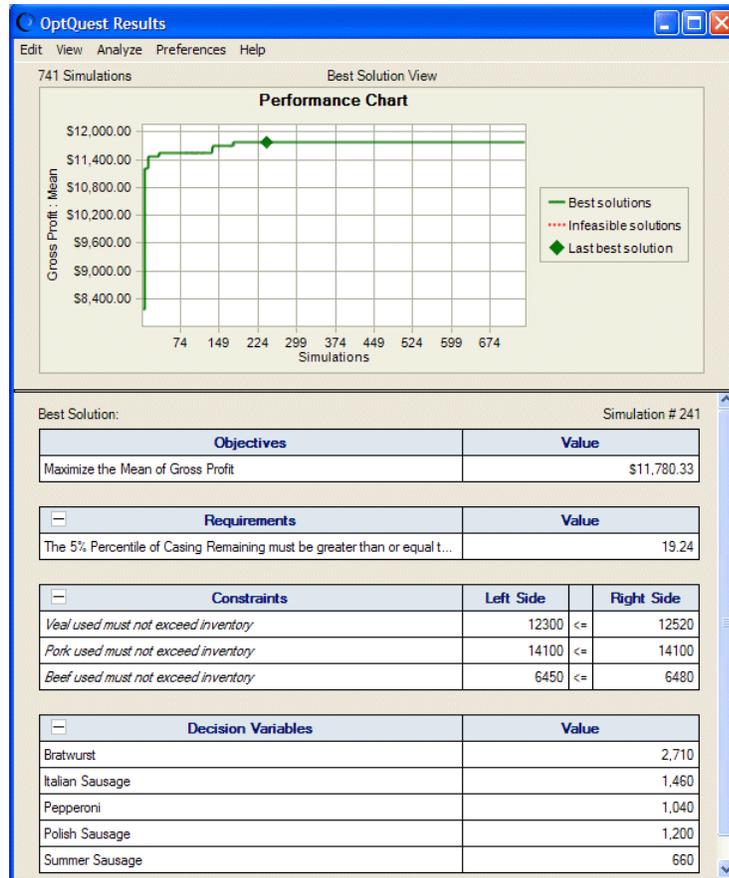
- [“Best Solution View” on page 37](#)

- “Solution Analysis View” on page 38
- “Efficient Frontier Chart” on page 39

Best Solution View

Figure 5 shows Best Solution view results for an OptQuest example model, Product Mix.xls.

Figure 5 OptQuest Results window, Best Solution view



In the Best Solution view, the OptQuest Results window shows a performance chart plotting best solutions found during analysis. It also shows the single best solution found for the objective, any requirements, any constraints, and all included decision variables.

Performance Chart

The performance chart displays the trajectory of the search; that is, the rate at which the best objective value has changed during the course of the search. This is shown as a plot of the best objective values as a function of the number of simulations (solutions). If any requirements have been specified, the line might initially be red, indicating that the corresponding solutions are not feasible according to the requirements. A green line indicates feasible solutions.

Once OptQuest finds a feasible solution, it is common for this line to show an exponential decay form (for minimization), where most improvements occur early in the search.

Best Solution Values

Each time OptQuest identifies a better solution (closer to feasibility or with a better objective) during the optimization, it plots new points in the performance chart and updates the tables below the chart.

If you have requested an Efficient Frontier analysis, you can also display the Efficient Frontier view. For more about this view, see [“Efficient Frontier Analysis” on page 19](#).

Menus

The OptQuest Results window has several menus you can use to copy results to your spreadsheet, copy charts, print results, view other charts, and more. For a list of menu commands and their shortcut keys, see [“OptQuest Results Window Menus” on page 125](#).

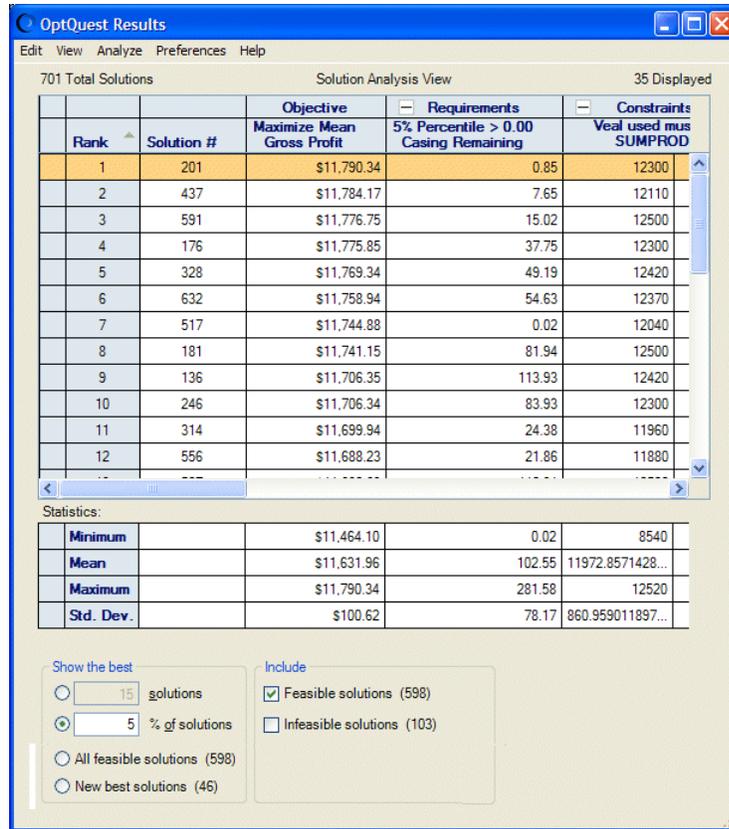
Solution Analysis View

In Solution Analysis view, the OptQuest Results window lists the best solutions found during the optimization. By default, the top 5% of solutions are sorted by the objective value. Controls at the bottom of the window indicate how many solutions to view. Statistics are calculated for the solutions shown.

Note: While OptQuest is running, Solution Analysis view shows the new best solutions, except for Efficient Frontier analyses. The top ten solutions still show when an Efficient Frontier analysis is running.

To display Solution Analysis view, choose View, Solution Analysis in the OptQuest Results window menubar.

Figure 6 Solution Analysis view



In the Show The Best group, indicate whether to show a specific number or percentage of the best solutions or all solutions. Your entry defines the analysis range. For example, if you want to examine the top 10% of all the solutions, check __ % of Solutions and enter 10 in the box.

You can choose whether to include feasible, infeasible, or all solutions. If you have requested an Efficient Frontier analysis, you can choose just the solutions for a particular efficient frontier test point.

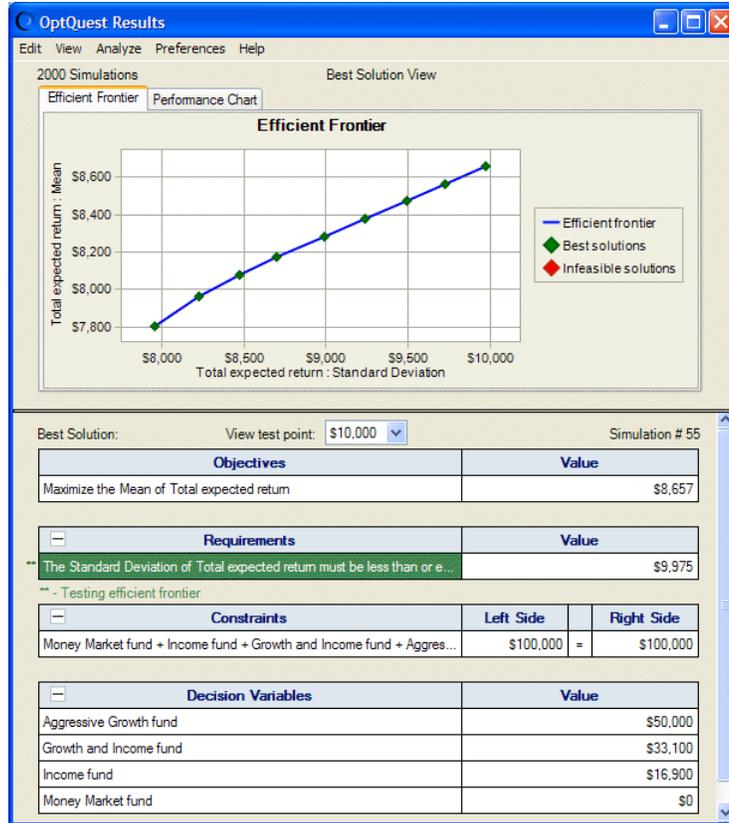
When you make your choices, statistics are calculated in the four rows at the bottom: the minimum, mean, maximum, and standard deviation values for all columns according to your display selections.

You can click the – or + at the left side of a column heading to condense sections and show more columns onscreen. You can also click in the right side of a heading to sort that heading. A small triangle is displayed. You can click it to sort the column in ascending or descending order.

Efficient Frontier Chart

If you have entered a variable requirement for the optimization, an Efficient Frontier tab opens with the Performance Chart tab in Best Solution view ([Figure 7](#)).

Figure 7 Efficient Frontier chart, Best Solution view



The Efficient Frontier window displays a plot of the objective value against the requirement or constraint that is being tested. The best solution for each test point is displayed as a green diamond in the chart. The table below the chart shows the best solution values for a specific test point. You can choose which best solution to view by selecting the View Test Point drop-down menu or by clicking the diamond symbol in the chart. For more information about Efficient Frontier analysis, see [“Setting Up Efficient Frontier Analysis in OptQuest”](#) on page 46.

Interpreting the Results

After solving an optimization problem with OptQuest, you can:

1. View a solution analysis to determine the robustness of the results.
2. Run a longer Crystal Ball simulation using the optimal values of the decision variables to more accurately assess the risks of the recommended solution.
3. Use Crystal Ball’s analysis features to further evaluate the optimal solution.

Viewing a Solution Analysis

► After the optimization is finished, interpret your optimization results:

- 1 Select View, Solution Analysis in the OptQuest Results window.

The Solution Analysis view opens with a partial listing of the solutions that OptQuest tried during the optimization. The solutions are shown row-wise in the upper grid with a smaller grid giving the statistics for each column.

Note that the OptQuest Results window has several menus you can use to copy results to your spreadsheet, copy charts, print results, view other charts, and more. For a list of menu commands and their shortcut keys, see [Appendix B, “Accessibility.”](#)

2 Choose which solutions to view.

Below both grids are groups of controls you can use to filter the solutions to view. All of the controls combine to filter the set of solutions. Some controls show the number of solutions that will be included in parenthesis.

- In the leftmost group, choose to view only the top number or percentage of best solutions (highest or lowest objective values), all of the solutions, or only the new best solutions (corresponding to up or down “jumps” in the performance chart).
- In the center group, choose whether to include feasible, infeasible, or both types of solutions.
- If you have requested an efficient frontier analysis, choose a test point from the dropdown menu in the rightmost group. Note that all of the solutions are considered for a particular test point, even if they were evaluated at a previous or later test point in the optimization.

Once you have chosen a set of solutions to analyze, you can click on the right-side of a column heading to sort the solution by that heading. The small triangle indicates the direction of the sort order. You can also click the + or – symbol at the left side of a group of columns to condense or expand the amount of information displayed.

Bounds Analysis

The Solution Analysis view is useful for determining how restrictive the bounds are for requirements or constraints, especially when there are multiple bounds involved. When viewing the best solutions for an optimization, if most of the values for a requirement or constraint are at or near a specific bound, this indicates that the requirement or constraint is having a significant effect on the values that are obtainable for the objective.

Sensitivity Analysis

The Solution Analysis view is useful for determining the sensitivity of decision variables with respect to the model’s objective. When viewing the best solutions for an optimization, compare the relative size of the ranges for each of your decision variables. Generally speaking, a decision variable with a smaller relative range indicates that it has a greater impact on the objective. This is because small changes in the decision variable can force the solutions to be less than optimal. Conversely, a decision variable with a wider relative range indicates that it has a lesser impact on the objective since different values do not seem to alter the set of best solutions.

These are general guidelines only. The results for your situation can be affected by the type and length of the optimization, the initial bounds defined for the decision variables, and other factors.

Running a Longer Simulation of the Results

- To more accurately assess the recommended solution, run a longer Crystal Ball simulation using the optimal values of the decision variables.
 - 1 If you did not choose to automatically copy OptQuest results to the model workbook (set in the Options panel), you can choose **Edit, Copy [Best] Solution to Spreadsheet** in the OptQuest Results window.
OptQuest copies the decision variables values from the selected solution into the Excel model.
 - 2 In Crystal Ball, reset the optimization, select **Run, Run Preferences**, and increase the maximum number of trials per simulation.
 - 3 Run the simulation.
 - 4 Use Crystal Ball analysis tools to analyze your results.

For more information on using these tools, see the *Oracle Crystal Ball User's Guide*.

Printing OptQuest Results

- To print results from any OptQuest results view:
 - 1 Run an OptQuest optimization and open the OptQuest Results window.
 - 2 Choose a view from the View menu in the OptQuest Results window menu bar.
 - 3 Choose Edit in the OptQuest Results window menu bar.
 - 4 Choose an appropriate command related to printing at the bottom of the Edit menu. Choices are **Page Setup**, **Print Preview**, and **Print**.

Viewing Charts in Crystal Ball

When an optimization completes, you can choose **Analyze, Forecast Charts** to view forecast charts and other charts based on the best solution results. However, if you copied a solution from the Solution Analysis view that is different from the best solution, you need to run a simulation in Crystal Ball before choosing a chart command from the Analyze menu. See the *Oracle Crystal Ball User's Guide* for further instructions.

Creating OptQuest Reports

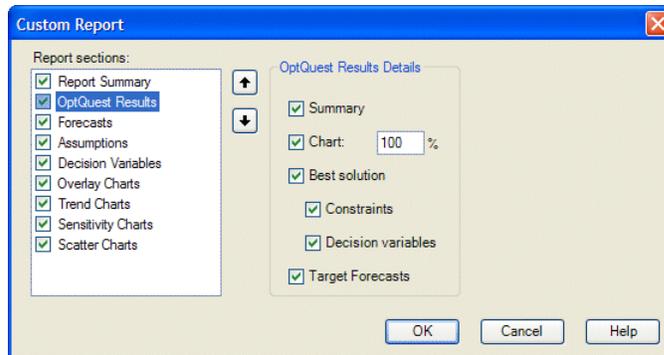
Following an optimization, you can create several different types of OptQuest reports.

- To create an OptQuest report:
 - 1 Run an optimization in OptQuest.
 - 2 Choose **Analyze, Create Report**.
 - 3 In the **Create Report Preferences** dialog, choose:

- Full, to create a complete OptQuest report including simulation results for the best solution
- OptQuest, to create a report with OptQuest results only
- Custom, to display the Custom Report dialog, where you can choose which information — including OptQuest results — to display in the report.

Figure 8 shows elements you can choose to include in the OptQuest Results section of a custom report.

Figure 8 OptQuest Results settings in the Custom Report dialog



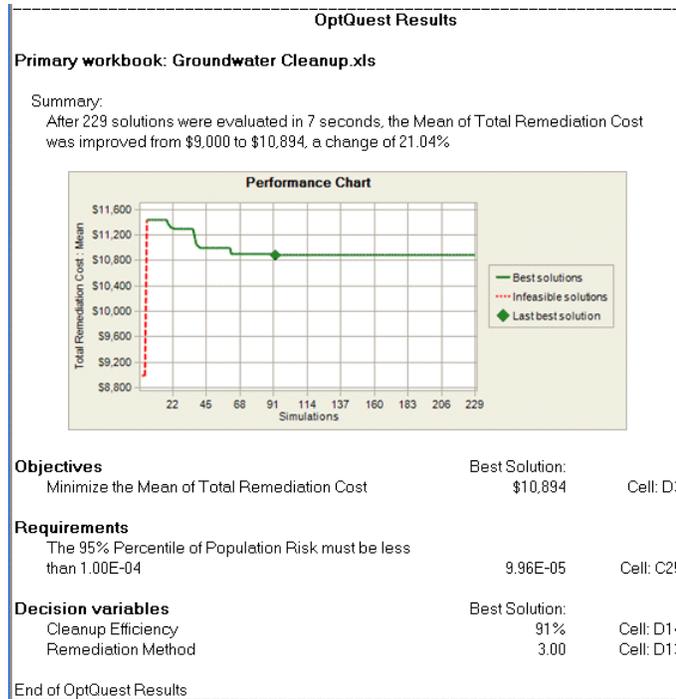
- 4 Click OK in the Create Report Preferences dialog to generate the report (Figure 9).

The first set of information is textual and numerical: related run preferences, run statistics, other statistics (such as number of infeasible solutions), and Crystal Ball data (the number of requirements, constraints, assumptions, decision variables, forecasts, and "frozen" items).

The second set of information is graphical, similar to that shown in Figure 9, and contains information displayed in the OptQuest Results window.

For more information about Crystal Ball reports, see the online *Oracle Crystal Ball User's Guide*.

Figure 9 Graphical OptQuest results in a custom report



Extracting OptQuest Data

► To extract several types of OptQuest data to worksheet cells for further analysis:

1 Run an optimization and choose Analyze, Extract Data.

The Extract Data Preferences dialog opens. By default, the OptQuest Data tab is selected.

2 Choose whether to extract OptQuest solutions, OptQuest statistics, or both, and then indicate whether to extract them for all decision variables or only the ones you choose.

3 Optionally, click the Simulation Data tab to extract simulation data for the best solution only, similar to that described in the *Oracle Crystal Ball User's Guide*.

4 Optionally, click the Options tab to indicate whether to extract data to a new workbook or worksheet and can specify the name to use for that data sheet.

5 When all settings are complete, click OK to extract the data.

Figure 10 shows what happens when you check OptQuest Solutions and OptQuest Statistics. Some OptQuest solutions data rows have been omitted to show the OptQuest statistics data.

Figure 10 Extracted data from Hotel Design.xls

	A	B	C	D	E	F	G
2			Maximize Mean	80% Percentile < 450			
3	Solution #		Total Revenue	Total room demand	Standard price	Platinum price	Gold price
4	302		\$40,406.61	450	\$80.00	\$133.00	\$110.00
5	65		\$40,387.99	450	\$80.00	\$134.00	\$109.00
6	61		\$40,364.52	449	\$80.00	\$134.00	\$110.00
7	64		\$40,344.47	449	\$80.00	\$135.00	\$109.00
8	57		\$40,321.00	448	\$80.00	\$135.00	\$110.00
9	156		\$40,299.50	449	\$80.00	\$136.00	\$109.00
10	40		\$40,293.91	450	\$80.00	\$137.00	\$107.00
11	87		\$40,276.04	448	\$80.00	\$136.00	\$110.00
12	62		\$40,274.53	449	\$80.00	\$137.00	\$108.00
13	134		\$40,253.10	448	\$80.00	\$137.00	\$109.00
14	798		\$40,250.17	447	\$81.00	\$129.00	\$104.00
15							
82							
83							
84			Objective	Requirements	Decision Variables		
85			Maximize Mean	80% Percentile < 450			
86	Statistics		Total Revenue	Total room demand	Standard price	Platinum price	Gold price
87	Minimum		\$40,085.37	440	\$80.00	\$120.00	\$102.00
88	Mean		\$40,308.74	447	\$80.78	\$127.58	\$107.48
89	Maximum		\$40,447.09	450	\$82.00	\$138.00	\$110.00
90	Std. Dev.		\$91.47	2	\$0.46	\$5.65	\$2.30
91							
92	Notes:						
93	Extracted data for top 10% of solutions						

The output is virtually identical to the information shown in the Solution Analysis view of the OptQuest Results window, including the filtering options and the column sort order. To see a different set of solutions, display the Solution Analysis view and change the options before you choose Analyze, Extract Data.

For more information about extracting data, see the online *Oracle Crystal Ball User's Guide*.

Saving optimization models and settings

When you run an optimization, current settings on the Options panel and Advanced Options dialog are automatically saved in a preference file and will be applied to future optimizations.

Other settings — such as objectives, requirements, and constraint definitions — are saved in the primary workbook selected in the dropdown list on the Objectives panel. They are saved to the workbook when the optimization runs, however they are not saved permanently until you save the primary workbook itself.

If you choose to copy optimization values to the model, these values appear as the new cell values and are also saved when the model is saved. Each workbook can have one set of optimization settings.

If you click Close in the OptQuest wizard before you run an optimization, OptQuest asks whether to save the settings. If you respond Yes, current settings are saved to the workbook. Otherwise, current settings are discarded and the last saved settings remain.

Closing OptQuest

To exit OptQuest without running an optimization, click Close in the OptQuest wizard.

If you have not saved changes to the optimization settings yet, OptQuest prompts you to save them to the primary workbook.

Setting Up Efficient Frontier Analysis in OptQuest

Efficient Frontier analysis calculates the curve that plots an objective value against changes to a requirement or constraint. A typical use is for comparing portfolio returns against different risk levels so that investors can maximize return and minimize risk. For a theoretical discussion, see [“Efficient Frontier Analysis” on page 19](#). For an illustration of an Efficient Frontier chart, see [Figure 7](#).

To request an Efficient Frontier analysis in OptQuest, you need to define a requirement or a constraint with a variable bound in either the Objectives or Constraints panel of the OptQuest wizard.

- To define a variable bound for Efficient Frontier analysis:
 - 1 In the Objectives panel, select an existing requirement to modify or add a new one and select it.
Alternately, select a constraint in the Constraints panel.
 - 2 Click **Efficient Frontier**.
 - 3 An **Efficient Frontier** row opens below the requirement or constraint. Adjust the underlined elements to define a range of values for one or both bounds of the requirement or constraint.

When you define a range for a requirement or constraint bound (instead of a single point), you also define a number of points to check within the range by setting the step amount. OptQuest runs one full optimization for each test point in the range, starting with the most limiting requirement test point. Then, you can see the effects of tightening or loosening a requirement.

Efficient Frontier Variable Bound Example

In [“Tutorial 2 — Portfolio Allocation Model” on page 54](#), the investor wants to impose a condition that limits the standard deviation of the total return. Because the standard deviation is a forecast statistic and not a decision variable, this restriction is a requirement.

However, if the investor wants to see if a small increase in the requirement could create a sharp increase in the investment return, the investor can set this as a requirement with a variable upper bound (since this limits the maximum standard deviation). The investor can define this upper bound with a lower limit of \$8,000 and an upper limit of \$10,000. For an example of this technique, see [Portfolio Revisited.xls](#).

Transferring Settings from .opt Files

OptQuest versions shipped with previous releases of Crystal Ball (prior to 11.1.1.x) stored optimization settings in .opt files. As described in [“Saving optimization models and settings” on page 45](#), this version of OptQuest saves settings in workbooks. An .opt file viewer is available to help you transfer settings from .opt files into this version.

► To use the .opt file viewer:

1 Open an optimization model created in a version of Crystal Ball earlier than 11.1.1. The model should have at least one forecast and one decision variable defined. They can be "dummy" data cells, and you can delete them later if you need to.

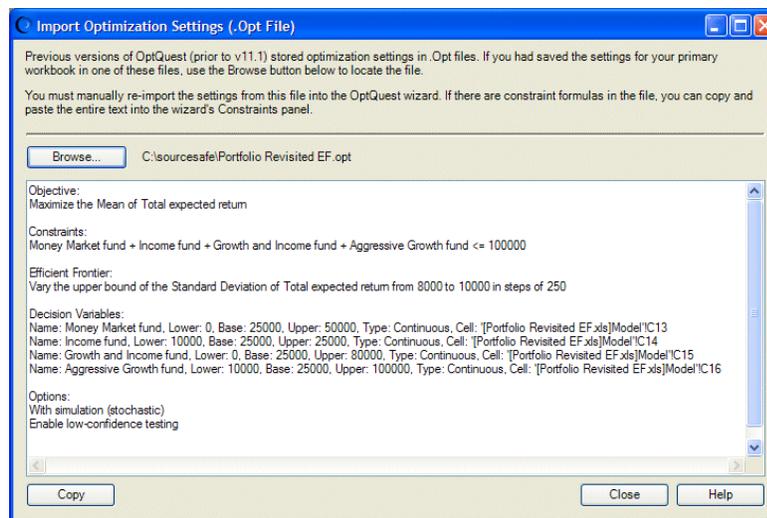
2 Choose Run, OptQuest, or click  .

3 When the Objectives panel opens, click Import.

The Import Optimization Settings dialog opens.

4 Click Browse to locate the .opt file. When you reach its folder, double-click the file. Its settings appear in the Import Optimization Settings dialog (Figure 11, following).

Figure 11 Imported settings for Portfolio Revisited EF.xls



The objective and any requirements or constraints appear at the top. Decision variables and options appear at the bottom.

If the .opt file contains a variable requirement for Efficient Frontier analysis, it is displayed above the decision variables and is labeled "Efficient Frontier," as shown in Figure 11, above.

The Options settings indicate whether the optimization is stochastic or deterministic and whether low-confidence testing is enabled to automatically stop the optimization when specified conditions are met.

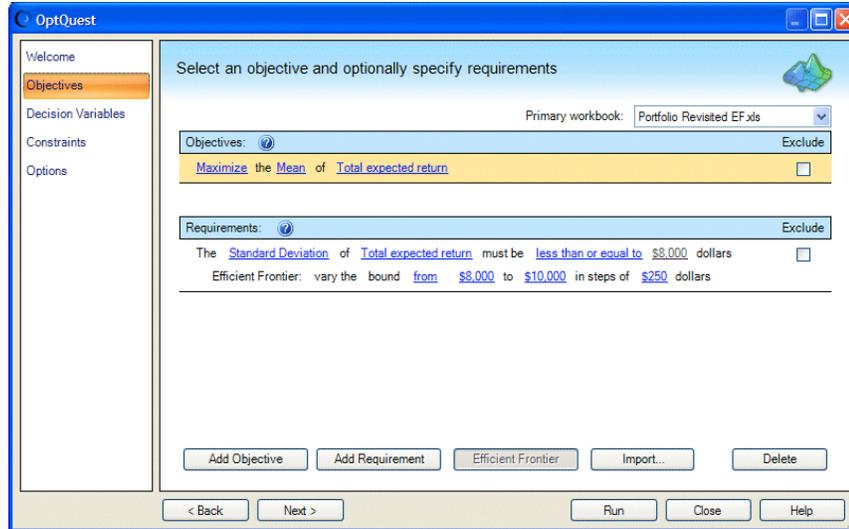
Once you have imported the .opt file into the viewer, you can start transferring the information to each panel of the OptQuest wizard.

5 On the Objectives panel, add an objective and set it to match the text in the viewer.

6 If there are any requirements or variable requirements, add those and edit them to match the text.

Figure 12 shows how to enter the objective and the requirement labeled Efficient Frontier in Figure 11.

Figure 12 Objectives panel, Portfolio Revisited EF.xls



- 7 Enter any constraints on the Constraints panel. You can select one or more constraint formulas in the viewer, click the Copy button, and then paste the constraint(s) into an empty constraint row using Ctrl-v. If you paste more than one constraint, each is automatically placed in a separate row.
- 8 If new decision variables are required, they must be added in Crystal Ball. If necessary, you can copy decision variables from the viewer into Notepad, print them, and then use the printout for a reference when creating the new ones.

When all the decision variables have been defined, start OptQuest again. Click the Decision Variables panel to confirm that all have been entered correctly.

- 9 Now you can run the optimization. All your new settings are stored in the workbook and will be saved permanently with the workbook the next time you save it.

If you have more than one .opt file for one workbook, you can store settings in additional workbooks and use them for a single model. For instructions, see [“Maintaining Multiple Optimization Settings for a Model”](#) on page 120.

Learning More About OptQuest

To learn more about OptQuest, complete the tutorials in [Chapter 4, “OptQuest Tutorials.”](#) Then, review the examples in [Chapter 5, “Examples Using OptQuest.”](#) For further information, check the Crystal Ball website for training opportunities:

<http://www.oracle.com/crystalball>.

4

OptQuest Tutorials

In This Chapter

Introduction.....	49
Tutorial 1 – Futura Apartments Model.....	49
Tutorial 2 – Portfolio Allocation Model	54

Introduction

The first tutorial, the Futura Apartments model, is an extension of the model used in the first Crystal Ball tutorial in the *Oracle Crystal Ball User's Guide* and finds the optimal rent for an apartment building. This model is virtually ready to run, so you can quickly see how OptQuest works.

The second tutorial, the Portfolio Allocation model, shows how to set up and define an optimization yourself. This model finds the optimal set of investments that balances the risk and the return of an investment portfolio.

Tutorial 1 – Futura Apartments Model

Suppose that you have recently purchased the Futura Apartments complex. One of your critical decisions is the amount of rent to charge. You have researched the situation and created a spreadsheet model to help you make a knowledgeable decision.

From the analysis of the price structures and occupancy rates of similar apartment complexes, you have estimated that demand for rental units is a linear function of the rent charged, and is expressed as:

$$\text{Number of units rented} = -.1(\text{rent per unit}) + 85$$

for rents between \$400 and \$600.

In addition, you have estimated that operating costs will average about \$15,000 per month for the entire complex.

Note: You can use Predictor, supplied with Crystal Ball, to find the linear relationship of a dependent variable to one or more independent variables.

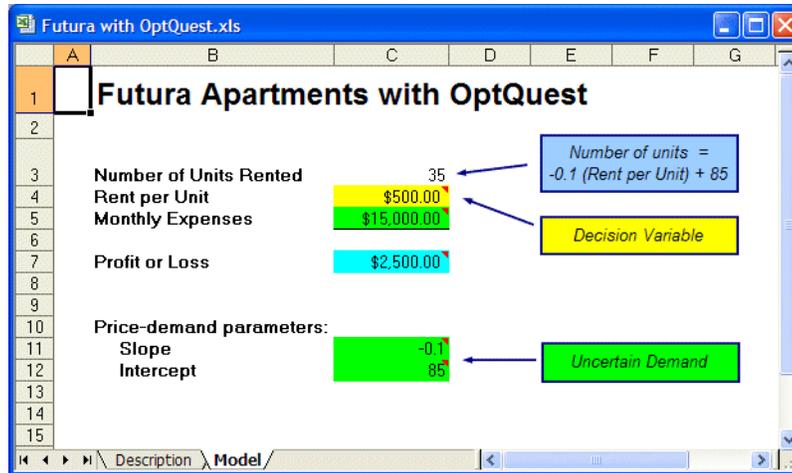
► To begin the tutorial:

- 1 Start Crystal Ball.
- 2 Open the Futura With OptQuest workbook from the Crystal Ball Examples folder.

This spreadsheet is an enhanced version of the original Futura Apartments example in Crystal Ball. This version contains decision variables.

The Futura Apartments worksheet opens as shown in [Figure 13](#).

Figure 13 Futura Apartments worksheet



In this spreadsheet, the rent is set to \$500, where:

$$\text{Number of units rented} = -.1(500) + 85 = 35$$

and the total profit will be \$2,500. If all the data were certain, the optimal value for the rent could be found using a simple data table. However, in a more realistic situation, monthly operating costs and the price-demand function parameters (-.1 and 85) are not certain (probability distributions for these assumptions are already defined for this example). Therefore, determining the best rental price is not a straightforward exercise.

- 3 Before running OptQuest, select Run, Run Preferences and set the following run preferences:
 - Maximum number of trials to run set to 1000 (the default)
 - Sampling method set to Latin Hypercube
 - Sample Size For Latin Hypercube set to 500
 - Random Number Generation set to Use Same Sequence Of Random Numbers with an Initial Seed Value of 999

Running OptQuest

► Use the following steps to start OptQuest and optimize the Futura Apartments model.

- 1 To start OptQuest choose Run, OptQuest, or click .

The OptQuest wizard starts.

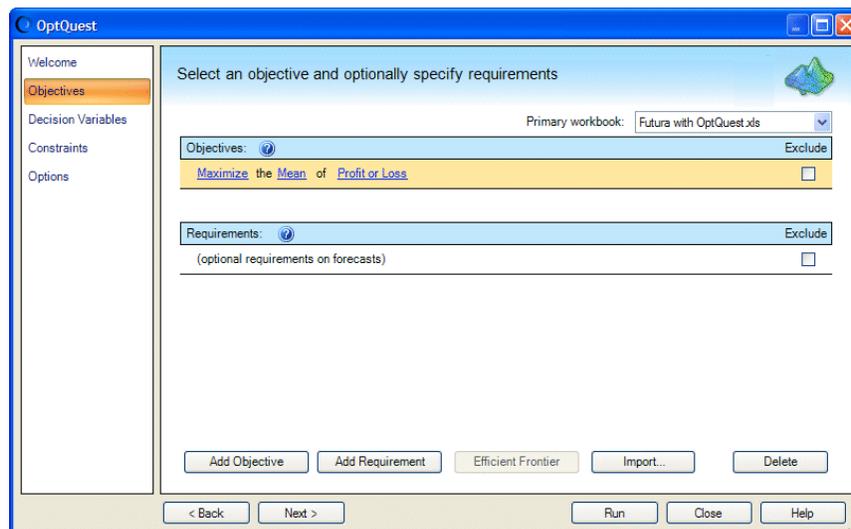
If this is the first time you have run OptQuest, the OptQuest Welcome panel opens. Otherwise, the Objectives panel opens.

Note: Notice the text at the bottom of the Welcome panel that says all OptQuest settings will be stored in the workbook when you run an optimization.

2 If the Welcome panel opens, click Next.

The Objectives panel opens (Figure 14).

Figure 14 Objectives panel, Futura with OptQuest example



The objective for this example is to maximize the mean of the Profit or Loss forecast.

3 To define an objective, click Add Objective. (For this example, the objective has already been added.) A default objective is displayed in the Objectives list:

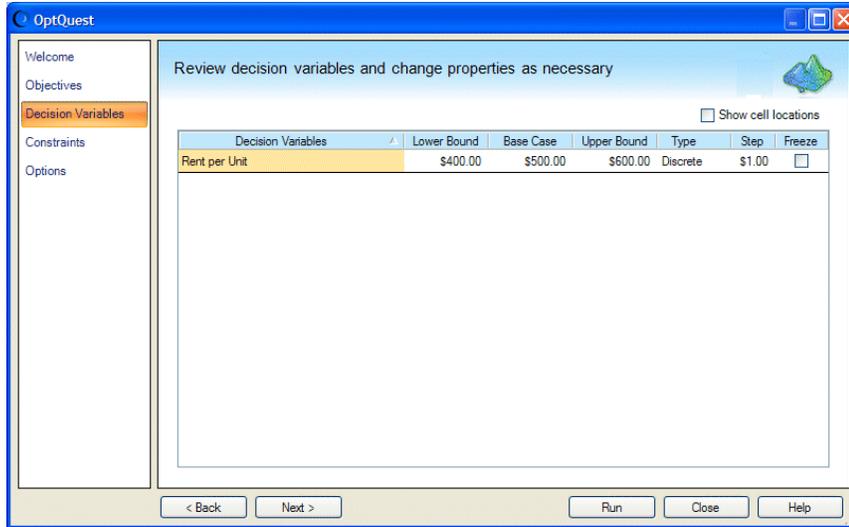
Maximize the Mean of Profit or Loss

This is the desired objective, so no further editing is necessary.

4 Click Next to continue.

The Decision Variables panel opens, as shown in Figure 15.

Figure 15 Decision Variables panel, Futura with OptQuest example



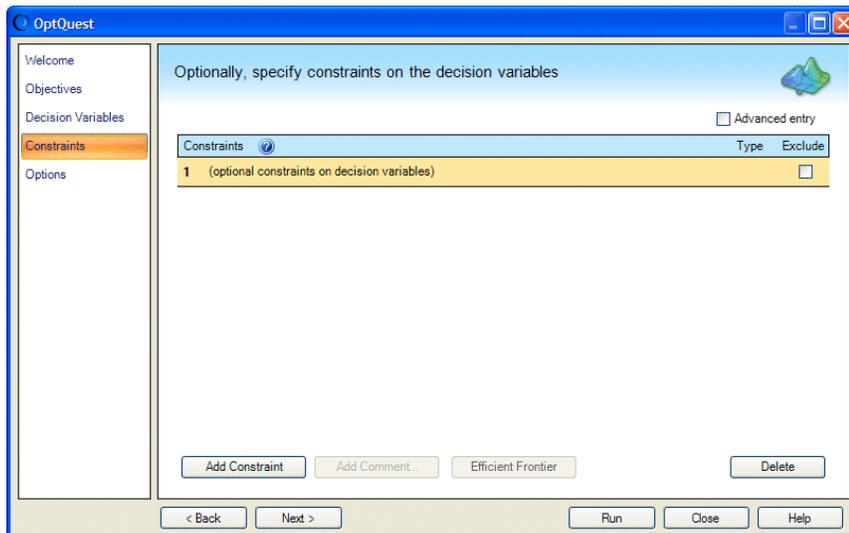
- 5 The Decision Variables panel shows one decision variable, Rent Per Unit.

The lower bound on the variable is 400, the upper bound is 600, and the base case is 500 (the current value in the worksheet). The variable type is listed as Discrete. Because Freeze is not checked, this decision variable will be included in the OptQuest simulation.

- 6 Click Next to continue.

The Constraints panel opens, as shown in Figure 16.

Figure 16 Constraints panel, Futura with OptQuest example

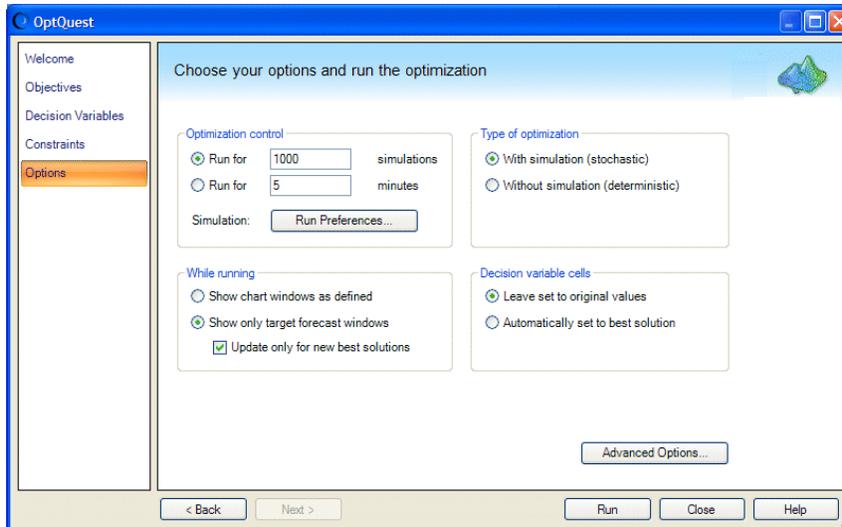


This example has no constraints on the decision variables, so do not add any here.

- 7 Click Next in the Constraints panel.

The Options panel opens.

Figure 17 Options panel, Futura with OptQuest example

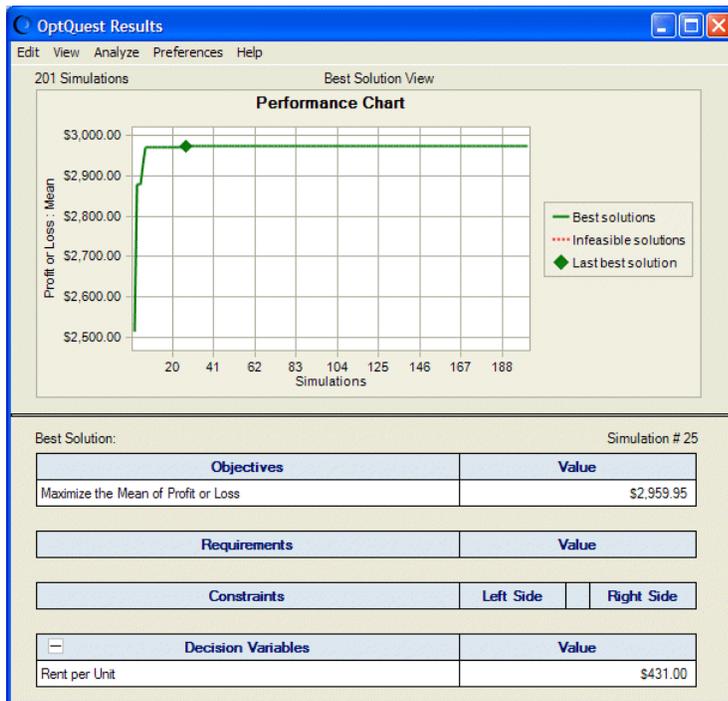


- 8 Set OptQuest to run for 1000 simulations, the default.
- 9 Click Run in the Options panel.

OptQuest systematically searches among the set of **feasible solutions** for ones that improve the mean value of the Profit Or Loss forecast.

In a short time, OptQuest finds the best solution and displays the OptQuest Results window (Figure 18).

Figure 18 OptQuest results for Futura Apartments model



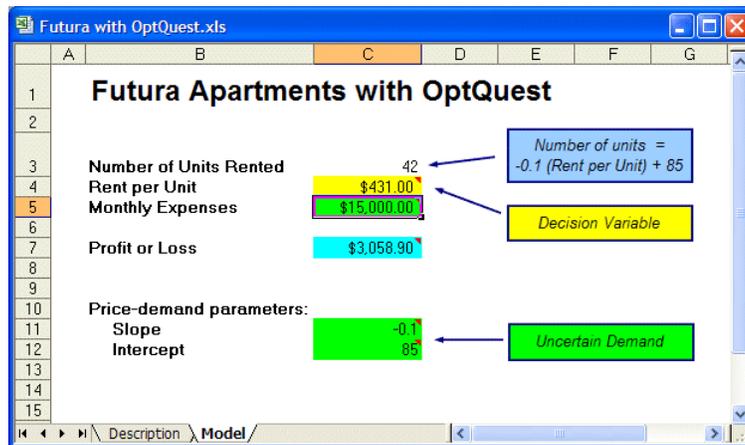
The performance chart shows solutions calculated by OptQuest. Numeric results appear in the table below the chart. For this optimization, the best solution was found at simulation 25. The optimum rent of \$431 per unit produced a maximum mean expected profit of \$2,959.95.

As you requested in the Options panel, a forecast chart for the best solution is displayed. If you choose View, Statistics in the forecast chart menubar, you can see that the mean of the displayed forecast distribution is equal to the maximum mean expected profit shown in the OptQuest Results window (\$2,959.95).

- Choose Edit, Copy Best Solution to Spreadsheet in the OptQuest Results window menu bar.

If you look at the Futura with OptQuest workbook, you can see that cell C4, the decision variable, is now set to the Rent per Unit value that OptQuest calculated, \$431. The spreadsheet holds deterministic calculations based on the optimal value of the decision variable. Because cells C3 and C7 contain formulas that include C4, the values of those cells have also changed, as shown in Figure 19. Now, you need to rent 42 units at \$431 each to obtain maximum profit of about \$3,059.

Figure 19 Futura with OptQuest optimized for maximum profit



Note: When you run an optimization, wizard settings are automatically saved to your workbook. For details, see “Saving optimization models and settings” on page 45.

Tutorial 2 – Portfolio Allocation Model

This is a more detailed tutorial that will guide you through setting up and running an optimization model using Crystal Ball Decision Optimizer with OptQuest. If you are not familiar with basic optimization terminology, such as "objectives" and "constraints," review [Chapter 2, “Overview.”](#)

Problem Description

An investor has \$100,000 to invest in four assets. Below is a list of the assets' expected annual returns and the minimum and maximum amounts the investor is comfortable allocating to each investment.

Table 3 Portfolio Allocation expected returns and investment bounds

Investment	Annual return	Lower bound	Upper bound
Money market fund	3%	\$0	\$50,000
Income fund	5%	\$10,000	\$25,000
Growth and income fund	7%	\$0	\$80,000
Aggressive growth fund	11%	\$10,000	\$100,000

The source of uncertainty in this problem is the annual return of each asset. The more conservative assets, the Income and Money Market funds, have relatively stable annual returns, while the Aggressive Growth fund has higher volatility.

The decision problem, then, is to determine how much to invest in each asset to maximize the total expected annual return while maintaining the risk at an acceptable level and keeping within the minimum and maximum limits for each investment.

Using OptQuest

Using OptQuest involves the following steps:

1. Create a Crystal Ball model of the problem.
2. Define the decision variables within Crystal Ball.
3. Start OptQuest.
4. In OptQuest, define a forecast objective and any requirements.
5. Select decision variables to optimize.
6. Specify any constraints on the decision variables.
7. Select optimization settings.
8. Run the optimization.
9. Interpret the results.

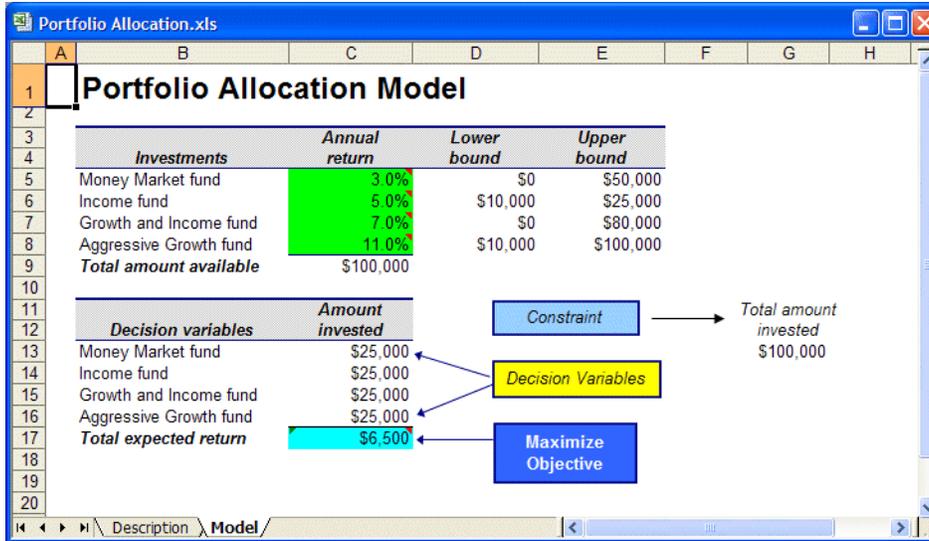
Creating the Crystal Ball Model

► In this case, the model has already been created for you. To review it:

- 1 In Excel, open the **Portfolio Allocation** workbook from the **Crystal Ball Examples** folder.

The worksheet for this problem is shown in [Figure 20](#), following.

Figure 20 Portfolio Allocation worksheet



In this example, problem data values are specified in rows 5 through 9. Model inputs (the values of the decision variables), the model output (the forecast objective), and the constraint (the total amount invested) are in the bottom half of the worksheet.

This model already has the assumptions and forecast cells defined in Crystal Ball. The decision variables are defined as part of this tutorial.

2 Make sure the assumptions are defined as follows:

Assumption	Cell	Distribution	Parameters
Money market fund	C5	uniform	minimum: 2% maximum: 4%
Income fund	C6	normal	mean: 5% standard deviation: 5%
Growth and income fund	C7	normal	mean: 7% standard deviation: 12%
Aggressive growth fund	C8	normal	mean: 11% standard deviation: 18%

If you need help viewing or defining assumptions or forecasts, see the *Oracle Crystal Ball User's Guide*.

3 Select Run, Run Preferences, , and set the following run preferences:

- Maximum number of trials to run set to 1000
- Sampling method set to Latin Hypercube
- Sample Size For Latin Hypercube set to 500

- Random Number Generation set to Use Same Sequence Of Random Numbers with an Initial Seed Value of 999

Defining Decision Variables

➤ The next step is to identify and define decision variables in the model. OptQuest models must have at least one decision variable.

1 Define the first decision variable.

- a. Select cell C13.
- b. Select Define, Define Decision, .
- c. Set the Variable Type to Continuous.
- d. Set the lower and upper bounds according to the problem data (columns D and E in the worksheet), as shown in [Table 3](#) and [Figure 20](#).

Notice that you can enter cell references for cells D5, E5, and the fund name (cell B5). After you complete an entry, the cell reference changes to its value.

2 Define the decision variables for cells C14, C15, and C16 according to the values in columns D and E of the worksheet, by following the process described in step 1.

If you used cell references for the name, lower bound, and upper bound of the decision variable defined in step 1, you can use Crystal Ball's Copy Data and Paste Data commands to define the remaining decision variables.

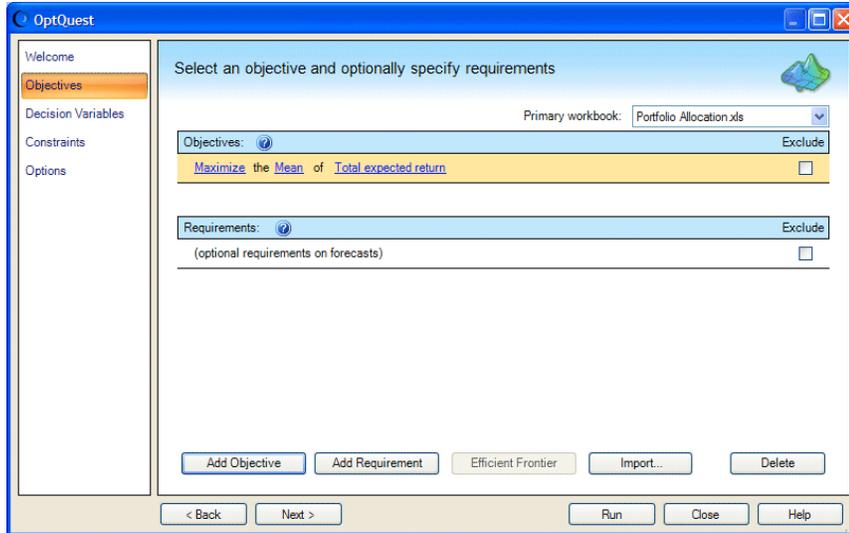
Starting OptQuest and Defining the Forecast Objective

➤ Before you can run an OptQuest simulation, you must define a forecast objective. To do this:

1 Start OptQuest by selecting Run, OptQuest.

You have probably already started OptQuest at least one time, so the Objectives panel opens ([Figure 21](#)).

Figure 21 Objectives panel, Portfolio Allocation example (objective added)



OptQuest requires that you select one forecast statistic to be the **objective** to minimize, maximize, or set to a target value. In addition to defining an objective, you can define optimization **requirements** (described in [“Editing the Optimization Settings” on page 64](#)).

As described earlier, the objective for this example problem is to maximize the total expected return. Since OptQuest, working with Crystal Ball, calculates forecasts as distributions (ranges of values), the mean of the Total Expected Return forecast provides a good representative statistic to use for the objective.

To define an objective, click Add Objective. A default objective is displayed. In this example, the default objective has already been added:

Maximize the Mean of Total Expected Return

This is the desired objective and needs no editing.

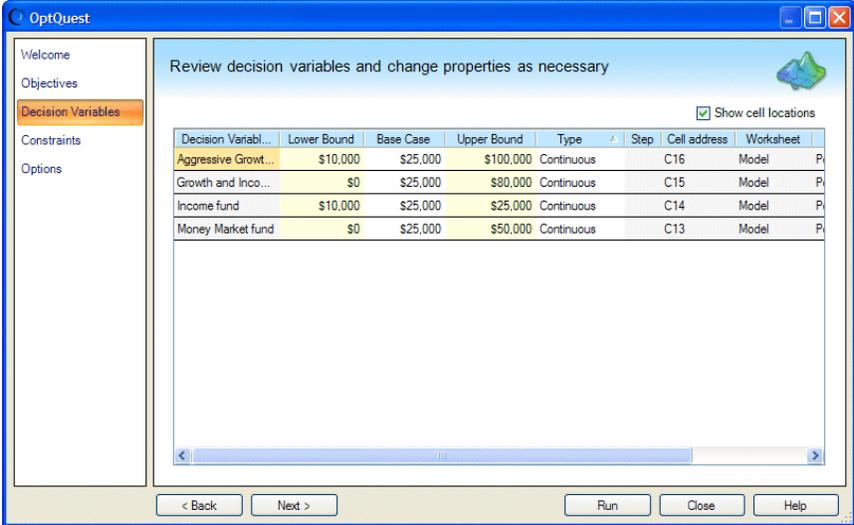
2 Click Next to continue.

The Decision Variables panel opens.

Selecting Decision Variables to Optimize

When you click Next, the Decision Variables panel opens, similar to [Figure 22](#).

Figure 22 Decision Variables panel with cell locations, Portfolio Allocation example



Every decision variable defined in the Crystal Ball model is displayed in the Decision Variables panel. The last column indicates whether the variable has been "frozen," or removed from the optimization. In Figure 22, Show Cell Locations is checked so cell addresses appear before the last column.

The other columns show the bounds, base case (current model value), type, and step for each variable.

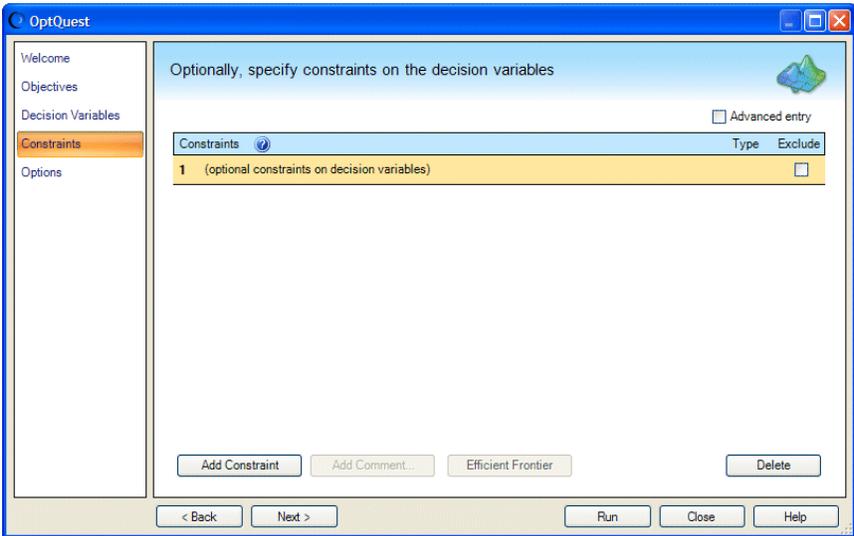
The settings are correct for this example, so click Next to continue.

The Constraints panel opens, as shown in Figure 23.

Specifying Constraints

When you click Next in the Decision Variables panel, the Constraints panel opens.

Figure 23 Constraints panel with no data, Simple Entry mode



Optionally, use the Constraints panel to specify any restrictions that you can define with the decision variables. The constraint in this model limits the initial investment to \$100,000.

By default, the Constraints panel opens in Simple Entry mode. In this mode, most of the constraint formula is entered into cells in your spreadsheet. You then complete the constraint formula on the Constraints panel using a simple conditional expression like `Sheet!A1 <= 100`.

For example, consider the constraint formula given previously as an example:

Money Market fund + Income fund + Growth and Income fund + Aggressive Growth fund = 100000

Each of the fund values is defined in Crystal Ball as a decision variable. In this example, these decision variables are defined in cells C13 through C16, as shown in [Figure 22](#).

The left-hand side of the constraint formula shown above is already entered into cell G13 of the Model worksheet of the Portfolio Allocation example:

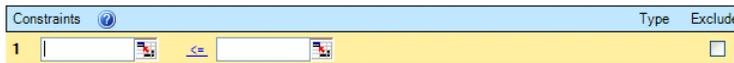
`=SUM(C13:C16)`

➤ To enter this into the Constraints panel:

1 Click Add Constraint.

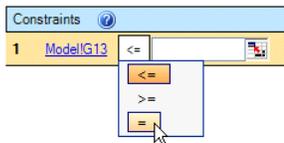
A row with two edit boxes is displayed as shown in [Figure 24](#), following.

Figure 24 Constraints editor in Simple Entry mode



- 2** In the first box from the left, enter the cell that contains the left-hand side of the constraint formula, in this case, cell G13. You can type `=G13` or you can use the cell selector to point to that cell. If the cell has a range name, you can use that instead of the cell address.
- 3** The default operator is `<=`. In this case, the formula calls for `=`. Click the underlined operator and select the one you want ([Figure 25](#)).

Figure 25 Changing the constraint operator



- 4** To enter the right-hand value for the equation, either type a number or reference a cell or range name that contains a value or formula. In [Figure 26](#), following, a number was entered.

Figure 26 A constraint entered in Simple Entry mode



5 At this point, you can:

- Add another constraint,
- Add a comment,
- Add a variable bound for Efficient Frontier analysis,
- Click Next to continue to the Options panel, or
- Click Run to run the optimization.

For more information about adding comments and variable bounds, see [“Constraints Editor and Related Buttons” on page 31](#).

As an alternative, you can enter the constraint formula directly, using Advanced Entry mode. For an example, see [“Specifying Constraints in Advanced Entry Mode” on page 29](#).

6 When constraints settings are complete, click Next to continue.

The Options panel opens, similar to [Figure 17](#).

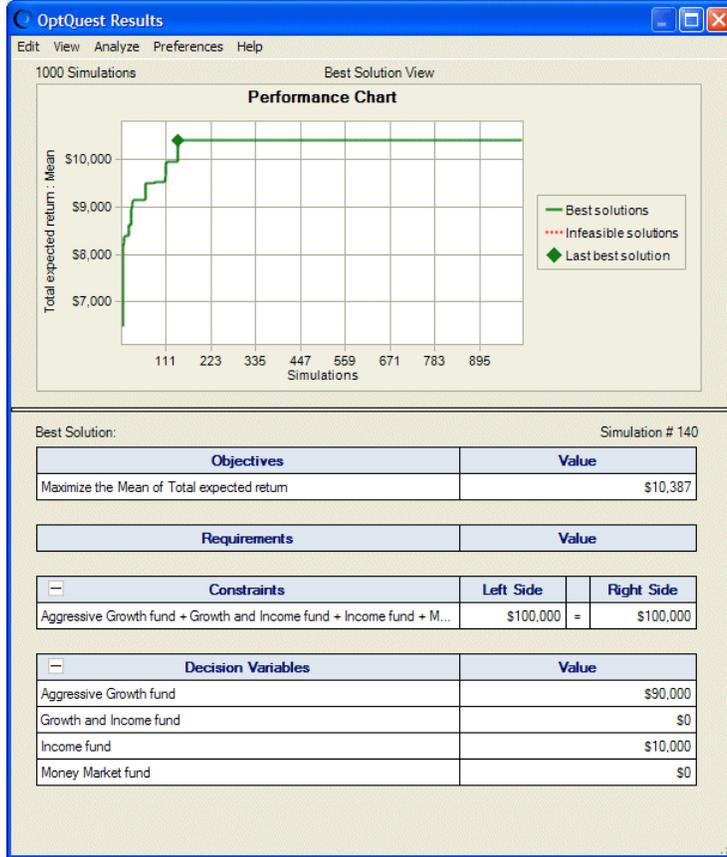
Setting Options and Running the Optimization

- In the Options panel, you set options for controlling the optimization process. For details, click the Help button.

- 1 For this tutorial, set the maximum number of simulations to 1000.**
- 2 Click Run.**

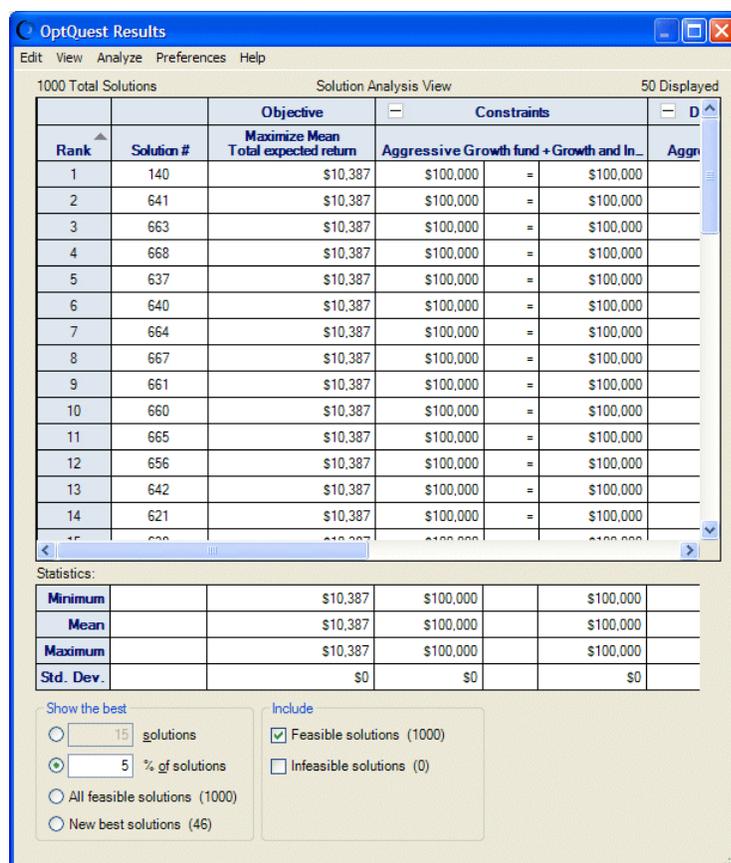
The OptQuest Results window opens ([Figure 27](#)). It is displayed in Best Solution view, which provides an overview of the best solution found during the optimization.

Figure 27 OptQuest Results window – Best Solution view, Portfolio Allocation model



The mean of the Total Expected Return forecast is displayed in the Objectives table. In the Decision Variables table, you can see the amount to allocate to each fund to achieve the objective. If you choose View, Solution Analysis in the menubar, the Solution Analysis tables appear.

Figure 28 OptQuest Results window – Solution Analysis view, Portfolio Allocation model



By default, the solutions list displays the best 5% of solutions ranked by the objective value. If you scroll to the right, you can see the sets of decision variable values that OptQuest tried during its search for the best solution. You can also see the values of requirements and constraint formulas that were calculated based on these decision variables.

The statistics table below the solutions list shows the minimum, mean, maximum, and standard deviation values for the objective, the constraint, and each decision variable (the columns in the table).

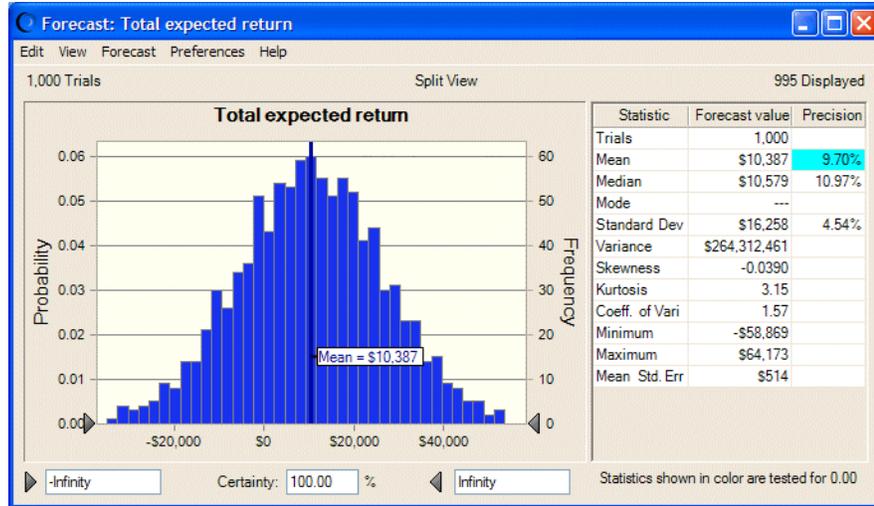
In this case, the investment strategy maximized the return of the portfolio, but at a price: high risk due to high volatility and little diversification. Is this really the best strategy? To find out, the investor must interpret the results.

Interpreting the Results

To interpret the OptQuest results, start by viewing the forecast chart for the best simulation. If it isn't already onscreen, choose Analyze, Forecast Charts and select Total Expected Return.

Figure 29, following, shows the forecast chart and statistics in Split view. Note that the standard deviation of the forecast is quite high, \$16,258, compared to the mean return of \$10,387. The ratio of these two values, the coefficient of variability, is shown as 1.57, or above 150%. Most of the money allocated was in the Aggressive Growth fund, and the uncertainty of returns for that fund was quite high, indicating the relative riskiness of the investment.

Figure 29 Portfolio allocation forecast chart, Split view



Editing the Optimization Settings

In portfolio management, controlling the variability of the solution to minimize risk can be just as important as achieving large expected returns. Suppose that this same investor wants to reduce the uncertainty of returns for the portfolio, while still attempting to maximize the expected return. You might want to find the best solution for which the standard deviation is much lower, say, below \$8,000.

You can edit the OptQuest settings to add this risk limitation and still maximize the total expected return.

➤ To edit OptQuest:

- 1 With **Portfolio Allocation.xls** open, choose **Run, OptQuest**. If you just ran an optimization, click **Reset** in the **OptQuest Control Panel**. When the **Reset** prompt is displayed, check **Launch OptQuest Wizard** and click **Yes**.
- 2 If it is not already open, click **Objectives** in the left pane of the **OptQuest wizard**.

The panel opens with **Maximize the Mean of Total Expected Return** listed as the objective.

- 3 Click **Add Requirement**.

This creates a new row in the **Requirements** area:



- 4 In the new row, click **Mean**. In the list, select **Standard Deviation**.
- 5 Click **greater than or equal to** and change it to **less than or equal to**.
- 6 Then, click **100** and change it to **8000**.

This adds a requirement that the standard deviation of the expected returns must be less than or equal to \$8,000 for a solution to be considered feasible.

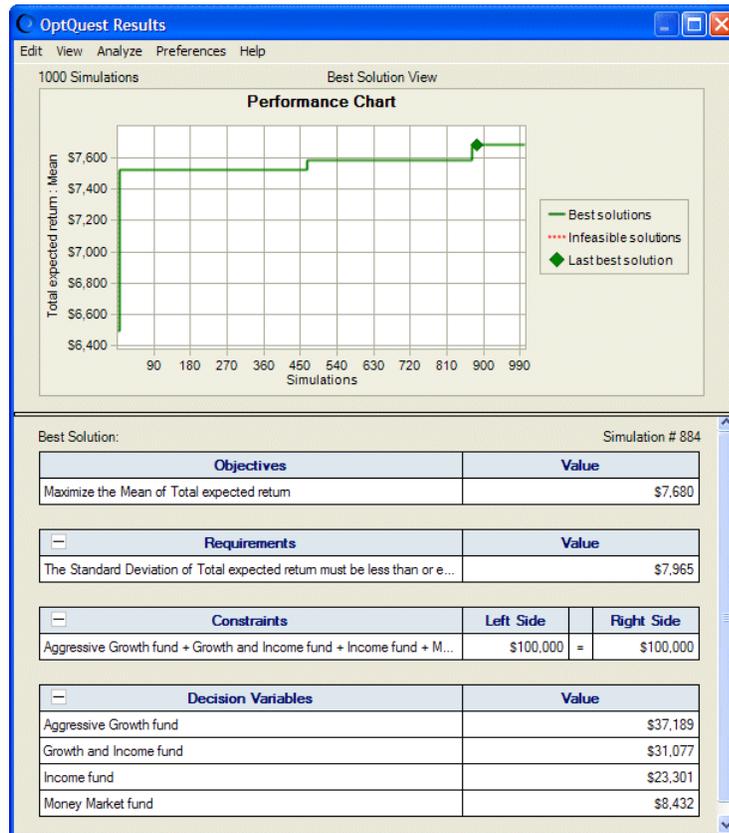
Figure 30 Objectives panel with the new requirement

Objectives: ?		Exclude
Maximize	the Mean of Total expected return	<input type="checkbox"/>
Requirements: ?		
The Standard Deviation of Total expected return must be less than or equal to \$8,000 dollars		<input type="checkbox"/>

7 Click Run.

The new results are shown in Figure 31.

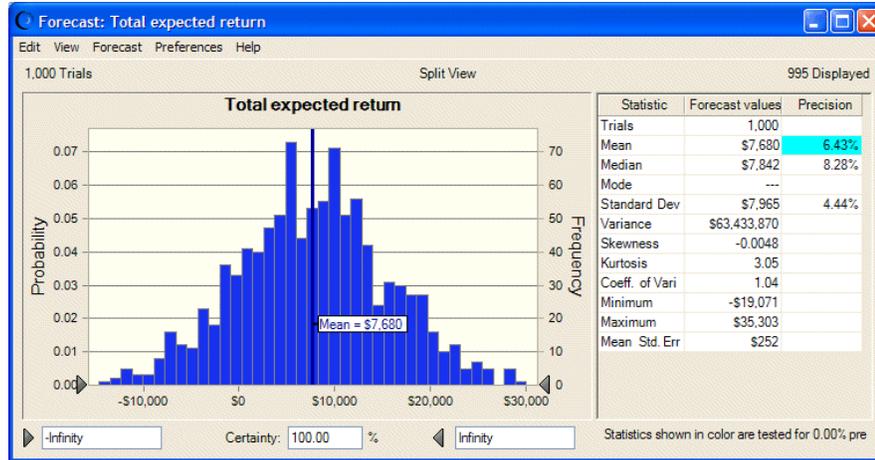
Figure 31 Portfolio allocation optimization results with risk



As shown in Figure 31, after several hundred simulations, OptQuest found a solution that meets the requirement well. The standard deviation of Total Expected Return is just below \$8,000. The objective value, though, is now significantly lower than the previous solution without the limit on risk (Figure 27).

If you return to the Portfolio Allocation model and display the resulting forecast chart in Split View (Figure 32), you can see that the new values appear. The standard deviation of Total Expected Return is just under \$8,000 and the coefficient of variability is slightly greater than 1.

Figure 32 Best optimization solution with lower risk requirement



Interpreting Results

This solution has significantly reduced the variability of the total expected return, even though it now has a lower mean return. The portfolio achieved this by finding the best diversification of conservative and aggressive investments. Thus, the investor must face the trade-off between higher returns with higher risk, and lower returns with lower risk.

How does this solution compare with the high-risk solution? You can compare [Figure 29](#) with [Figure 32](#) to answer that question. The mean return is lower in [Figure 32](#), but the standard deviation, variance, and coefficient of variability — the risk indicators — are also lower.

Portfolio Allocation Optimization Summary

The best OptQuest solution identified might not be the true optimal solution to the problem, but should be close to the true optimal solution. The accuracy of the results depends on the time limit you select for searching, the number of trials per simulation, the number of decision variables, and the complexity of the problem. With more decision variables, you need a larger number of simulations. Further details of the search procedure can be found in [Appendix A](#), “[Optimization Tips and Notes](#),” and “[References](#)” on page 129.

After solving an optimization problem with OptQuest, run a longer Crystal Ball simulation using the optimal values of the decision variables to more accurately compute the risks of the recommended solution.

5

Examples Using OptQuest

In This Chapter

Overview	67
Product Mix	69
Hotel Design and Pricing Problem	72
Budget-constrained Project Selection	75
Groundwater Cleanup	79
Oil Field Development	83
Portfolio Revisited	86
Tolerance Analysis	92
Inventory System Optimization	96
Drill Bit Replacement Policy	102
Gasoline Supply Chain	105

Overview

This chapter presents a variety of examples using OptQuest. These examples illustrate how to use spreadsheets to model optimization problems, the key features of OptQuest, and the variety of applications for which you can use OptQuest.

Each section includes a problem statement, a description and explanation of the spreadsheet model, the OptQuest solution, and optionally additional practice exercises using the model. All Excel model files and associated OptQuest files are in the Examples folder under the main Crystal Ball installation folder. You can also display an index to the examples by choosing one of the following command sequences and selecting from the Examples Guide:

- Help, Crystal Ball, Examples Guide (Excel 2003 or earlier)
- Help, Resources, Examples Guide (Excel 2007)
- Start, All Programs, Crystal Ball, Examples (Windows Start menu)

Table 4, following, summarizes the examples in this chapter and the features illustrated.

Table 4 OptQuest examples

Application	Decision Variables	Type	Constraints	Requirements	Illustrated Methods
Product mix	5	discrete	3	1	Classic optimization example.
Hotel design and pricing	3	discrete	0	1	Uses a percentile requirement; shows the risk of using a deterministic solution instead of a probabilistic one.
Budget-constrained project selection	8	binary (0-1)	1	0	Uses binary decision variables for Yes/No decisions.
Groundwater cleanup	2	mixed	0	1	Uses a category decision variable to select different sets of assumptions.
Oil field development	3	mixed	0	0	Uses a percentile objective and a lookup table based on a decision variable.
Portfolio revisited (including Portfolio Revisited EF)	4	discrete, step = \$100	1	0	Combines several objective functions into one multiobjective using extracted statistics and uses the Arbitrage Pricing Theory for incorporating risk. Example of Efficient Frontier.
Tolerance analysis	7	continuous	0	2	Uses process capability metrics.
Inventory system	2	discrete	0	0	Searches a wide solution space with large steps, and then refines the search.
Drill bit replacement	1	continuous	0	0	Defines time as a decision variable.
Gasoline supply chain	8	discrete	0	1	Classic optimization example.

Note: Most of the examples included here use one of the Advanced Options settings for automatically stopping the optimization when either a solution confidence level or certain number of non-improving solutions is reached. If you follow along with these examples, your results should be similar but may not always be identical.

Product Mix

The following sections describe this problem and its OptQuest solution:

- “Product Mix Problem Statement” on page 69
- “Product Mix Spreadsheet Model” on page 70
- “Product Mix OptQuest Solution” on page 70

Product Mix Problem Statement

Ray's Red Hots, Inc. manufactures five types of sausages. The number of pounds of four ingredients—veal, pork, beef, and casing—used per unit of product and the profit generated per unit are given in the table below.

Table 5 Ray's Red Hots data summary

Products	Veal	Pork	Beef	Casing	Profit Per Unit
Summer Sausage	0.00	2.50	1.00	1.00	\$1.25
Bratwurst	4.00	1.00	0.00	1.50	\$1.80
Italian Sausage	1.00	3.00	1.50	1.00	\$1.40
Pepperoni	0.00	4.00	0.00	2.00	\$2.10
Polish Sausage	0.00	1.00	3.00	1.50	\$1.70

Limited amounts of ingredients are available for the next production cycle. Specifically, only 12,520 pounds of veal, 14,100 pounds of pork, 6,480 pounds of beef, and 10,800 pounds of casing are available.

Complicating this situation is:

- The unit profits are only estimates because all customer contracts have not been finalized.
- The amount of casing used per unit might be more than anticipated because of production losses due to tearing or partial rejections during inspection.

The problem is to determine how many pounds of each product to produce in order to maximize gross profit without running out of meat ingredients or casing during the manufacturing run.

Product Mix Spreadsheet Model

The Product Mix.xls file, shown in Figure 33, is a spreadsheet model for this problem. The input data and model outputs are straightforward.

Figure 33 Product mix problem spreadsheet model

Product Mix - Ray's Red Hots, Inc.							
Products:	Lbs of Ingredients Needed Per Unit				Profit per Unit	Quantity to Produce	
	Veal	Pork	Beef	Casing			
Summer Sausage	0.00	2.50	1.00	1.00	\$1.25	1,000	
Bratwurst	4.00	1.00	0.00	1.50	\$1.80	1,000	
Italian Sausage	1.00	3.00	1.50	1.00	\$1.40	1,000	
Pepperoni	0.00	4.00	0.00	2.00	\$2.10	1,000	
Polish Sausage	0.00	1.00	3.00	1.50	\$1.70	1,000	

Inventory:	Veal	Pork	Beef	Casing
On Hand	12,520.00	14,100.00	6,480.00	10,800.00
Used	5,000.00	11,500.00	5,500.00	7,000.00
Remaining	7,520.00	2,600.00	980.00	3,800.00

Gross Profit: \$8,250.00

Product Mix OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See “Setting Crystal Ball Run Preferences” on page 25.

➤ To run the optimization:

- 1 With Product Mix.xls open in Crystal Ball, set the number of trials in Crystal Ball to 2000, since tail-end percentile requirements need more accuracy.
- 2 Start OptQuest from the Crystal Ball Run menu and click Next to view each wizard panel:
 - The objective is to maximize the mean of gross profit.
 - The only requirement ensures that at most a 5% chance exists of exceeding the casing limitation.
 - This problem has five decision variables (one for each product), and three constraints (one each for availability of veal, pork, and beef).
- 3 Run the optimization.

Figure 34 shows the OptQuest solution. The optimal mean profit is \$11,780.33, obtained by producing 2710 pounds of bratwurst, 1460 pounds of Italian sausage, 1040 pounds of pepperoni, 1200 pounds of Polish sausage, and 660 pounds of summer sausage.

Figure 34 Product mix model optimization results

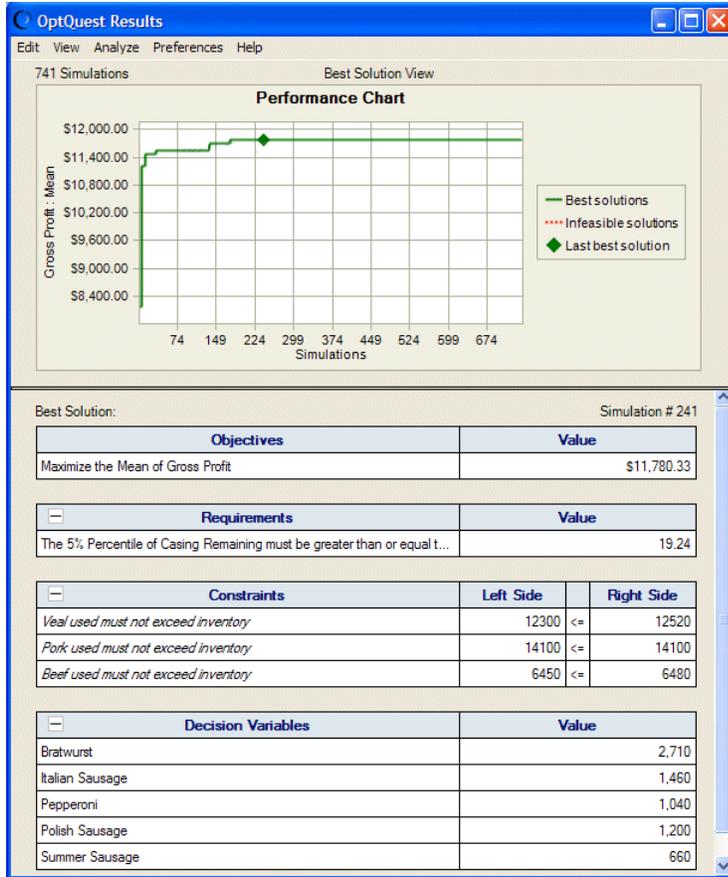
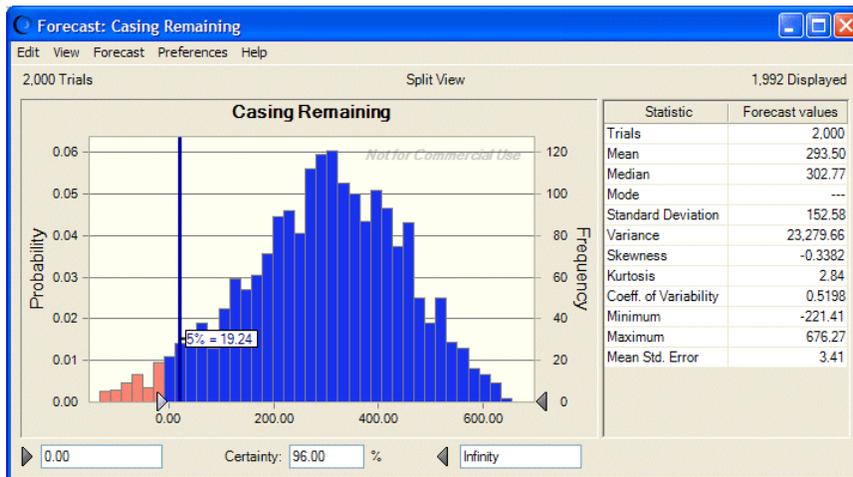


Figure 35 shows the Casing Remaining forecast chart for these decision variables, verifying that the chance of running out of casing is indeed at most 5%.

Figure 35 Product mix - remaining casing forecast chart



Hotel Design and Pricing Problem

A downtown hotel is considering a major remodeling effort and needs to determine the best combination of rates and room sizes to maximize revenues.

The following sections describe this problem and its OptQuest solution:

- [“Hotel Design Problem Statement” on page 72](#)
- [“Hotel Design Spreadsheet Model” on page 73](#)
- [“Hotel Design OptQuest Solution” on page 73](#)

Hotel Design Problem Statement

Currently the hotel has 450 rooms with the following history:

Table 6 Hotel example data summary

Room Type	Rate	Daily Avg. No. Sold	Revenue
Standard	\$85	250	\$21,250
Gold	\$98	100	\$9,800
Platinum	\$139	50	\$6,950

Each market segment has its own price/demand elasticity. Estimates are:

Room Type	Elasticity
Standard	-3
Gold	-1
Platinum	-2

This means, for example, that a 1% decrease in the price of a standard room will increase the number of rooms sold by 3%. Similarly, a 1% increase in the price will decrease the number of rooms sold by 3%. For any proposed set of prices, the projected number of rooms of a given type sold can be found using the formula:

$$\text{rooms sold} = H + \frac{E \cdot H \cdot (N - C)}{C}$$

where variables are:

Variable	Description
H	Historical average number of rooms sold
E	Elasticity
N	New price

Variable	Description
C	Current price

The hotel owners want to keep the price of a standard room between \$70 and \$90, a gold room between \$90 and \$110, and a platinum room between \$120 and \$149. All prices are in whole dollar increments (discrete). Although the rooms may be renovated and reconfigured, there are no plans to expand beyond the current 450-room capacity.

Hotel Design Spreadsheet Model

To follow this example, open the Hotel Design example shown in [Figure 36](#).

Figure 36 Hotel pricing problem spreadsheet model

The screenshot shows a spreadsheet titled "Hotel Pricing Problem" with the following data:

Room type	Rate	Average Daily Sold	Revenue	Elasticity	New Price	Projected Rooms Sold	Projected Revenue
Standard	\$ 85.00	250	\$ 21,250.00	-3	\$ 85.00	250	\$ 21,250.00
Gold	\$ 98.00	100	\$ 9,800.00	-1	\$ 98.00	100	\$ 9,800.00
Platinum	\$ 139.00	50	\$ 6,950.00	-2	\$ 139.00	50	\$ 6,950.00
Total						400	\$ 38,000.00
Capacity						450	

Annotations in the spreadsheet include a yellow box pointing to the "New Price" column with the text "Determine best price of hotel rooms" and a blue box pointing to the "Total" and "Capacity" rows with the text "Maximize revenue subject to capacity".

The decision variables correspond to cells G7 through G9.

Hotel Design OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See “[Setting Crystal Ball Run Preferences](#)” on page 25.

► With Hotel Design.xls open in Crystal Ball:

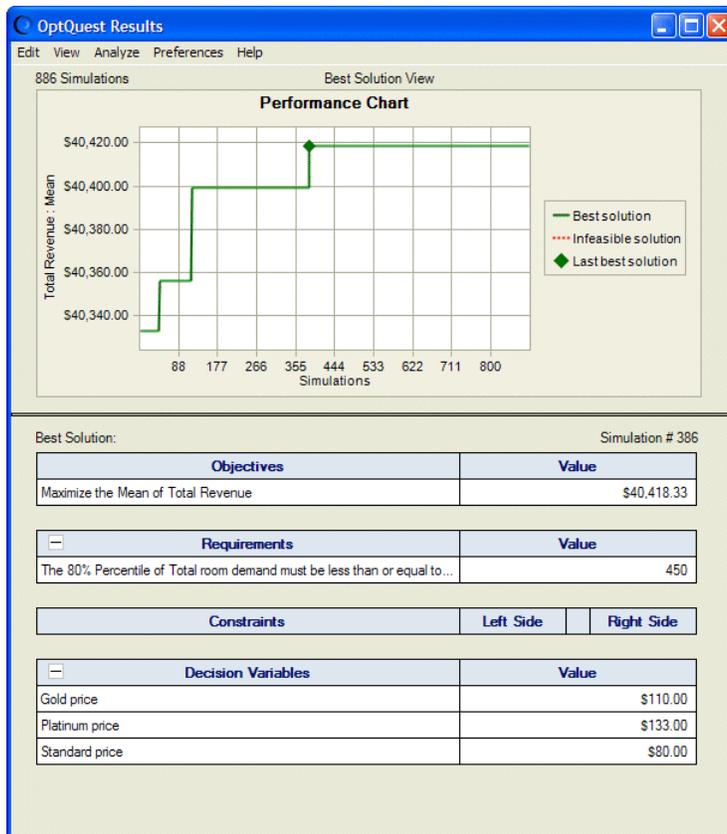
1 Start the OptQuest wizard.

As you click Next to step through the problem, note:

- The objective is to maximize the mean of total revenue.
- To ensure that the probability of demand exceeding capacity does not exceed 20%, the projected number of rooms sold (cell H12) is a forecast in the Crystal Ball model, with a requirement added in the Objectives panel. Specifically, the total room demand is limited by a requirement using the forecast statistic Percentile (80), with an upper bound of 450.

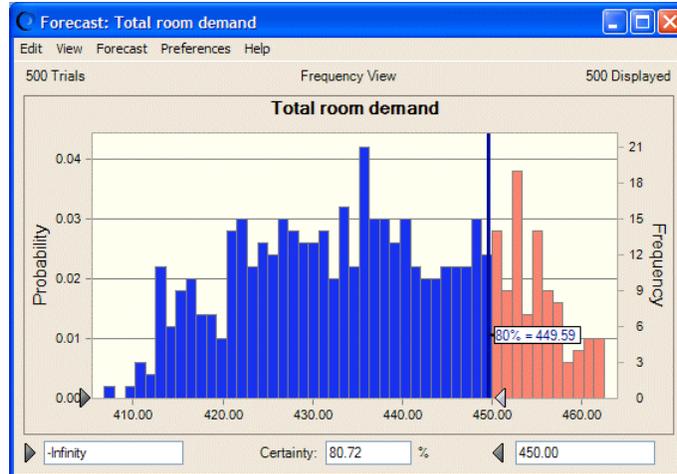
- This problem has three decision variables and no constraints.
- 2 Run the optimization.
- The results appear in [Figure 37](#).

Figure 37 Hotel pricing model optimization results



The Crystal Ball simulation of this solution in [Figure 38](#) verifies that the chance of demand exceeding capacity is just slightly less than 20% (100% – 80.72%).

Figure 38 Hotel pricing solution (percentiles view)



Budget-constrained Project Selection

This example concerns project selection for maximum profitability.

The following sections describe this problem and its OptQuest solution:

- [“Project Selection Problem Statement” on page 75](#)
- [“Project Selection Spreadsheet Model” on page 76](#)
- [“Project Selection OptQuest Solution” on page 77](#)

Project Selection Problem Statement

The R&D group of a major public utility has identified eight possible projects. A net present value analysis has computed:

- The expected revenue for each if it is successful
- The estimated probability of success
- The initial investment required for each project

Using these figures, the finance manager has computed the expected return and the expected profit for each project as shown in the following table.

Table 7 Project analysis example data summary

Project	Expected Revenue	Success Rate	Expected Return	Initial Investment	Expected Profit
1	\$750,000	90%	\$675,000	\$250,000	\$425,000
2	\$1,500,000	70%	\$1,050,000	\$650,000	\$400,000
3	\$600,000	60%	\$360,000	\$250,000	\$110,000

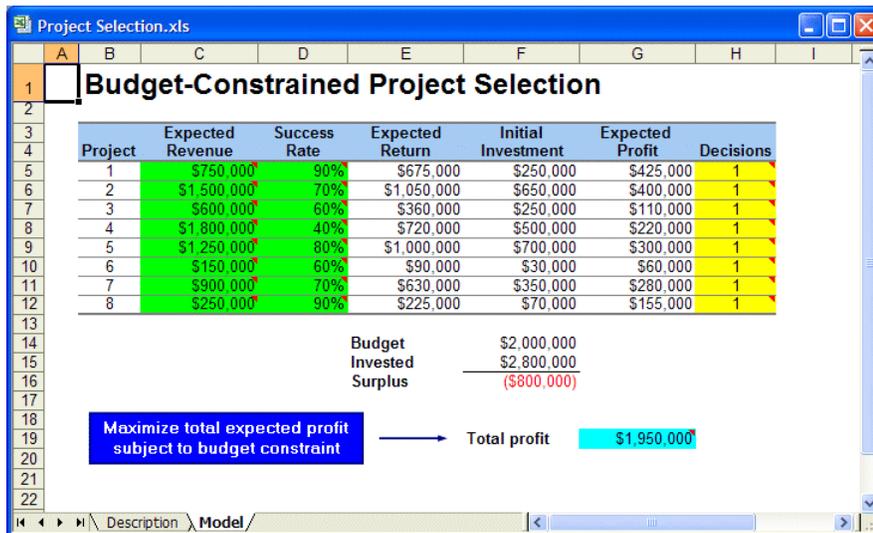
Project	Expected Revenue	Success Rate	Expected Return	Initial Investment	Expected Profit
4	\$1,800,000	40%	\$720,000	\$500,000	\$220,000
5	\$1,250,000	80%	\$1,000,000	\$700,000	\$300,000
6	\$150,000	60%	\$90,000	\$30,000	\$60,000
7	\$900,000	70%	\$630,000	\$350,000	\$280,000
8	\$250,000	90%	\$225,000	\$70,000	\$155,000
			Total invested	\$2,800,000	Total profit
			Budget	\$2,000,000	\$1,950,000

Unfortunately, the available budget is only \$2.0 million, and selecting all projects would require a total initial investment of \$2.8 million. Thus, the problem is to determine which projects to select to maximize the total expected profit while staying within the budget limitation. Complicating this decision is the fact that both the expected revenue and success rates are highly uncertain.

Project Selection Spreadsheet Model

Figure 39 shows a spreadsheet model for this problem, which you can view by opening the Project Selection.xls file. The decision variables in column H are binary; that is, they can assume only the values zero and one, representing the decisions of either not selecting or selecting each project. The total investment in cell F15 is the required investment in column F multiplied by the respective decision variable in column H.

Figure 39 Project selection problem spreadsheet model



The expected revenue and success rates are assumption cells in the Crystal Ball model. The expected revenues have various distributions, while the success rates are modeled using a binomial distribution with one trial. During the simulation, the outcomes in column D will be either 0% or 100% (not successful or successful) with the probabilities initially specified. Thus, for each simulated trial, the expected returns will either equal the expected revenue generated in column C or zero. Consequently, the expected profits can be positive or negative.

Although good solutions might be identified by inspection or by trial and error, basing a decision on expected values can be dangerous because it doesn't assess the risks. In reality, selecting R&D projects is a one-time decision; each project will be either successful or not. If a project is not successful, the company runs the risk of incurring the loss of the initial investment. Thus, incorporating risk analysis within the context of the optimization is a very useful approach.

Project Selection OptQuest Solution

- ▶ With Project Selection.xls open in Crystal Ball, start OptQuest from the Crystal Ball Run menu. Then:

1 Start the OptQuest wizard.

As you click Next to step through the problem, note that:

- The objective is to maximize the final value of total profit.
- There are no requirements.
- There are eight decision variables.
- There is one constraint representing the budget limitations. Note the use of Excel's SUMPRODUCT function in the constraint to create a linear combination of the decision variables and investment amounts.

2 Run the optimization.

Figure 40 shows the results of an OptQuest optimization. The best solution identified selects all the projects except for 3, 4, and 8.

Figure 40 Project selection model optimization results

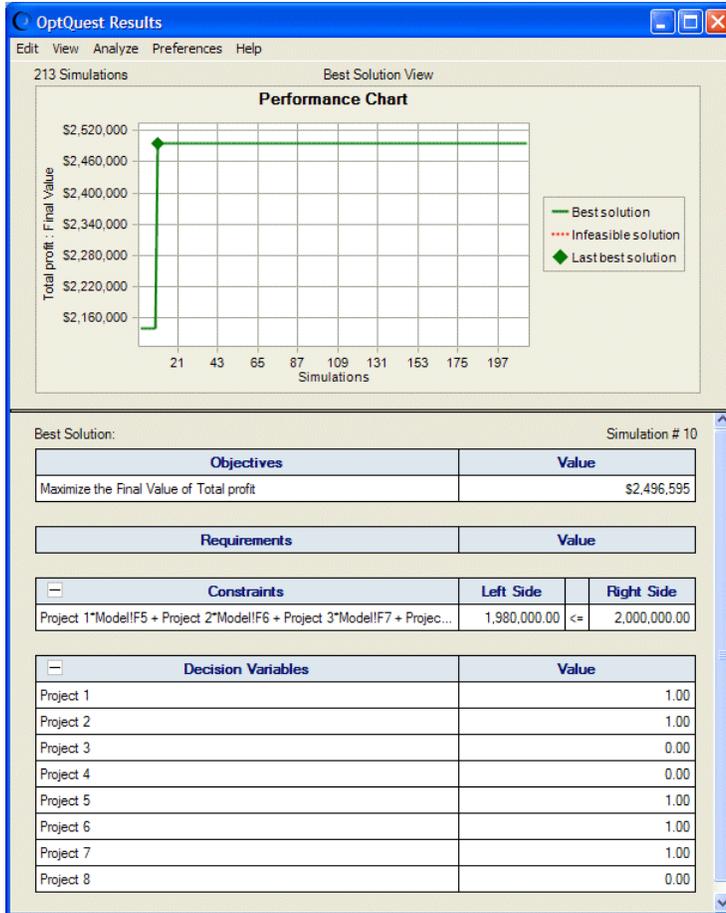
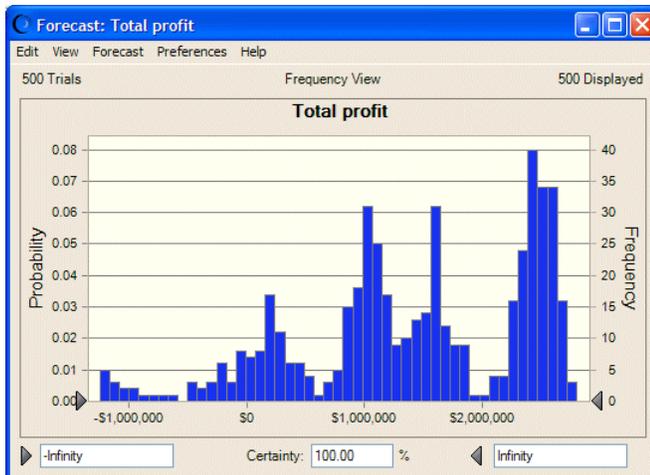


Figure 41, the forecast chart for Total Profit, shows that the distribution of profits is highly irregular, and depends on the joint success rate of the chosen projects. There is a risk of realizing a loss. You might wish to evaluate the risks associated with some of the other solutions identified during the search.

Figure 41 Project selection solution forecast chart



Groundwater Cleanup

This example concerns choosing a method for cleaning up groundwater contamination.

The following sections describe this problem and its OptQuest solution:

- “Groundwater Cleanup Problem Statement” on page 79
- “Groundwater Cleanup Spreadsheet Model” on page 79
- “Groundwater Cleanup OptQuest Solution” on page 81

Groundwater Cleanup Problem Statement

A small community gets its water from wells that tap into an old, large aquifer. Recently, an environmental impact study found toxic contamination in the groundwater due to improperly disposed chemicals from a nearby manufacturing plant. Since this is the community’s only source of potable water and the health risk due to exposure to these chemicals is potentially large, the study recommends that the community reduce the overall risk to below a 1 in 10,000 cancer risk with 95% certainty (95th percentile less than 1E-4).

A task force narrowed down the number of appropriate treatment methods to three. It then requested bids from environmental remediation companies to reduce the level of contamination down to recommended standards, using one of these methods.

Your remediation company wants to bid on the project. The costs for the different cleanup methods vary according to the resources and time required for each (cleanup efficiency). With historical and site-specific data available, you want to find the best process and efficiency level that minimizes cost and still meets the study’s recommended standards with a 95% certainty.

Complicating the decision-making process:

- You have estimates of the contamination levels of the various chemicals. Each contaminant’s concentration in the water is measured in micrograms per liter.
- The cancer potency factor (CPF) for each chemical is uncertain. The CPF is the magnitude of the impact the chemical exhibits on humans; the higher the cancer potency factor, the more harmful the chemical is.
- The population risk assessment must account for the variability of body weights and volume of water consumed by the individuals in the community per day.

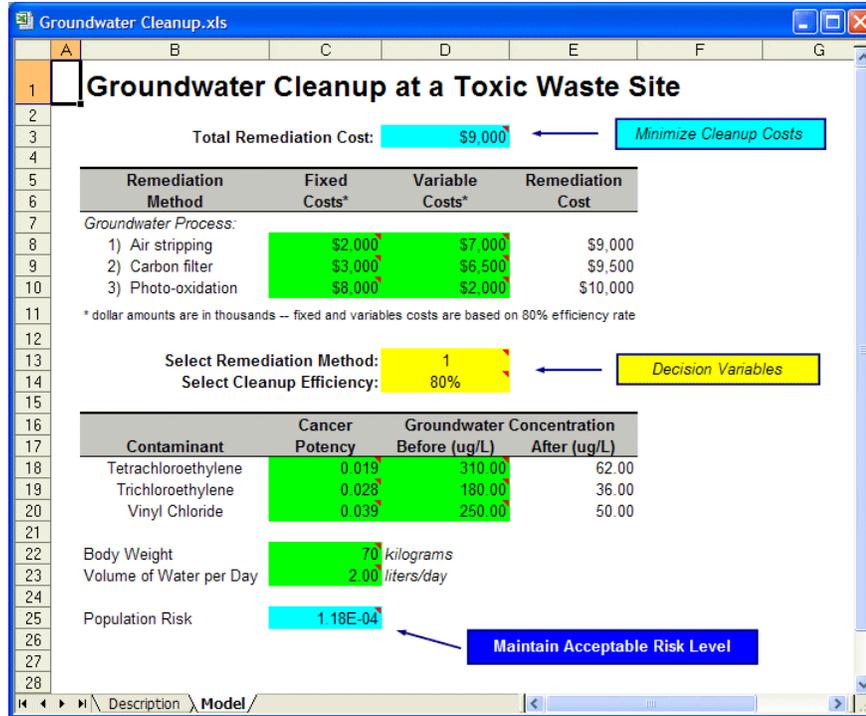
All these factors lead to the following equation for population risk:

$$\text{population risk} = \frac{\text{cancer potencies} \bullet \text{contaminant concentrations} \bullet \text{water consumed per day}}{\text{body weight} \bullet \text{conversion factor}}$$

Groundwater Cleanup Spreadsheet Model

Open the file Groundwater Cleanup.xls (Figure 42)

Figure 42 Groundwater cleanup spreadsheet model



This model shows the population risk (cell C25), which is the overall contamination risk to the people in the community as a function of the factors shown in Table 8, following:

Table 8 Groundwater Cleanup population risk factors

Risk factors	Cells	Description	Distribution
Cancer Potency	C18:C20	Cancer potency of each contaminant.	Lognormal
Concentration Before	D18:D20	Concentration of each contaminant before cleanup.	Normal
Volume Of Water Per Day	C23	Interindividual variability of volume of water consumed each day.	Normal, with lower bound of 0.
Body Weight	C22	Interindividual variability of body weights in the community.	Normal, with lower bound of 0.

Remediation costs of the various cleanup methods (cells E8:E10) are a function of factors shown in Table 9, following.

Table 9 Groundwater Cleanup remediation cost factors

Remediation cost factors	Cells	Description	Distribution
Fixed Costs	C8:C10	Flat costs for each method to pay for initial setup.	Triangular
Variable Costs	D8:D10	Costs for each method based on how long the cleanup takes.	Uniform
Efficiency	D14	Percent of contaminants that the cleanup process removes. Each remediation method has a different cost for different efficiency levels.	None

Groundwater Cleanup OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See [“Setting Crystal Ball Run Preferences”](#) on page 25.

► To run the optimization:

- 1 Be sure **Groundwater Cleanup.xls** is open in Crystal Ball.
- 2 In Crystal Ball, set the number of trials per simulation to 2000, since tail-end percentile requirements need more accuracy.
- 3 Start OptQuest.

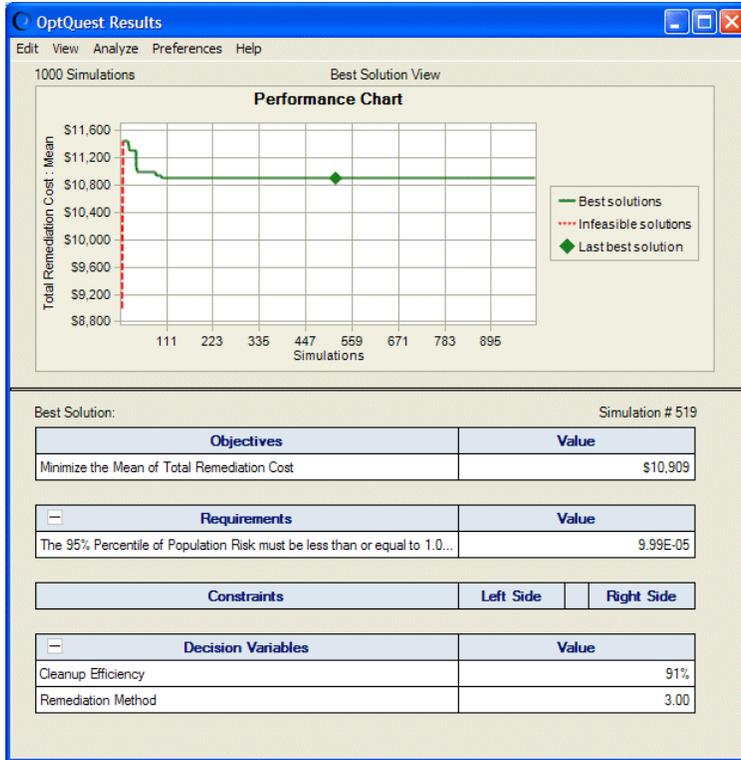
As you click OK to step through the problem, note:

- The objective is to minimize the remediation cost while requiring that the population risk be under 1E-4 with 95% certainty.
- There are two decision variables: Remediation Method (cell D13), and Cleanup Efficiency (cell D14). Notice that the Category type was chosen for Remediation Method since it acts as an "index" variable for selecting one of the methods.
- This problem has no constraints.

- 4 Run the optimization.

The results are shown in [Figure 43](#), following. The solution in [Figure 43](#) minimizes costs at \$10,909 while keeping the risk level at 9.99E-5, rounded.

Figure 43 Groundwater cleanup optimization results



The distributions for the total remediation cost and the population risk are shown in Figure 44 and Figure 45.

Figure 44 Groundwater cleanup total remediation cost forecast chart

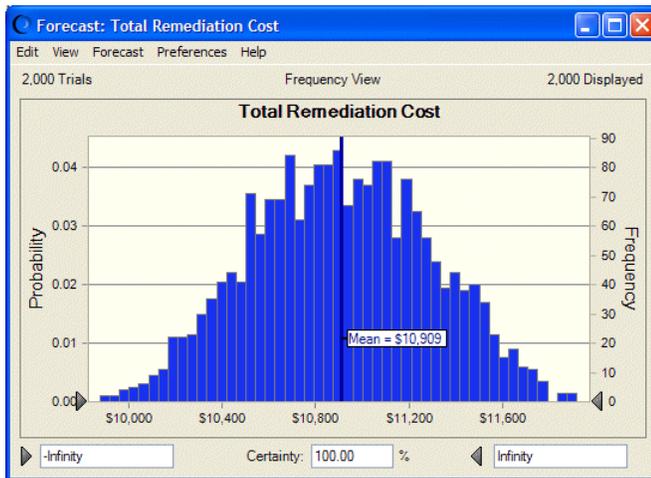
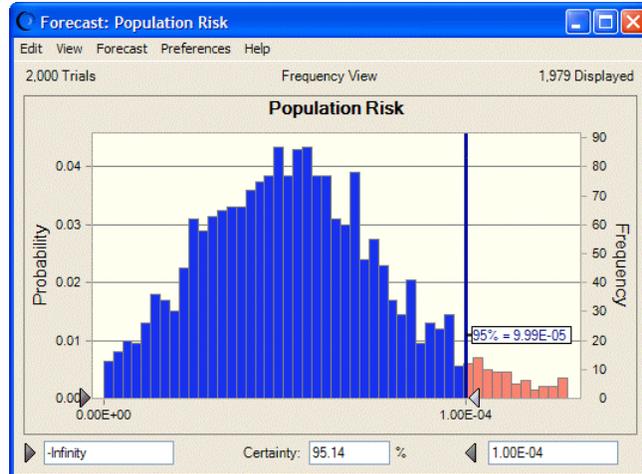


Figure 45 Groundwater cleanup population risk forecast chart



Oil Field Development

This example concerns an oil company analysis of Net Present Value for a new asset.

The following sections describe this problem and its OptQuest solution:

- “Oil Field Development Problem Statement” on page 83
- “Oil Field Development Spreadsheet Model” on page 84
- “Oil Field Development OptQuest Solution” on page 85

Oil Field Development Problem Statement

Oil companies need to assess new fields or prospects where very little hard data exists. Based on seismic data, analysts can estimate the probability distribution of the reserve size. With little actual data available, the discovery team wants to quantify and optimize the Net Present Value (NPV) of this asset. You can simplify this analysis by representing the production profile by three phases, shown in Table 10.

Table 10 Oil production phases

Phase	Description
Build up	The period when you drill wells to gain enough production to fill the facilities.
Plateau	After reaching the desired production rate (plateau), the period when you continue production at that rate as long as the reservoir pressure is constant and until you produce a certain fraction of the reserves. In the early stages of development, you can only estimate this fraction, and production above a certain rate influences plateau duration.
Decline	The period when production rates, P , decline by the same proportion in each time step, leading to an exponential function: $P(t) = P(0) \exp(-c*t)$

Phase	Description
	where t is the time since the plateau phase ended and c is some constant.

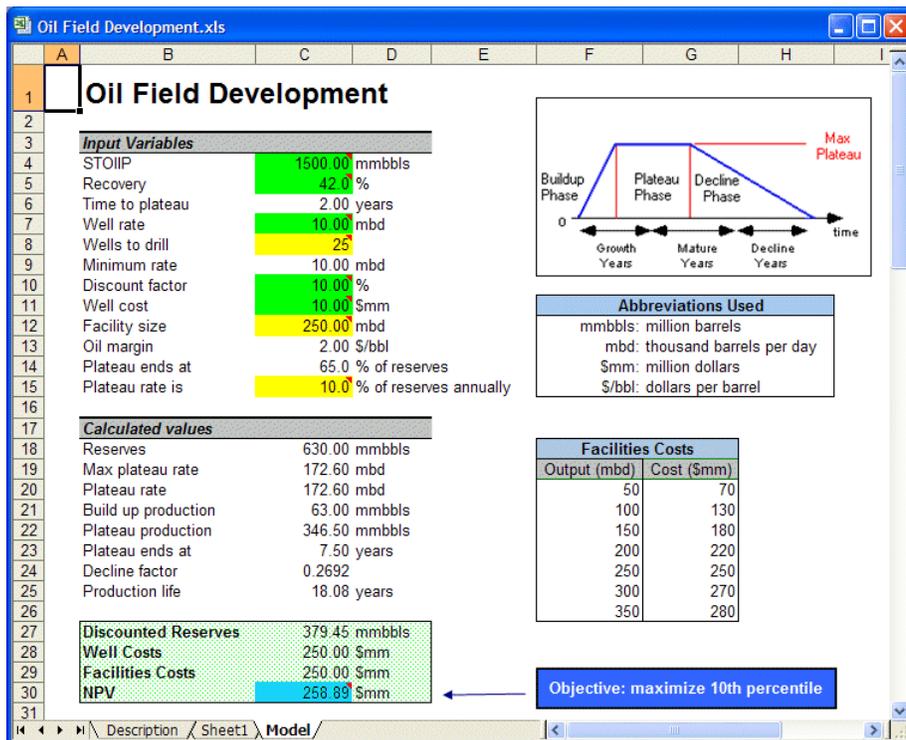
With only estimates for the total Stock Tank Oil Initially In Place (STOIIP = reserve size) and percent recovery amounts, the objective is to select a production rate, a facility size, and well numbers to maximize some financial measure. In this example the measure used is the 10th percentile (P90) of the NPV distribution. In other words the oil company wants to optimize an NPV value which they are 90% confident of achieving or exceeding.

As described, the problem is neither trivial nor overly complex. A high plateau rate doesn't lose any reserves, but it does increase costs with extra wells and larger facilities. However, facility costs per unit decrease with a larger throughput, so choosing the largest allowed rate and selecting a facility and number of wells to match might be appropriate.

Oil Field Development Spreadsheet Model

Open the Oil Field Development.xls workbook found in the Crystal Ball Example folder (Figure 46).

Figure 46 Oil field development problem spreadsheet model



Net present value (cell C30) of this oil field is based on:

- Total discounted reserves (cell C27)
- Oil margin (cell C13), which is equivalent to oil price minus operating costs

- Well costs (cell C28)
- Facilities cost (cell C29), which is determined for various production levels by a look-up table

Facility capacity places a maximum limit on production rate, while the production rate of the wells is defined as a normal distribution (cell C7).

The Production Profile table at the bottom of the model shows that the production phase determines annual production rates. Cumulative oil production is calculated per year and is then discounted at 10% (lognormal distribution in cell B10), resulting in a total discounted reserves value. The model gives an oil (or profit) margin of \$2.00 per barrel (bbl) and converts total discounted reserves to present value dollars. Total well and facilities costs are then subtracted for total project NPV.

Oil Field Development OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See [“Setting Crystal Ball Run Preferences” on page 25](#).

► Be sure Oil Field Development.xls is open in Crystal Ball. Then:

1 Start the OptQuest wizard.

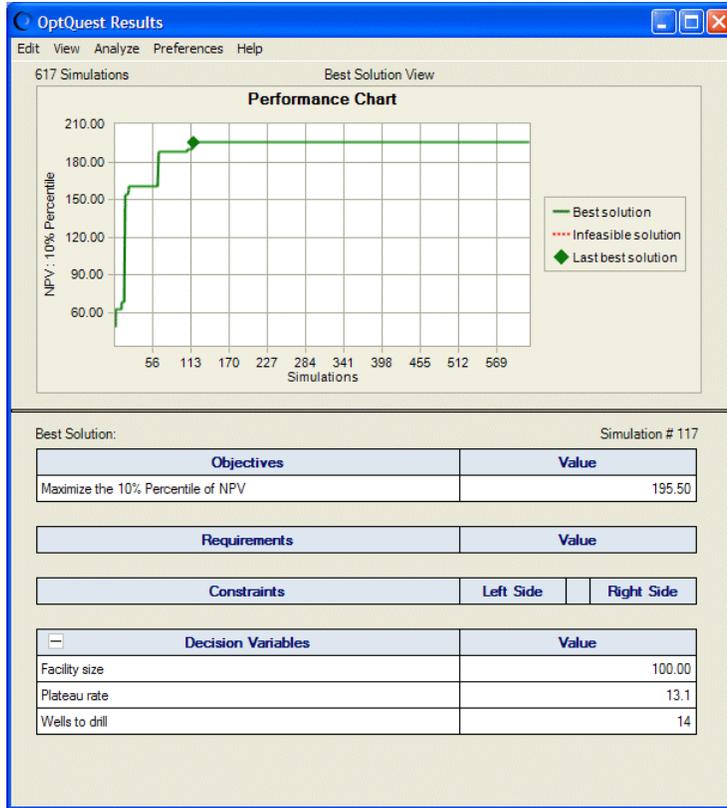
As you click OK to step through the problem, note:

- The objective is to maximize the 10th percentile (P90) of the NPV.
- There are three decision variables: Wells to drill (cell C8), Facility size (cell C12), and Plateau rate (cell C15).
- This problem has no constraints.

2 Run the optimization.

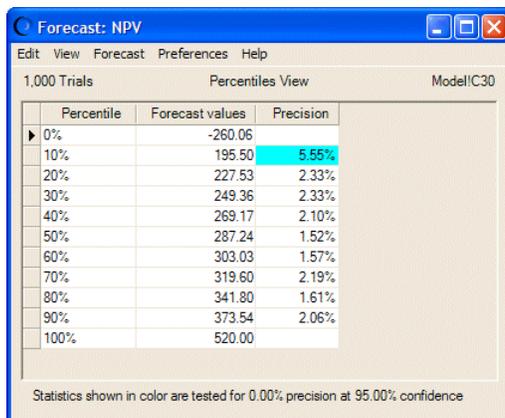
The results are shown in [Figure 47](#).

Figure 47 Oil field development optimization results



The Crystal Ball simulation of this solution in Figure 48 maximizes the 10th percentile (P90) of the NPV.

Figure 48 Oil field development solution (percentile view)



Portfolio Revisited

This example concerns analysis of an investment portfolio with respect to risk as well as return. The following sections describe this problem and two ways to solve it using OptQuest:

- [“Portfolio Revisited Problem Statement” on page 87](#)
- [“Portfolio Revisited Method 1: Efficient Frontier Optimization” on page 87](#)
- [“Portfolio Revisited Method 2: Multi-objective Optimization” on page 89](#)

Portfolio Revisited Problem Statement

The investor from [“Tutorial 2 — Portfolio Allocation Model” on page 54](#) has \$100,000 to invest in four assets. Below is a relisting of the investor’s expected annual returns, and the minimum and maximum amounts the investor is comfortable allocating to each investment.

Table 11 Sample investment requirements

Investment	Annual return	Lower bound	Upper bound
Money market fund	3%	\$0	\$50,000
Income fund	5%	\$10,000	\$25,000
Growth and income fund	7%	\$0	\$80,000
Aggressive growth fund	11%	\$10,000	\$100,000

When the investor maximized the portfolio return without regard to risk, OptQuest allocated almost all the money to the investment with the highest return. This strategy didn’t result in a portfolio that maintained risk at a manageable level. Only limiting the standard deviation of the total expected return generated a more diversified portfolio.

[“Efficient Frontier Analysis” on page 19](#) examines the reasons for this.

Portfolio Revisited Method 1: Efficient Frontier Optimization

OptQuest has a feature that creates an efficient frontier for you automatically. To use the Efficient Frontier function in OptQuest, you need only define a requirement with a variable upper or lower bound. OptQuest then calculates points within the variable requirement range.

The following sections describe the model for this solution method and its OptQuest solution:

- [“Efficient Frontier Spreadsheet Model” on page 87](#)
- [“Efficient Frontier OptQuest Solution” on page 88](#)

Efficient Frontier Spreadsheet Model

Open the Portfolio Revisited EF.xls workbook found in the Crystal Ball Examples folder. The total expected return forecast, assumptions, and decision variables are the same as in the original model, with the decision variables already defined.

Efficient Frontier OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See “Setting Crystal Ball Run Preferences” on page 25.

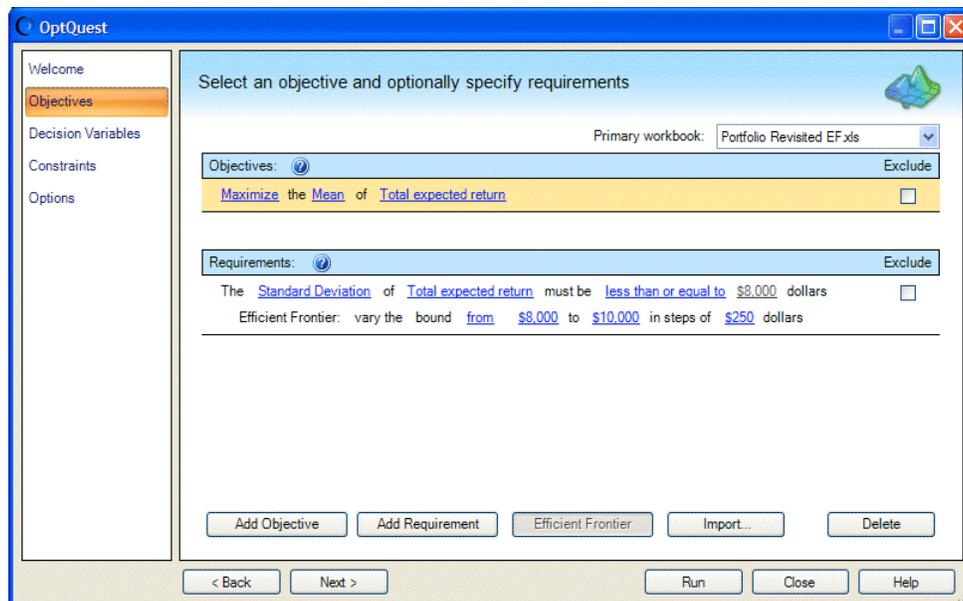
► Perform these steps:

- 1 With *Portfolio Revisited EF.xls* open in Crystal Ball, set the number of trials per simulation to 2000 in the Run Preferences dialog.
- 2 Start OptQuest from the Crystal Ball Run menu.

As you click Next to step through the problem, note that the objective, decision variables, and constraints are the same as for the original example (“Tutorial 2 — Portfolio Allocation Model” on page 54).

Figure 49 shows the Objectives panel with the variable requirement needed for efficient frontier testing.

Figure 49 Objectives panel with a variable requirement



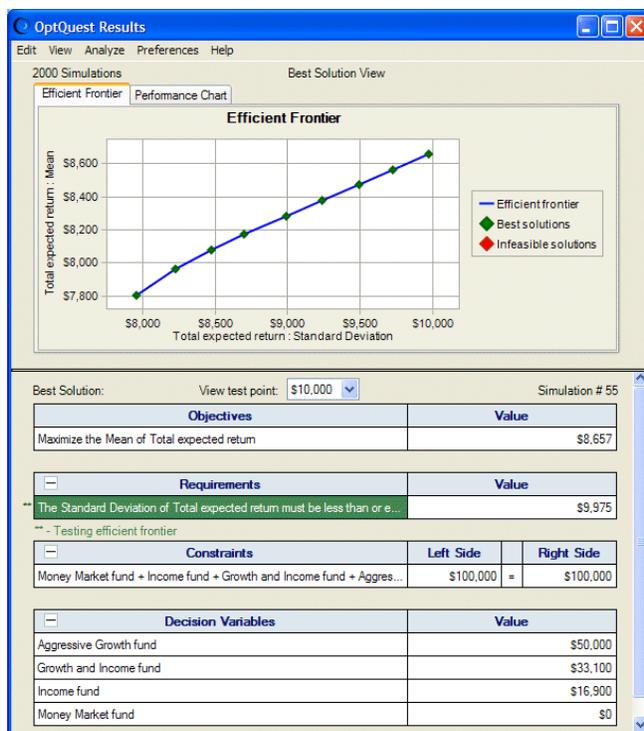
The requirement has a variable upper bound for the standard deviation statistic.

The variable requirement bounds are \$8,000 for the lower bound and \$10,000 for the upper bound.

- 3 Run the optimization for 2000 simulations.

The results are shown in Figure 50.

Figure 50 Portfolio Revisited Efficient Frontier optimization results



When should you use the Efficient Frontier function? This method is useful when it is difficult to determine reasonable lower or upper bounds for requirement statistics.

Portfolio Revisited Method 2: Multi-objective Optimization

Another technique for finding efficient portfolios is called multi-objective (or multi-criteria) optimization. You can use this technique to optimize multiple, often conflicting objectives, such as maximizing returns and minimizing risks, simultaneously. Other examples of multi-objective optimization include:

- Aircraft design, requiring simultaneous optimization of weight, payload capacity, airframe stiffness, and fuel efficiency
- Public health policies, requiring simultaneous minimization of risks to the population, direct taxpayer costs, and indirect business regulation costs
- Electric power generation, requiring simultaneous optimization of operating costs, reliability, and pollution control

Most forms of multi-objective optimization are solved by minimizing or maximizing a weighted combination of the multiple objectives. In the portfolio example, a weighted combination of the return and risk objectives might be:

$$\text{mean return} - (k * \text{standard deviation})$$

where $k > 0$ is a risk aversion constant, and the objective is to maximize the function. The relationship between return and risk for the investor is captured entirely by this one function; no additional requirements are necessary.

Geometrically, the optimal solution for a multi-objective function occurs in the saddle point between the optimal endpoints of the individual objectives. In the case of the two-objective function above, the optimal solution occurs somewhere on the efficient frontier between the maximum-return portfolio and the minimum-risk portfolio.

For $k = 0.5$, the optimal solution occurs at the point where the return minus one-half the standard deviation has the highest value.

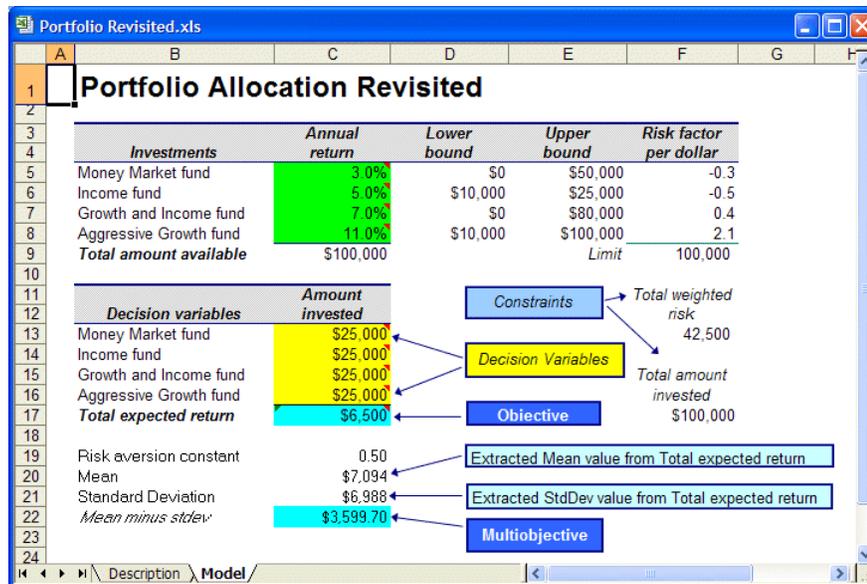
The following sections describe the model for this problem and its OptQuest solution:

- “Multi-objective Optimization Spreadsheet Model ” on page 90
- “Multi-objective Optimization OptQuest Solution” on page 90

Multi-objective Optimization Spreadsheet Model

Open the Portfolio Revisited.xls workbook found in the Crystal Ball Examples folder. The total expected return forecast, assumptions, and decision variables are the same as in the original model. Scroll down to see the new items added as shown in Figure 51.

Figure 51 Portfolio Revisited Spreadsheet Model



This new function (cell C22) contains the multi-objective relationship described by mean return – ($k * \text{standard deviation}$) with the risk aversion constant (cell C19) broken out into a separate cell. The mean return and standard deviation variables in this equation are automatically extracted at the end of the simulation from the Total Expected Return forecast (cell C17). See the *Oracle Crystal Ball User's Guide* for more information on the Auto Extract feature.

Multi-objective Optimization OptQuest Solution

To follow this example:

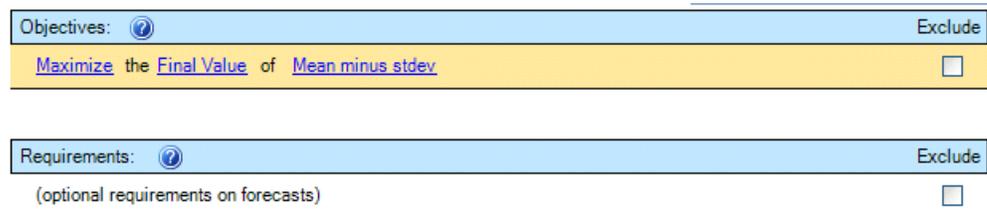
1. Open Portfolio Revisited.xls in Crystal Ball.

This example uses the Crystal Ball run preferences recommended in “[Setting Crystal Ball Run Preferences](#)” on page 25.

2. Start the OptQuest wizard.

As you click OK to step through the problem, note:

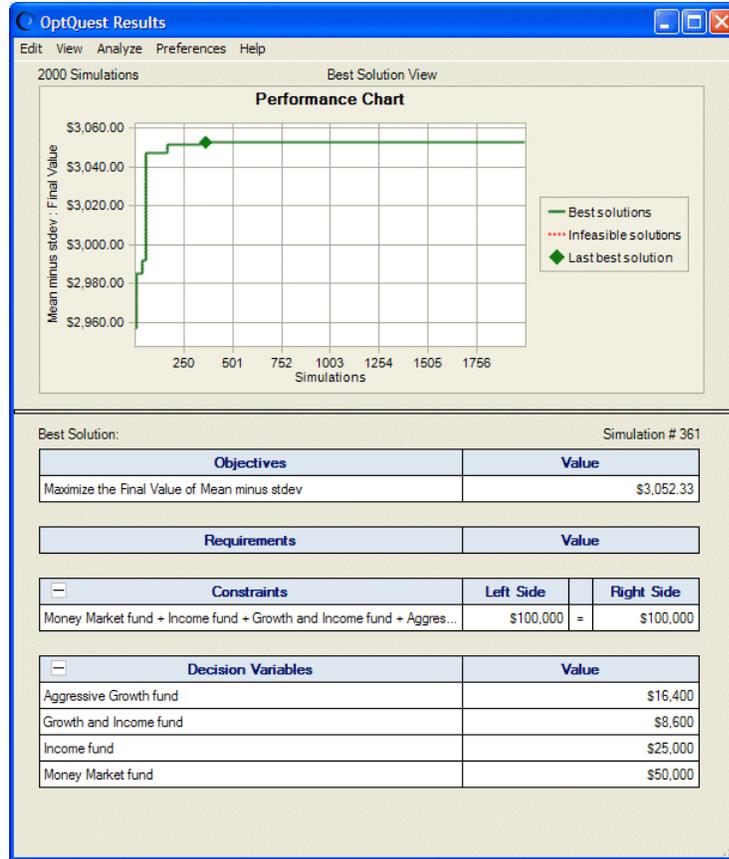
- The objective refers to the new multi-objective function. The statistic to optimize is Final Value, to calculate only the statistical values for the total expected return forecast at the end of the simulation. The following figure illustrates the objective and shows that there are no requirements:



- The decision variables and constraints are the same as previous Portfolio Allocation examples.
3. Run the optimization for 2,000 simulations.

The results appear in [Figure 52](#). After reviewing the results, close Portfolio Revisited.xls without saving it.

Figure 52 Portfolio Revisited Multi-objective Optimization Results



When should you use multi-objective optimization, and when should you use single objectives with requirements? The former method is especially useful when it is difficult to determine reasonable lower or upper bounds for requirement statistics. This method is also recommended for situations where OptQuest has trouble finding feasible solutions that satisfy many requirements. The latter method is generally easier to implement and understand.

Tolerance Analysis

An engineer at an automobile design center needs to specify components for piston and cylinder assemblies that work well together. To do this, he needs the dimensions of the components to be within certain tolerance limits, while still choosing the most cost-efficient methods. This is called an optimal stack tolerance analysis.

The following sections describe this problem and its OptQuest solution:

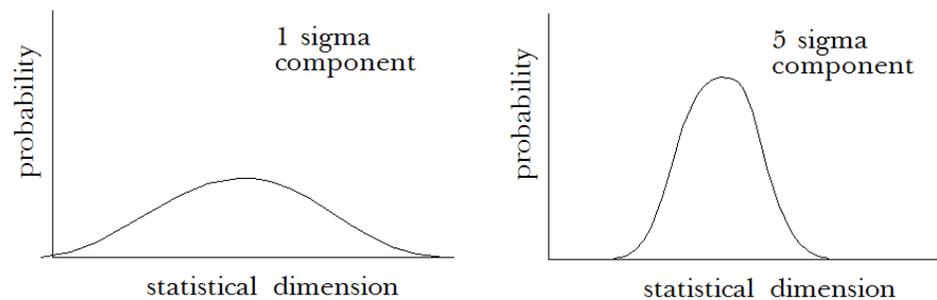
- [“Tolerance Analysis Problem Statement” on page 93](#)
- [“Tolerance Analysis Spreadsheet Model” on page 93](#)
- [“Tolerance Analysis OptQuest Solution” on page 95](#)

Note: This example involves concepts used only by Six Sigma and similar quality programs. If you are not familiar with Crystal Ball's process capability features, consider reviewing the process capability appendix in the *Oracle Crystal Ball User's Guide*.

Tolerance Analysis Problem Statement

The piston assembly consists of five components, and the cylinder assembly consists of two, each with certain nominal dimensions. These components are then stacked to create the assembly. The difference in length between the two, called the assembly gap, must be between 0.003 and 0.02 inches. This might seem like a simple problem, but since milling processes are not exact and quality control has a direct effect on prices, components have an error associated with each, called tolerance. When stacked, these errors compile or add together to create a cumulative tolerance.

When a batch of components is milled and measured, the components' actual dimensions form a distribution around the desired, or nominal, dimension. Standard deviation, or sigma, is a measure of the variation present in a batch of components. The components then have a statistical dimension based on this distribution. The quality of the component and the associated tolerance is described in terms of sigmas, with 1 sigma component having the largest tolerance and a 5 sigma component the smallest. This is called the quality specification.



One simplified solution takes the total tolerance allowed and divides it by the number of components. But, due to individual component complexity and process differences in manufacturing, each component of the assembly has a different cost function associated with the quality specification. This then becomes a juggling act to balance cumulative tolerance and associated cost.

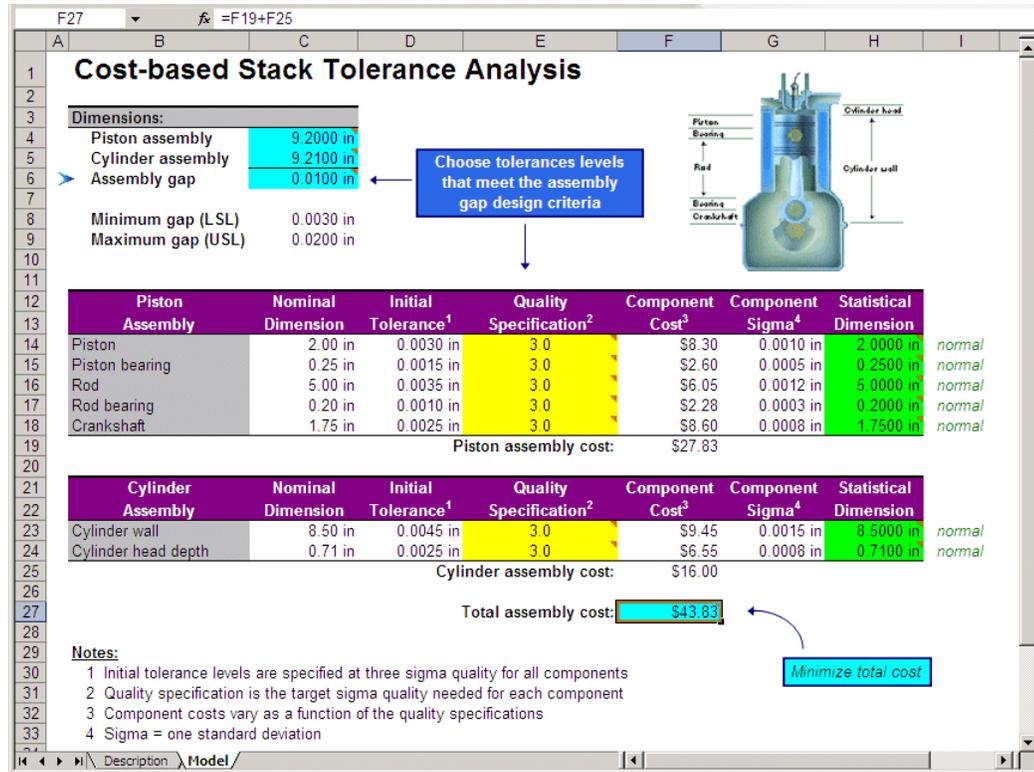
Crystal Ball supports quality programs such as Six Sigma by calculating a set of process capability metrics for forecasts when the process capability features are activated and at least one specification limit (LSL or USL) is entered for the forecasts. OptQuest then includes these metrics in the list of statistics that can be optimized. For more information, see [“OptQuest and Process Capability” on page 21](#).

This example assumes that the process capability metrics have been activated in Crystal Ball. Then, the capability metrics are available in the Forecast Statistic list of the Objectives panel.

Tolerance Analysis Spreadsheet Model

Open the Tolerance Analysis.xls file ([Figure 53](#)).

Figure 53 Tolerance analysis spreadsheet model



A drawing of the assembly is in the upper right corner. In this example:

- The nominal dimensions are in cells C14:C18 and C23:C24.
- Initial tolerances of each 3-sigma component are in cells D14:D18 and D23:D24.
- The relationship between the initial tolerance and the quality specifications (cells E14:E18 and E23:E24) yields a component sigma (cells G14:G18 and G23:G24).
- The statistical dimension (cells H14:H18 and H23:H24) of each component is defined as an assumption with a normal distribution having a mean equal to the nominal dimension and a standard deviation equal to the component sigma. Note that the mean and standard deviation are cell references to these cells.

The dimensions of the assemblies are a cumulation of their respective components' statistical dimensions. The difference in length between the cylinder assembly (cell C5) and the piston assembly (cell C4) is the assembly gap (cell C6).

Component cost (cells F14:F18 and F23:F24) is a nonlinear function of quality specification. The higher the specification, the higher the cost. Also note that each component has a different cost function associated with it.

In addition to the recommended options, before running OptQuest, in Crystal Ball select Run, Run Preferences and set:

- The maximum number of trials run to 2000
- The sampling method to Latin Hypercube
- The sample size to 2000 for Latin Hypercube

Since the model is heavily dependent on the tails of the forecast distribution, these settings will provide higher accuracy and will be adequate for this example. In actual practice, to gain better accuracy, the engineer might want to run longer simulations of 5000 or 10,000 trials.

Tolerance Analysis OptQuest Solution

The goal of the following solution is to maximize quality while minimizing cost.

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See “Setting Crystal Ball Run Preferences” on page 25.

- ▶ To run the optimization:
 - 1 Be sure Tolerance Analysis.xls is open in Crystal Ball and the maximum trials and sample sizes have been set to 2000 as described above.
 - 2 Start OptQuest.

As you click Next to step through the problem, note:

- The objective is to minimize the mean of total assembly cost. Note that the total cost function does not depend on any assumption cells, and thus has a deterministic value. You can use the final value statistic in these cases to retrieve the deterministic value.
 - Two requirements ensure that the assembly gap is between 0.003 and 0.02 inches.
 - This problem has seven decision variables, one for the quality specification for each assembly component, with a continuous range between 1 and 5 sigmas.
 - The problem has no constraints.
- 3 Return to the Objectives panel and enter the objective and requirements shown in Figure 54.

Figure 54 Objectives settings to optimize quality and cost

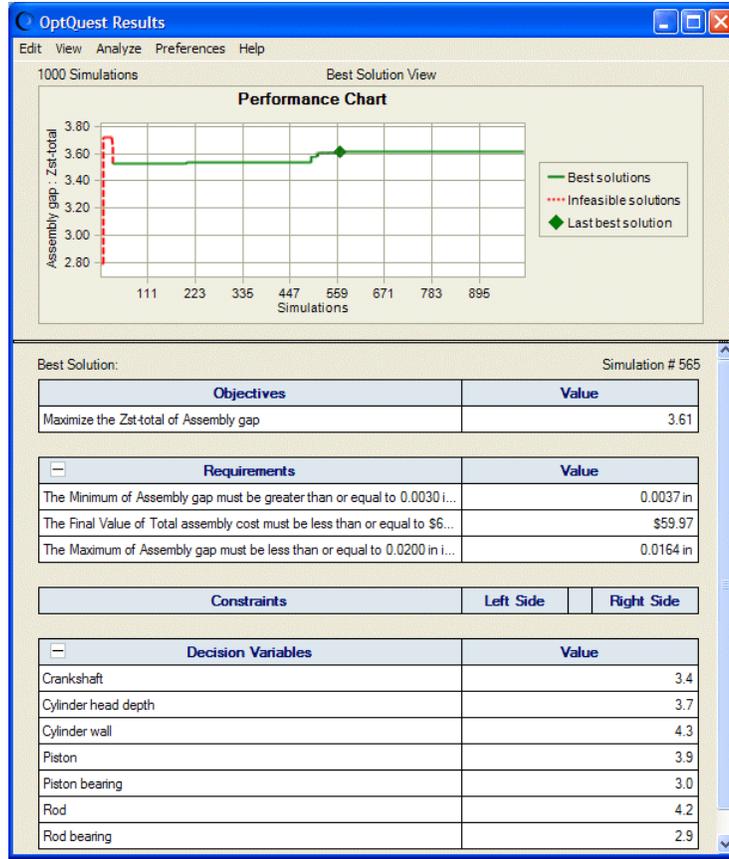
Objectives:	Exclude
Maximize the Zst-total of Assembly gap	<input type="checkbox"/>

Requirements:	Exclude
The Minimum of Assembly gap must be greater than or equal to 0.0030 in inches	<input type="checkbox"/>
The Maximum of Assembly gap must be less than or equal to 0.0200 in inches	<input type="checkbox"/>
The Final Value of Total assembly cost must be less than or equal to \$60.00 dollars	<input type="checkbox"/>

- 4 Run the optimization.

The cost and quality solution values appear in Figure 55.

Figure 55 OptQuest solution for maximum quality with a cost requirement



Inventory System Optimization

This example is adapted from James R. Evans and David L. Olson, *Introduction to Simulation and Risk Analysis*. New York: Prentice-Hall, 1998.

The following sections describe this problem and its OptQuest solution:

- “Inventory System Problem Statement” on page 96
- “Inventory System Spreadsheet Model” on page 98
- “Inventory System OptQuest Solution” on page 99

Inventory System Problem Statement

The two basic inventory decisions that managers face are:

- How much additional inventory to order or produce
- When to order or produce it

Although it is possible to consider these two decisions separately, they are so closely related that a simultaneous solution is usually necessary. Typically, the objective is to minimize total

inventory costs. Total inventory costs typically include holding, ordering, shortage, and purchasing costs.

In a continuous review system, managers continuously monitor the inventory position. Whenever the inventory position falls at or below a level R , called the reorder point, the manager orders Q units, called the order quantity. (Note that the reorder decision is based on the inventory position including orders and not the inventory level. If managers used the inventory level, they would place orders continuously as the inventory level fell below R until they received the order.) When you receive the order after the lead-time, the inventory level jumps from zero to Q , and the cycle repeats.

In inventory systems, demand is usually uncertain, and the lead-time can also vary. To avoid shortages, managers often maintain a safety stock. In such situations, it is not clear what order quantities and reorder points will minimize expected total inventory cost. Simulation models can address this question.

In this example, demand is uncertain and is Poisson distributed with a mean of 100 units per week. Thus, the expected annual demand is 5,200 units.

Note: For large values of the rate parameter, λ , the Poisson distribution is approximately normal. Thus, this assumption is tantamount to saying that the demand is normally distributed with a mean of 100 and standard deviation of $\sqrt{100} = 10$. The Poisson distribution is discrete, thus eliminating the need to round off normally distributed random variates

Additional relationships that hold for the inventory system are:

- Each order costs \$50 and the holding cost is \$0.20 per unit per week (\$10.40 for one year).
- Every unfilled demand is lost and costs the firm \$100 in lost profit.
- The time between placing an order and receiving the order is 2 weeks. Therefore, the expected demand during lead-time is 200 units. Orders are placed at the end of the week, and received at the beginning of the week.

The traditional economic order quantity (EOQ) model suggests an order quantity:

$$Q = \sqrt{\frac{2 \times 5200 \times 50}{10.4}} = 224$$

For the EOQ policy, the reorder point should equal the lead-time demand; that is, place an order when the inventory position falls to 200 units. If the lead-time demand is exactly 200 units, the order will arrive when the inventory level reaches zero.

However, if demand fluctuates about a mean of 200 units, shortages will occur approximately half the time. Because of the high shortage costs, the manager would use either a larger reorder point, a larger order quantity, or both. In either case, the manager will carry more inventory on average, which will result in a lower total shortage cost but a higher total holding cost. A higher order quantity lets the manager order less frequently, thus incurring lower total ordering costs. However, the appropriate choice is not clear. Simulation can test various reorder point/order quantity policies.

Inventory System Spreadsheet Model

Before examining the spreadsheet simulation model, step through the logic of how this inventory system operates. Assume that no orders are outstanding initially and that the initial inventory level is equal to the order quantity, Q . Therefore, the beginning inventory position will be the same as the inventory level. At the beginning of the week, if any outstanding orders have arrived, the manager adds the order quantity to the current inventory level.

Next, determine the weekly demand and check if sufficient inventory is on hand to meet this demand. If not, then the number of lost sales is the demand minus the current inventory. Subtract the current inventory level from the inventory position, set current inventory to zero, and compute the lost sales cost. If sufficient inventory is available, satisfy all demand from stock and reduce both the inventory level and inventory position by the amount of demand.

The next step is to check if the inventory position is at or below the reorder point. If so, place an order for Q units and compute the order cost. The inventory position is increased by Q , but the inventory level remains the same. Schedule a receipt of Q units to arrive after the lead-time.

Finally, compute the holding cost based on the inventory level at the end of the week (after demand is satisfied) and the total cost.

Open the file Inventory System.xls. This spreadsheet model, shown in Figure 56, implements this logic. The basic problem data are shown in the upper left corner. The decision variables are the order quantity (cell E3) and the reorder point (cell E4). The initial inventory is set equal to the chosen order quantity. This example assumes the specified lead-time is constant.

Figure 56 Inventory system problem spreadsheet model

The screenshot shows a spreadsheet titled 'Inventory System.xls' with the following data:

Parameter	Value	Unit	Cost
Order Quantity	250	units	\$ 50
Reorder Point	250	units	\$ 0.20
Initial Inventory	250	units	\$ 100
Lead time	2	weeks	
Total Annual Costs			
	\$1,040	\$1,050	\$5,000
			\$ 7,090

Week	Beg Inv Pos	Beg Inv	Order Rec'd	Units Rec'd	Dmd	End Inv	Lost Sales	Order Placed?	Ending Inv Pos	Week Due	Hold Cost	Order Cost	Short Cost	Total Cost
1	250	250		0	100	150	0	TRUE	400	4	\$30.00	\$ 50	\$ -	\$ 80
2	400	150		0	100	50	0	FALSE	300		\$10.00	\$ -	\$ -	\$ 10
3	300	50	FALSE	0	100	0	50	TRUE	500	6	\$ -	\$ 50	\$5,000	\$ 5,050
4	500	0	TRUE	250	100	150	0	FALSE	400		\$30.00	\$ -	\$ -	\$ 30
5	400	150	FALSE	0	100	50	0	FALSE	300		\$10.00	\$ -	\$ -	\$ 10
6	300	50	TRUE	250	100	200	0	TRUE	450	9	\$40.00	\$ 50	\$ -	\$ 90
7	450	200	FALSE	0	100	100	0	FALSE	350		\$20.00	\$ -	\$ -	\$ 20
8	350	100	FALSE	0	100	0	0	TRUE	500	11	\$ -	\$ 50	\$ -	\$ 50
9	500	0	TRUE	250	100	150	0	FALSE	400		\$30.00	\$ -	\$ -	\$ 30
10	400	150	FALSE	0	100	50	0	FALSE	300		\$10.00	\$ -	\$ -	\$ 10
11	300	50	TRUE	250	100	200	0	TRUE	450	14	\$40.00	\$ 50	\$ -	\$ 90
12	450	200	FALSE	0	100	100	0	FALSE	350		\$20.00	\$ -	\$ -	\$ 20

In the actual simulation, the beginning inventory position and inventory level for each week equals the ending levels for the previous week, except for the first week, which is specified in the problem data. The demand is in column F as Crystal Ball assumptions.

Since all shortages are lost sales, the inventory level cannot be negative. Thus, the ending inventory each week is:

$$\text{ending inventory} = \max \left\{ \begin{array}{l} \text{beginning inventory level} - \text{demand} + \text{orders received} \\ 0 \end{array} \right\}$$

Lost sales are computed by checking if demand exceeds available stock and computing the difference.

The spreadsheet simulates 52 weeks, or one year of operation of the inventory system. Since the objective is to minimize the mean total annual cost, cell O6 is defined as a forecast cell.

Column I determines whether the manager should place an order by checking if the beginning inventory position minus the weekly demand is at or below the reorder point. The ending inventory position is:

$$\begin{array}{l} \text{ending} \\ \text{inventory} \\ \text{position} \end{array} = \begin{array}{l} \text{beginning} \\ \text{inventory} \\ \text{position} \end{array} - \text{weekly demand} + \text{lost sales} + \text{weekly orders}$$

This formula might not appear to be obvious. Clearly, if there are no lost sales, the ending inventory position is simply the beginning position minus the demand plus any order that may have been placed. If lost sales occur, computing the ending inventory position this way reduces it by the unfulfilled demand, which is incorrect. Thus, you must add back the number of lost sales to account for this.

In the ordering process, the manager places orders at the end of the week and receives orders at the beginning of the week. Thus, in [Figure 56](#), the order placed at the end of the first week with a lead-time of 2 weeks will arrive at the beginning of the fourth week. Column K determines the week an order is due to arrive, and a MATCH function is used in column D to identify whether an order is scheduled to arrive.

Inventory System OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See [“Setting Crystal Ball Run Preferences” on page 25](#).

► With Inventory System.xls open in Crystal Ball:

1 Start the OptQuest wizard.

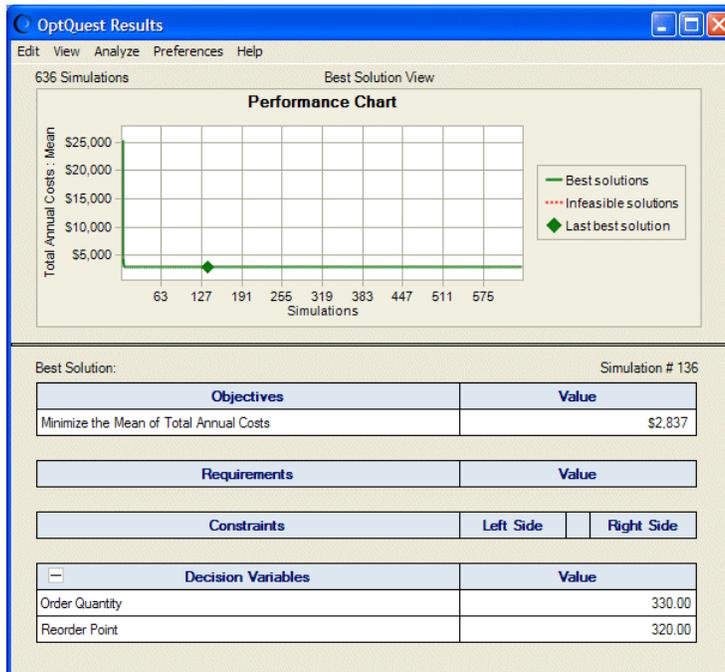
As you click Next to step through the problem, note:

- The objective is to minimize the total annual costs.
- There are no constraints or requirements.
- This problem has two decision variables.
- The initial search limits are set between 200 and 400 for both variables using a Discrete decision variable type.
- This optimization runs more slowly than some. You might want to run fewer than 1,000 simulations or use the Advanced Options settings to automatically stop the optimization when certain criteria are met (see [“Advanced Options” on page 35](#)).

2 Run the optimization.

Figure 57, following, shows optimization results. OptQuest identified the best solution as having an order quantity of 330 and a reorder point of 320. The Performance Chart shows that OptQuest quickly found a good solution value.

Figure 57 Inventory system model optimization results



Because this optimization used a step size of 5, you can fine-tune the solution by searching more closely around the best solution using a smaller step size while also increasing the number of trials per simulation for better precision. This is a good practice, since choosing too small a step size initially consumes a lot of time or, if time is restricted, OptQuest might not find a good solution. Thus, as the number of decision variables and range of search increases, use larger step sizes and fewer trials initially. Later, refine the search around good candidates.

Figure 58 shows the results of an optimization with Order Quantity and Reorder Point bounded to the range 300 to 360, with a step size of 1, and 1000 trials per simulation. OptQuest identified the best solution as Order Quantity = 332 and Reorder Point = 325. There was very little change from the initial solution.

Figure 58 Inventory system—second optimization results

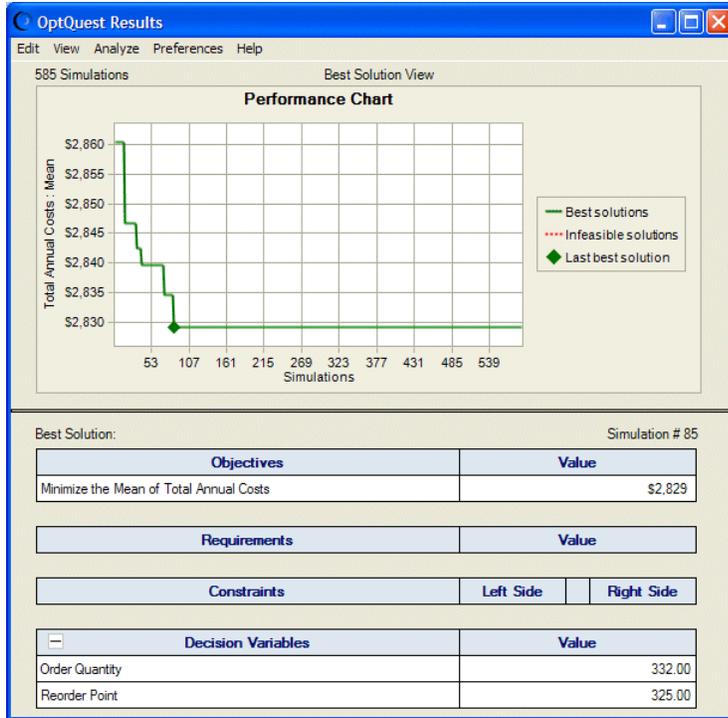


Figure 59 shows the Crystal Ball forecast chart for the annual total costs. You can see that the distribution of total annual cost is highly concentrated around the mean, but is also skewed far to the right, indicating that very high values of cost are possible, although not very likely. For such highly skewed distributions, run more trials than usual, since statistics like the mean and tail-end percentiles can be susceptible to extreme outliers.

Figure 59 Inventory system final (best) solution forecast chart



Drill Bit Replacement Policy

This example was suggested from an example in Kenneth K. Humphreys, *Jelen's Cost and Optimization Engineering*. 3rd ed. New York: McGraw-Hill, 1991. 257-262.

The following sections describe this problem and its OptQuest solution:

- “Drill Bit Replacement Problem Statement” on page 102
- “Drill Bit Replacement Spreadsheet Model” on page 103
- “Drill Bit Replacement OptQuest Solution” on page 104

Drill Bit Replacement Problem Statement

When drilling wells in certain types of terrain, the performance of a drill bit erodes with time because of wear. After T hours, the drilling rate can be expressed as:

$$\frac{dM}{dT} = \frac{15}{\sqrt{T/10}} \quad \text{meters per hour}$$

For example, after 5 hours of consecutive use (starting with a new drill bit), the drill is able to penetrate the terrain at a rate of:

$$\frac{15}{\sqrt{5/10}} = 21.21 \quad \text{meters per hour}$$

While after 50 hours, the penetration rate is only:

$$\frac{15}{\sqrt{50/10}} = 6.71 \quad \text{meters per hour}$$

Eventually, the bit must be replaced as the costs exceed the value of the well being drilled. The problem is to determine the optimum replacement policy; that is, the drilling cycle, T hours, between replacements.

T hours after replacing the bit, the total drilled depth in meters, M , is given by the integral of Equation 4.2 from 0 to T , or:

$$M = 300\sqrt{T/10} \quad \text{meters}$$

where 300 is a drilling depth coefficient.

The revenue value per meter drilled is calculated to be \$60. Drilling expenses are fixed at \$425 per hour, and it generally requires $R = 7.5$ hours to install a new drill bit, at a cost of \$8,000 + \$400 R .

If all drilling parameters were certain, calculating the optimal replacement policy would be straightforward. However, several of the drilling parameters are uncertain, and knowledge about their values must be assumed:

- Because of variations in the drilling process and terrain, the depth coefficient, C , is characterized by a normal distribution with a mean of 300 and a standard deviation of 20.

- The drill bit replacement time, R , varies and is determined by a triangular distribution with parameters 6.5, 7.5, and 9.
- The number of 10-hour days available per month, D , also varies due to the weather and the number of days in a month, and is assumed to be triangular with parameters 24, 28, and 30.

With these assumptions, the profit/drilling cycle if the bit is replaced after T hours equals the revenue obtained from drilling minus drilling expenses and replacement costs:

$$\text{profit/drilling cycle} = \$60M - \$425T - (\$8,000 + \$400R)$$

Assuming D ten-hour days per month, the average number of cycles per month is $10D/(T + R)$. Therefore, the average profit per month is:

$$\frac{\text{average profit}}{\text{month}} = \frac{10D \left[\$60 \left(C \sqrt{\frac{T}{10}} \right) - \$425T - \$8000 - \$400R \right]}{T + R}$$

The objective is to find the value of T that maximizes the average profit per month.

Drill Bit Replacement Spreadsheet Model

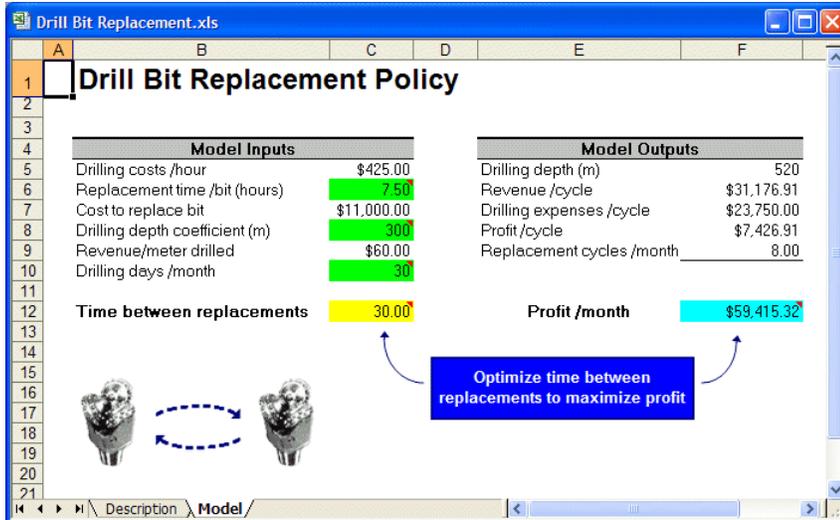
Open the Drill Bit Replacement example, shown in [Figure 60](#), below. This workbook has Crystal Ball assumptions defined for:

Table 12 Drill Bit Replacement model assumptions

Cell	Assumption
C6	Replacement time, R .
C8	Drilling depth function coefficient, C .
C10	Number of days available per month, D .

One decision variable is defined in cell C12: the cycle time between replacements of the drill bit, T .

Figure 60 Drill bit replacement problem spreadsheet model



The model outputs are computed using the formulas developed in the previous section. The drilling expenses in cell F7 include both the drilling costs and the replacement costs. The forecast cell is F12, profit per month.

Drill Bit Replacement OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See “Setting Crystal Ball Run Preferences” on page 25.

➤ With Drill Bit Replacement.xls open in Crystal Ball:

1 Start the OptQuest wizard.

As you click Next to step through the problem, note:

- The objective is to maximize the mean profit per month.
- The problem has no requirements or constraints.
- This problem has one decision variable, whose search limits are 1 and 50.

2 Run the optimization.

Figure 61, following, shows the OptQuest results. The best solution is to replace the drill bit approximately every 19.9 hours.

Figure 61 Drill bit replacement model optimization results

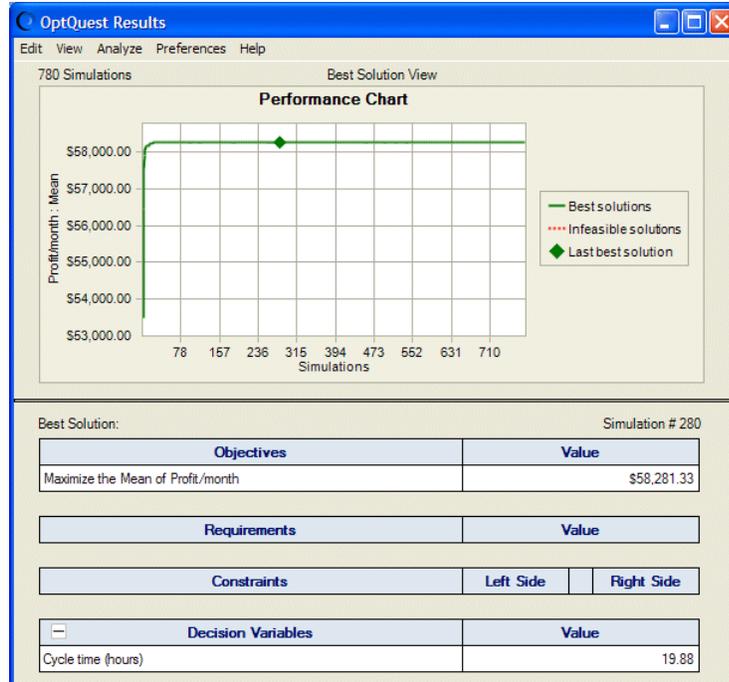
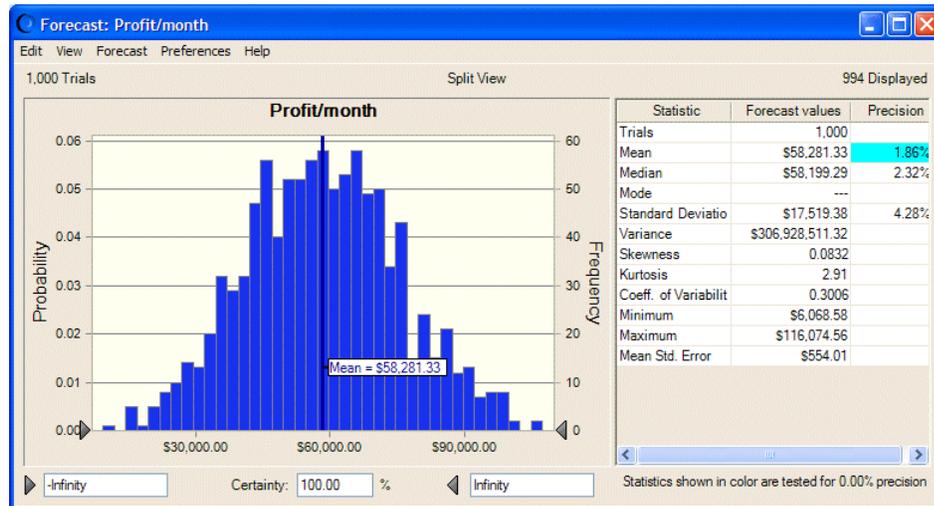


Figure 62, following, shows the OptQuest forecast chart and statistics for the simulation of this solution. The profit per month has a relatively large standard deviation compared to the mean (coefficient of variability=0.30); thus, it is likely that the true profit/month is significantly higher or lower than the mean objective value.

Figure 62 Drill bit replacement forecast chart and statistics



Gasoline Supply Chain

This example shows how to determine the optimum amount of gasoline to transport between different levels of a gasoline supply chain. The objective is to minimize the total cost, which

includes transportation costs and inventory holding costs at various points in the supply chain. It is also important to minimize stockouts at various retail outlets. The complexity of the problem arises from the fact that there is stochastic (variable) production at the refinery level and stochastic demand at the retail outlet level.

The following sections describe this problem and its OptQuest solution:

- [“Gasoline Supply Chain Statement of Problem” on page 106](#)
- [“Gasoline Supply Chain Spreadsheet Model” on page 107](#)
- [“Gasoline Supply Chain OptQuest Solution” on page 108](#)

Gasoline Supply Chain Statement of Problem

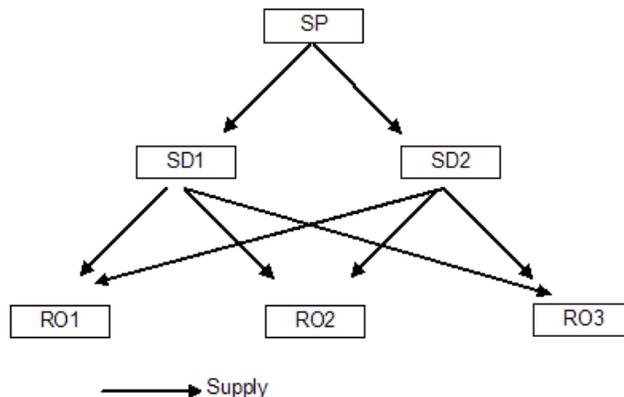
The supply chain illustrated here is simplified. It consists of one refinery (SP), two supply depots (SD), and three retail outlets (RO).

A weekly snapshot of this supply chain is as follows:

- The refinery produces a variable amount of gasoline every week, which it transports to SDs for cross-docking.
- SDs supply gasoline to ROs, which realize stochastic demand from end customers.
- All three supply chain levels (Refinery, SD, and RO) face inventory holding costs.
- In addition, the RO's face the risk of stockouts for not fulfilling customer demands.

The problem is to determine the amount of gasoline to transport between each level of the supply chain to minimize the total operating cost, which is computed as the sum of transportation costs and inventory holding costs. For business reasons, it is helpful to minimize stockouts at the ROs, to a certain extent.

The following is a schematic diagram of the supply chain:



Assumptions about the supply chain are as follows:

- The weekly supply of gasoline from the refinery (SP) follows a normal distribution with a mean of 2000 gallons and standard deviation (s.d.) of 450 gallons.

- The weekly demands at ROs are distributed lognormally with means and standard deviations of 400 gallons and 50 gallons, 500 gallons and 75 gallons, 650 gallons and 100 gallons respectively at RO1, RO2, RO3.
- The inventory holding cost is a dollar for every five gallons.
- The transportation costs in dollars per gallon are as follows (note that these costs include transportation distances):

SP to SD1 = \$15

SP to SD2 = \$12.5

SD1 to RO1 = \$6.5

SD1 to RO2 = \$7.5

SD1 to RO3 = \$9.0

SD2 to RO1 = \$9.0

SD2 to RO2 = \$8.0

SD2 to RO3 = \$7.0

- Existing inventories in gallons are:
Refinery: 200 gallons, SD1: 50 gallons, SD2: 100 gallons, RO1: 120 gallons, RO2: 180 gallons, RO3: 80 gallons.

Other assumptions include:

- There is no capacity limit on transportation links and supply chain points.
- There is an implicit constraint that the SDs do not have any stockouts. This mathematically implies that:

$$\text{Existing Inventory} + \text{Supply Received} - \text{Demand Fulfilled} \geq 0$$

Gasoline Supply Chain Spreadsheet Model

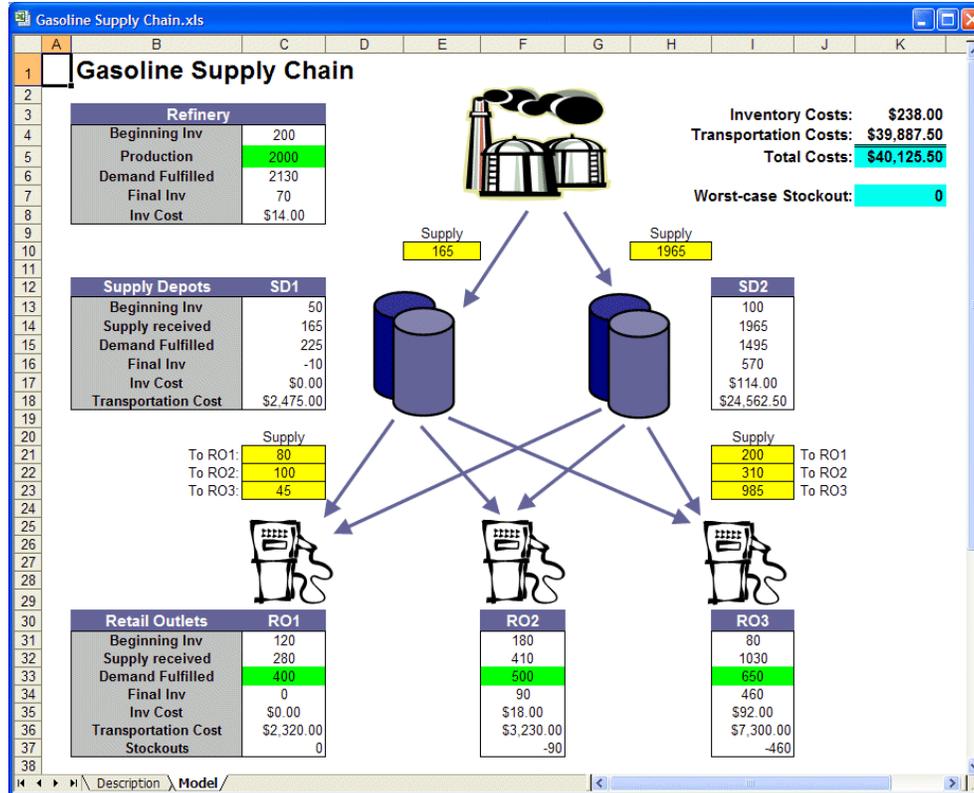
Open the spreadsheet model for this example, Gasoline Supply Chain.xls, as shown in [Figure 63](#).

This model includes:

- Four Crystal Ball assumptions in cells C5, C33, F33, and I33. These represent stochastic output from the refinery and stochastic demand at the retail outlets.
- Two Crystal Ball forecasts in cells K5 and K7 to represent total costs and the worst-case stockout situation.
- Eight decision variables that represent transportation costs from the refinery to the two supply depots and from each depot to each retail outlet. These appear in cells E10, H10, C21, C22, C23, I21, I22, and I23.

For this example, OptQuest can determine how much to supply at each of the SDs and ROs to minimize the total expected cost while maintaining stockouts at ROs at an acceptable level.

Figure 63 Gasoline supply chain spreadsheet model



Gasoline Supply Chain OptQuest Solution

Note: Except where indicated, this example uses the recommended Crystal Ball run preferences. See “Setting Crystal Ball Run Preferences” on page 25.

➤ With Gasoline Supply Chain.xls open in Crystal Ball:

1 Start the OptQuest wizard.

As you click Next to step through the problem, note:

- The objective is to minimize the mean of total costs.
- The problem has one requirement: the 95th percentile of the worst-case stockout forecast must be less than 0 gallons.
- This problem has eight discrete decision variables, with bounds of 0 to 2000. These represent transportation costs among the various elements of the supply chain.
- The problem has two constraints, shown in Figure 64. These constraints specify that both links of the supply chain (running through SD1 and SD2) must have sufficient inventory of gasoline.

Figure 64 Gasoline supply chain constraints

Constraints	Type	Exclude
1 Supply Depot 1 must have sufficient inventory of gasoline ModelIC13 + Refinery to SD1 - (SD1 to RO1 + SD1 to RO2 + SD1 to RO3) >= 0	Linear	<input type="checkbox"/>
2 Supply Depot 2 must have sufficient inventory of gasoline ModelII13 + Refinery to SD2 - (SD2 to RO1 + SD2 to RO2 + SD2 to RO3) >= 0	Linear	<input type="checkbox"/>

2 Run the optimization.

Figure 65 shows sample OptQuest results.

Figure 65 Gasoline supply chain model optimization results

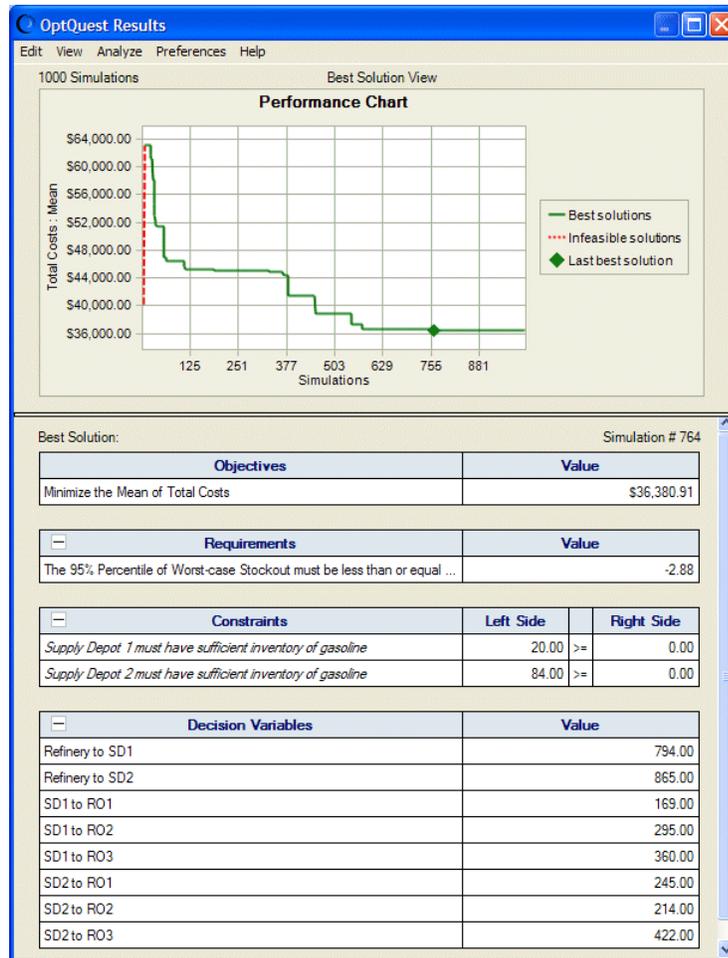
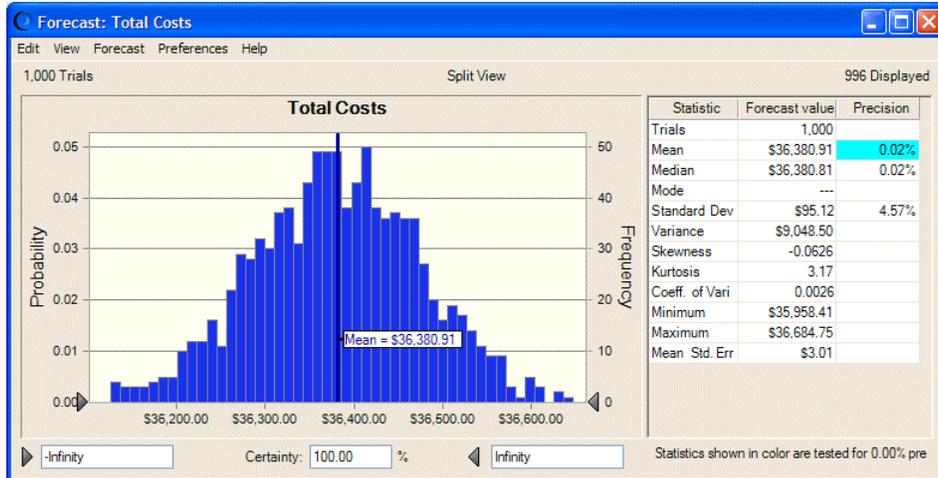
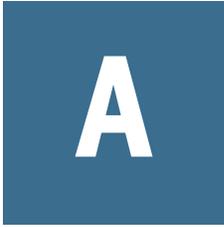


Figure 65 shows that if the quantities of gasoline shown for each decision variable are transported between the indicated destinations, the mean total cost will be \$36,380.91 and the 95th percentile of the worst-case stockout will be -2.88, less than 0.

Figure 66, following, shows the Total Costs forecast chart and statistics for the simulation of this solution. The Total Costs standard deviation is quite small relative to the total costs, suggesting that this cost forecast is an accurate representation of the true weekly costs.

Figure 66 Total Costs forecast chart and statistics





Optimization Tips and Notes

In This Appendix

Introduction.....	111
Model Types	111
Factors That Affect Optimization Performance.....	114
Sensitivity Analysis Using a Tornado Chart	119
Maintaining Multiple Optimization Settings for a Model.....	120
Other OptQuest Notes	121

Introduction

This appendix describes the different factors that affect how OptQuest searches for optimal solutions, including model types. Understanding how these factors affect the optimization helps you control the speed and accuracy of the search.

This appendix also includes discussion of the Crystal Ball Tornado Chart tool and how you can use it to analyze the sensitivity of the variables in your model and screen out minor decision variables.

These tips and suggestions are followed by some notes to help you avoid unexpected results when using OptQuest. They can also help you troubleshoot any difficulties that might occur.

Model Types

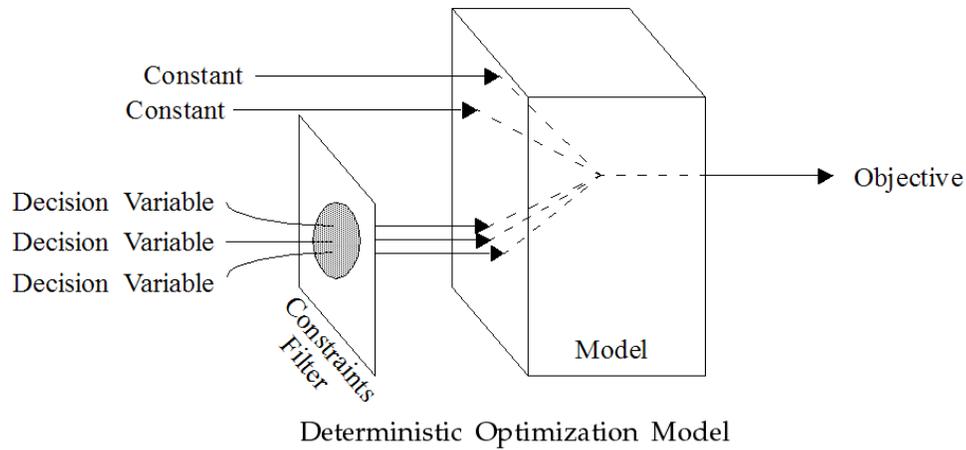
Selecting the right model for your scenario is essential for obtaining optimal results. These types of models are described here:

- “Optimization Models Without Uncertainty” on page 111
- “Optimization Models With Uncertainty” on page 112
- “Discrete, Continuous, or Mixed Models ” on page 113
- “Linear or Nonlinear Models” on page 114

Optimization Models Without Uncertainty

Conceptually, an optimization model might look like [Figure 67](#).

Figure 67 Schematic of an optimization model without uncertainty

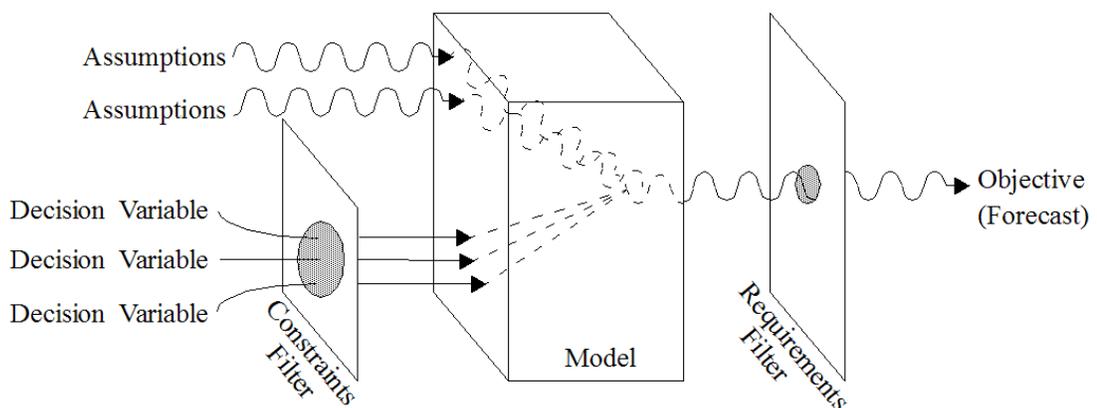


The solution to an optimization model provides a set of values for the decision variables that optimizes (maximizes or minimizes) the associated objective. If the world were simple and the future were predictable, all data in an optimization model would be constant, making the model *deterministic*.

Optimization Models With Uncertainty

In many cases, however, a deterministic optimization model can't capture all the relevant intricacies of a practical decision environment. When model data are uncertain and can only be described probabilistically, the objective will have some probability distribution for any chosen set of decision variables. You can find this probability distribution by simulating the model using Crystal Ball. This type of model is called *stochastic*.

Figure 68 Schematic of an optimization model with uncertainty



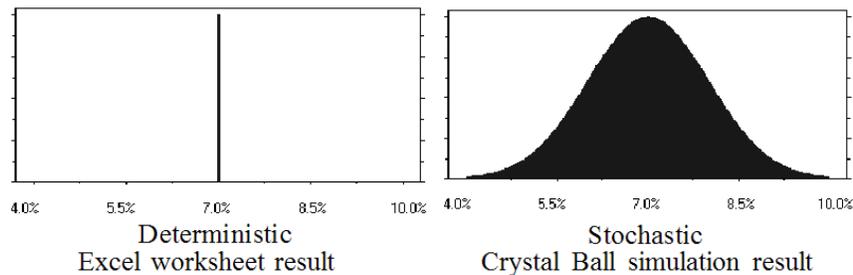
A stochastic optimization model has several additional elements:

- Assumptions — Capture the uncertainty of model data using probability distributions.
- Forecasts — Are frequency distributions of possible results for the model.

- Forecast statistics — Are summary values of a forecast distribution, such as the mean, standard deviation, or variance. You control the optimization by maximizing or minimizing forecast statistics, or setting them to a target.
- Requirements — Are additional restrictions on forecast statistics. You can set upper and lower limits for any statistic of a forecast distribution.

Stochastic models are much more difficult to optimize because they require simulation to compute the objective. While Crystal Ball is designed to solve stochastic models using Crystal Ball, it is also capable of solving deterministic models. [Figure 69](#) shows that deterministic results are a single value, while stochastic results are distributed over a curve.

Figure 69 Comparison of deterministic and stochastic results



Discrete, Continuous, or Mixed Models

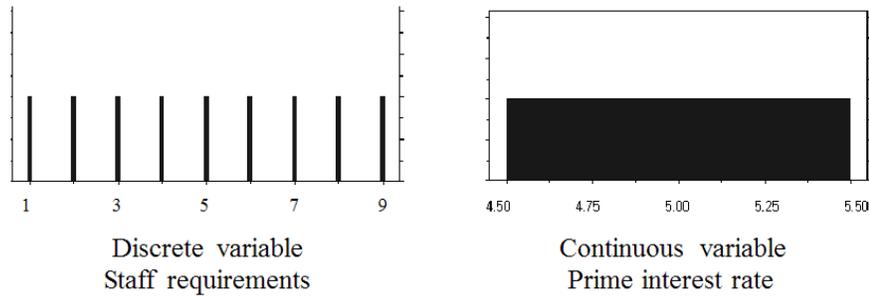
Optimization models can be classified as:

- Discrete — Contain only discrete decision variables.
- Continuous — Contain only continuous decision variables.
- Mixed — Contain both discrete and continuous decision variables, or any of the other decision variable types: binary, category, or custom.

For more information on discrete and continuous decision variables, see [“Decision Variables” on page 17](#).

[Figure 70](#) shows that discrete variable distributions are a series of individual values while continuous variable distributions are an infinite range of values without distinctive bounds except the end points.

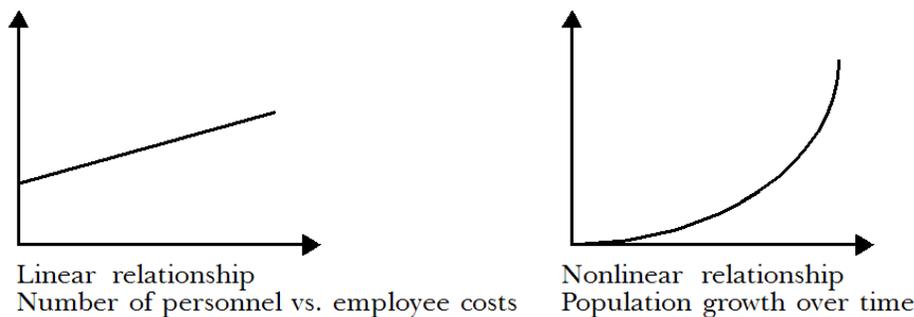
Figure 70 Comparison of discrete and continuous decision variables



Linear or Nonlinear Models

An optimization model can be linear or nonlinear, depending on the form of the mathematical relationships used to model the objective and constraints. Figure 71, following, illustrates linear and nonlinear relationships. In a linear relationship, all terms in the formulas only contain a single variable multiplied by a constant. For example, $3x - 1.2y$ is a linear relationship since both the first and second term only involve a constant multiplied by a variable. Terms such as x^2 , xy , $1/x$, or $3.1x$ make nonlinear relationships. Any models that contain such terms in either the objective or a constraint are classified as nonlinear.

Figure 71 Comparison of linear and nonlinear relationships



Crystal Ball can handle both linear and nonlinear objectives and constraints. For information on defining linear or nonlinear constraints, see [“Specifying Constraints”](#) on page 28.

Factors That Affect Optimization Performance

There are many factors that influence the performance of OptQuest. For example, consider two optimization methods, A and B, applied to an investment problem with the objective of maximizing expected returns. When you evaluate the performance of each method, you must look at which method:

- Finds an investment portfolio with a larger expected return
- Jumps to the range of high-quality solutions more quickly

Below is the performance graph for the two hypothetical methods.

Figure 72 Performance comparison

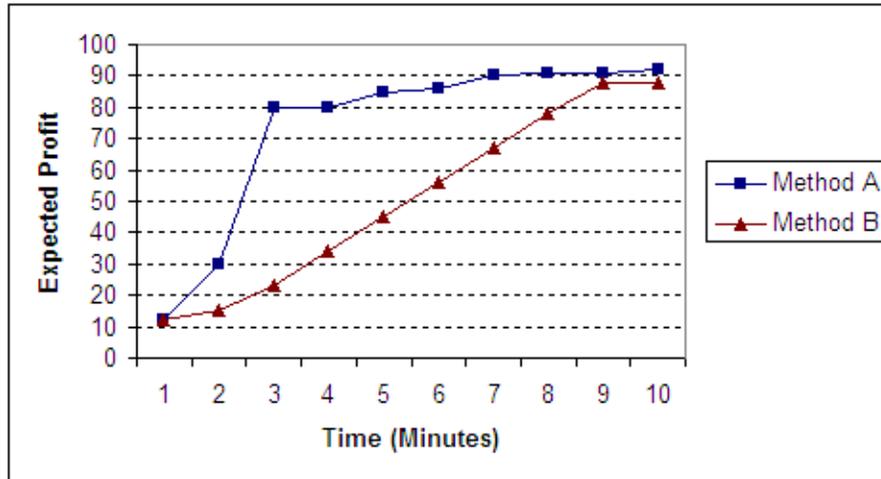


Figure 72 shows that although both methods find solutions with a similar expected profit after 10 minutes of searching, method A jumps to the range of high-quality solutions faster than B. For the criteria listed above, method A performs better than method B.

While using OptQuest, you will obtain performance profiles similar to method A. OptQuest’s search methodology (see the references in Appendix B) is very aggressive and attempts to find high-quality solutions immediately, causing large improvements (with respect to the initial solution) early in the search. This is critical when OptQuest can perform only a limited number of simulations within the available time limit.

However, several factors affect OptQuest’s performance, and the importance of these factors varies from one situation to another. The following is a list of the relevant factors that directly affect the search for an optimal solution:

- “Simulation Accuracy” on page 115
- “Number of Decision Variables” on page 116
- “Base Case Values” on page 116
- “Bounds and Constraints” on page 117
- “Requirements” on page 117
- “Complexity of the Objective ” on page 118
- “Simulation Speed” on page 118
- “Precision Control” on page 118

Simulation Accuracy

For sufficient accuracy, set the number of simulation trials to the minimum number necessary to obtain a reliable estimate of the statistic being optimized. For example, you can reliably estimate the mean with fewer trials than the standard deviation or a percentile.

General guidelines for determining the number of simulation trials necessary to obtain good estimates are:

- 200 to 500 trials is usually sufficient for obtaining accurate estimates for the mean.
- At least 1000 trials are necessary for obtaining reasonable estimates for tail-end percentiles.

Empirical testing with the simulation model using the Crystal Ball Bootstrap tool (see the *Oracle Crystal Ball User's Guide*) can help you find the appropriate number of trials for a given situation.

Number of Decision Variables

The number of decision variables greatly affects OptQuest's performance. OptQuest has no physical limit on the number of decision variables you can use in any given problem. As the number of decision variables increases, you need more simulations to find high-quality solutions. General guidelines for the minimum number of simulations required for a given number of decision variables in a problem are:

Decision variables	Minimum number of simulations
Fewer than 10	100
Between 10 and 20	500
Between 20 and 50	2000
Between 50 and 100	5000

For very large numbers of decision variables, you might try running more simulations by lowering the number of trials per simulation, at least initially. After you find an approximate solution, you can rerun the optimization by using the approximate solution as suggested values, further restricting the bounds on the decision variables, and increasing the number of trials to find more accurate results.

Recommended Number of OptQuest Elements

For best results, keep the number of OptQuest elements of each type below these limits:

- Decision variables < 4,096
- Constraints < 512
- Requirements < 512

Base Case Values

The base case values are the initial cell values listed in the Base Case column of the Decision Variables panel in the OptQuest wizard. The base case values are important because the closer they are to the optimal value, the faster OptQuest might find the optimal solution. If the values are constraint-infeasible, they will be ignored.

For potentially large models with many decision variables, you might find it helpful to first run a deterministic optimization to search for good base case values. Then, use the results as your

base case values and run a stochastic optimization. This technique, however, might not work well if you have objectives or requirements defined with other than central tendency statistics.

Bounds and Constraints

You can significantly improve OptQuest's performance by selecting meaningful bounds for the decision variables. Suppose, for example, that the bounds for three variables (X, Y, and Z) are:

$$0 \leq X \leq 100$$

$$0 \leq Y \leq 100$$

$$0 \leq Z \leq 100$$

And in addition to the bounds, there is the following constraint:

$$10 * X + 12 * Y + 20 * Z \leq 200$$

Although the optimization model is correct, the variable bounds are not meaningful. A better set of bounds for these variables would be:

$$0 \leq X \leq 20$$

$$0 \leq Y \leq 16.667$$

$$0 \leq Z \leq 10$$

These bounds take into consideration the values of the coefficients and the constraint limit to determine the maximum value for each variable. The new "tighter" bounds result in a more efficient search for the optimal values of the decision variables.

Since constraints limit the decision variables you are optimizing, OptQuest can eliminate sets of decision variable values that are constraint-infeasible before it spends the time running the simulation. Therefore, limiting the optimization with constraints is very time-effective.

Requirements

While the search process benefits from the use of constraints and tight bounds, performance generally suffers when you include requirements in the optimization model for two reasons:

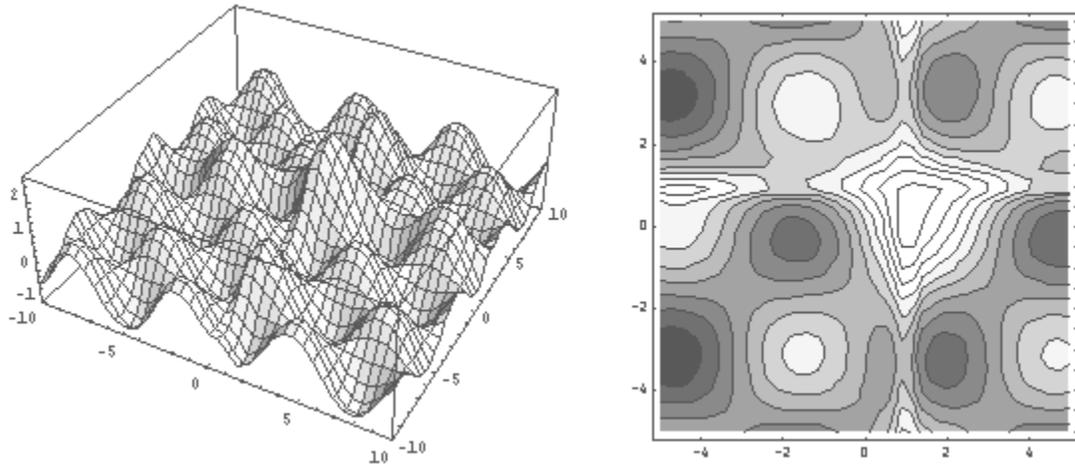
- Requirements are very time-consuming to evaluate, since OptQuest must run an entire simulation before determining whether the results are requirement-infeasible.
- To avoid running requirement-infeasible simulations, OptQuest must identify the characteristics of solutions likely to be requirement-feasible. This makes the search more complex and requires more time.

When you use requirements, you should increase the search time by at least 50% (based on the time used for an equivalent problem without requirements).

Complexity of the Objective

A complex objective has a highly nonlinear surface with many local minimum and maximum points.

Figure 73 Graphs of complex objectives



OptQuest is designed to find global solutions for all types of objectives, especially complex objectives like this one. However, for more complex objectives, you generally need to run more simulations to find high-quality global solutions.

Simulation Speed

By increasing the speed of each simulation, you can increase the number of simulations that OptQuest runs in a given time period. Some suggestions to increase speed are:

- Use Extreme Speed when practical.
- Use precision control in Crystal Ball to stop simulations as soon as they reach a satisfactory accuracy
- Reduce the size of your model
- Increase your system's RAM
- Reduce the number of assumptions and forecasts
- Quit other applications

The *Oracle Crystal Ball User's Guide* discusses these suggestions in more detail.

Precision Control

For some models, the accuracy of the statistics is highly dependent on the values of the decision variables. In these cases, you can use Crystal Ball's precision control feature to run a sufficient number of trials for each simulation to achieve the necessary level of accuracy.

You can use Crystal Ball's precision control feature for several purposes:

- When you are unsure of how to set the number of trials used for Crystal Ball simulations
- If you believe that the stability of the forecast statistics varies greatly depending on the decision variable values

Precision control periodically calculates the accuracy of the forecast mean, standard deviation, and any indicated percentile during the simulation. When the simulation reaches a desired accuracy, it stops, regardless of the number of trials already run.

This feature is especially useful for optimization models such as Portfolio Allocation, where the forecast statistics are highly sensitive to the decision variables. When OptQuest selects conservative investments, the variability of the expected return is low and the statistics are relatively stable. When OptQuest selects aggressive investments, the variability is high and the statistics are relatively less stable. Using precision control increases your forecast statistic accuracy while avoiding running too many trials when a simulation reaches this accuracy quickly.

Note that finding the appropriate precision control settings might require some trial and error. It can be challenging to decide whether to use absolute or relative precision, what is the best precision value in either case, and which statistics should receive precision control. For more information on setting the precision control feature, see the *Oracle Crystal Ball User's Guide*.

► To see the effects of using precision control with the Portfolio Allocation model:

1 In Crystal Ball, select Run, Run Preferences and change the maximum number of trials from 1000 to 5000.

This maximum limit is always in effect, even when precision control is turned on. Therefore, when using precision control, you must increase the maximum number of trials to let precision control achieve the appropriate accuracy.

2 Turn on Precision Control.

- a. Select cell C17.
- b. Select Define, Define Forecast.
- c. Click the More button in the Define Forecast dialog, then click the Precision tab.
- d. Check the Specify The Desired Precision For Forecast Statistics option.
- e. Check the Mean checkbox.
- f. Use an absolute precision of 1000 units.

3 Start OptQuest and reload the optimization settings file you saved earlier.

4 Run another optimization.

Experiment with various other precision control settings to see the difference in the results.

Sensitivity Analysis Using a Tornado Chart

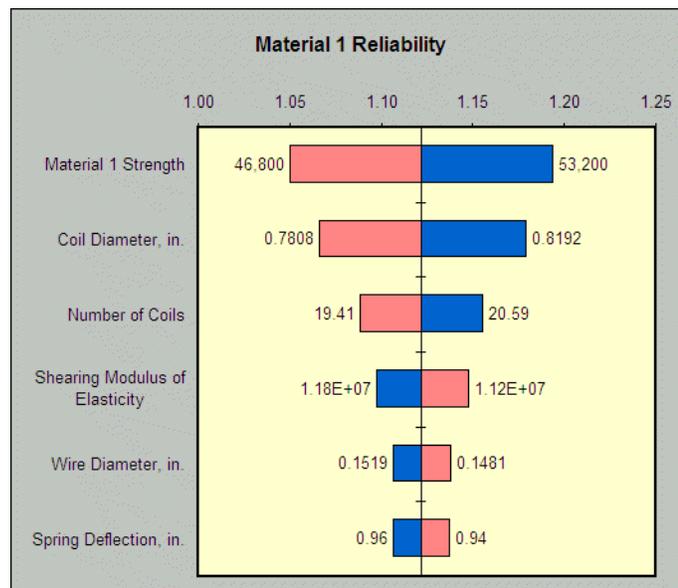
One of the easiest ways to increase the effectiveness of your optimization is to remove decision variables that require a lot of effort to evaluate and analyze, but that do not affect the objective very much. If you are unsure how much each of your decision variables affects the objective,

you can use the Tornado Chart tool in Crystal Ball (see the *Oracle Crystal Ball User's Guide* for more information on the Tornado Chart).

The Tornado Chart tool shows how sensitive the objective is to each decision variable as they change over their allowed ranges. The chart shows all the decision variables in order of their impact on the objective.

Figure 74 shows a Crystal Ball tornado chart. When you view a tornado chart, the most important variables are at the top. This arrangement makes it easier to see the relative importance of all the decision variables. The variables listed at the bottom are the least important in that they affect the objective the least. If their effect is significantly smaller than those at the top, you can probably eliminate them as variables and just let them assume a constant value.

Figure 74 Crystal Ball tornado chart



Before running the Tornado Chart tool, run an initial optimization so that the base case values of the decision variables are close to the optimal solution for your model. You can use the Tornado Chart tool to measure the impact of your decision variables. For information, see [“Viewing a Solution Analysis” on page 40](#).

Maintaining Multiple Optimization Settings for a Model

In this version of OptQuest, optimization settings are stored in workbooks instead of separate .opt files. Only one group of settings can be stored in each workbook. This is convenient for using and transferring models. However, there are times when you might want to have more than one group of optimization settings for a model. In that case, you can create different blank workbooks with one group of settings stored in each. Then, you can open a "profile" workbook with appropriate settings and use it as the primary workbook in the OptQuest wizard. As long as your main model workbook is also open, OptQuest will use the settings in the blank workbook and your model will still run as you intended.

Other OptQuest Notes

The following notes in this section can help you avoid problems while using OptQuest and can also assist in troubleshooting any difficulties that might happen:

- [“Automatic Resets of Optimizations” on page 121](#)
- [“Constraint Formula Limitations” on page 121](#)
- [“Minor Limit Violations With Continuous Forecasts” on page 122](#)
- [“Solutions Still Ranked Even With No Feasible Solution” on page 122](#)
- [“Referenced Assumption and Forecast Cells” on page 122](#)
- [“Decision Variables and Ranges With the Same Name” on page 122](#)
- [“Linear Constraints Can Be Evaluated As Nonlinear” on page 122](#)

Automatic Resets of Optimizations

If the first simulation does not run for some reason or generates an error, the entire optimization is reset. Otherwise, if the user stops a running optimization or an error occurs after an optimization starts successfully, the results to that point are kept and the optimization is not reset.

Constraint Formula Limitations

The following sections describe several limitations in defining constraint formulas:

- [“Array Formulas” on page 121](#)
- [“Date Formatting” on page 121](#)

Array Formulas

Array formulas with brackets are supported by Excel but are not allowed in OptQuest. For example, suppose you enter a constraint as follows, referencing a named range:

```
MyRange > {0}
```

An error about an unrecognized range or variable name is displayed.

Date Formatting

It is possible to reference a decision variable cell formatted as a date (such as 2/19/1900) and enter a constraint as follows in the OptQuest wizard's Constraints panel:

```
E2 > 2/19/1900
```

If you do this, OptQuest interprets it as 2 divided by 19 divided by 1900 and does not display an error message.

This behavior is consistent with that of the Excel formula bar. For best results, use the Excel DATE() function.

Minor Limit Violations With Continuous Forecasts

Slight violations of bounds can occur in requirements or constraints when evaluating small, continuous forecast values. If present, these violations should be ignored since the differences are very small compared to the relative magnitude of the forecast values.

Solutions Still Ranked Even With No Feasible Solution

If OptQuest fails to find a feasible solution, the Solution Analysis table still ranks solutions in order from best to worst objective value, even though none are feasible.

Referenced Assumption and Forecast Cells

If a constraint formula references an assumption or forecast cell, that cell is evaluated before each simulation runs. It is therefore not possible to enter a constraint on random values in the assumption cells or on the statistics in forecast cells. In general, you should avoid referencing these in constraint formulas.

In requirements, assumption and forecast cells are evaluated at the end of the simulation.

Decision Variables and Ranges With the Same Name

It is possible to define a decision variable and a cell address or range with the same name. If you do that, only the decision variable is accessible on the Constraints panel of the OptQuest wizard.

For best results, always give decision variables and ranges different names and avoid naming decision variables combinations of letters and numbers that resemble cell addresses.

Linear Constraints Can Be Evaluated As Nonlinear

If you have a cell reference in an OptQuest constraint that is more than seven levels of formulas removed from a decision variable, any constraint based on that cell will be evaluated as nonlinear even though it might be linear.



Accessibility

In This Appendix

Introduction	123
Accessibility Notes	123
OptQuest Wizard Keyboard Command Equivalents	124
OptQuest Results Window Menus	125
OptQuest Control Panel Keyboard Shortcuts	126

Introduction

This appendix describes the accessibility features of OptQuest, including keyboard equivalents for OptQuest commands. For details on accessibility features of Crystal Ball Decision Optimizer, see the *Oracle Crystal Ball User's Guide*. The following sections summarize accessibility features and keyboard shortcuts:

- [“Accessibility Notes” on page 123](#)
- [“OptQuest Wizard Keyboard Command Equivalents” on page 124](#)
- [“OptQuest Results Window Menus” on page 125](#)
- [“OptQuest Control Panel Keyboard Shortcuts” on page 126](#)

Accessibility Notes

The following sections discuss Crystal Ball and OptQuest accessibility features:

- [Accessibility of Code Examples in Documentation](#)
- [Accessibility of Links to External Web Sites in Documentation](#)
- [“Enabling Accessibility for Crystal Ball” on page 124](#)
- [“Using the Tab and Arrow Keys in the Crystal Ball Decision Optimizer User Interface” on page 124](#)
- [“TTY Access to Oracle Support Services” on page 124](#)

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Enabling Accessibility for Crystal Ball

You do not need to enable accessibility specifically for Oracle Crystal Ball Decision Optimizer, Fusion Edition, including OptQuest; it is always in accessible mode. OptQuest charts and other output can be extracted to Excel spreadsheets and pasted into PowerPoint slides, which are accessible through Microsoft Office. For information about Excel or PowerPoint accessibility, refer to Microsoft Office product documentation.

Using the Tab and Arrow Keys in the Crystal Ball Decision Optimizer User Interface

The main 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 wizard panels and 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.

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, seven days a week. For TTY support, call 800.446.2398.

OptQuest Wizard Keyboard Command Equivalents

Each panel of the OptQuest wizard has buttons that navigate through the wizard panels, run an optimization, close OptQuest, get online help, and perform other operations only available on a particular panel. When you click Alt in each panel, OptQuest Results view, or the OptQuest

Control Panel, shortcut keys are highlighted in each menu or button label. [Table 13](#) lists keyboard equivalents (shortcut keys) for OptQuest wizard buttons.

Table 13 OptQuest wizard keyboard shortcuts

Command	Panel	Keystrokes
Add Comment	Constraints	Alt-c
Add Objective	Objectives	Alt-o
Add Requirement	Objectives	Alt-r
Back	All but Welcome	Alt-b
Close	All	Alt-c
Delete	Objectives, Constraints	Alt-d
Efficient Frontier	Objectives, Constraints	Alt-e
Help	All	Alt-h
Insert Reference	Constraints	Alt-r
Insert Variable	Constraints	Alt-v
Next	All but Options	Alt-n
Run	All but Welcome	Alt-r

OptQuest Results Window Menus

The OptQuest Results window has the following menus, listed with the operations they perform. [Table 15](#) shows shortcut keys for commands on each menu.

Table 14 OptQuest Results window menus

Menu	Actions
Edit	Copies solutions to spreadsheet, copies charts, sets up pages for printing, prints
View	Switches between Best Solution (with Efficient Frontier) and Solution Analysis view
Analyze	Creates reports and extracts data
Preferences	Shows all solutions
Help	Displays help for the Best Solution, Solution Analysis, and Efficient Frontier windows

[Table 15](#) lists Alt-key combinations available in the OptQuest Results window to execute the listed menu commands without using the mouse. Commands are listed by menu in the order they appear on that menu. Not all commands are available in every view.

Table 15 OptQuest Results window keyboard shortcuts

Menu	Command	Keystrokes
Edit	Copy Best Solution to Spreadsheet	Alt-e, c
Edit	Copy Chart	Alt-e, o
Edit	Page Setup	Alt-e, g
Edit	Print Preview	Alt-e, r
Edit	Print	Alt-e, p
View	Best Solution	Alt-v, b
View	Solution Analysis	Alt-v, s
Analyze	Create Report	Alt-a, r
Analyze	Extract Data	Alt-a, e
Preferences	Show All Solutions	Alt-p, s
Help	Best Solution Help	Alt-h, b
Help	Solution Analysis Help	Alt-h, s
Help	Efficient Frontier Help	Alt-h, e

OptQuest Control Panel Keyboard Shortcuts

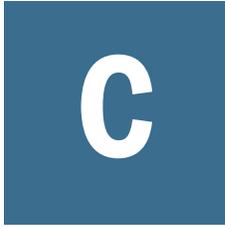
The OptQuest Control Panel controls OptQuest runs (Start, Stop, Continue, Reset), or you can click the Run Preferences button to control the maximum length of time or number of simulations for an optimization.

The Control Panel has three menus—Run, Analyze, and Help—to further control the performance of OptQuest and Crystal Ball. [Table 16](#), following, lists the commands and shortcut keys for each of the Control Panel menus.

Table 16 OptQuest Control Panel keyboard shortcuts

Menu	Command	Keystrokes
Run	Continue Optimization	Alt-r, c
Run	Stop Optimization	Alt-r, s
Run	Reset Optimization	Alt-r, r
Run	OptQuest	Alt-r, o
Run	Predictor	Alt-r, p
Run	Tools	Alt-r, t

Menu	Command	Keystrokes
Run	Save Results (Crystal Ball)	Alt-r, v
Run	Restore Results (Crystal Ball)	Alt-r, e
Run	Run Preferences	Alt-r, u
Analyze	Assumption Charts	Alt-a, a
Analyze	Forecast Charts	Alt-a, f
Analyze	Overlay Charts	Alt-a, o
Analyze	Trend Charts	Alt-a, t
Analyze	Sensitivity Charts	Alt-a, s
Analyze	Scatter Charts	Alt-a, e
Analyze	OptQuest Charts	Alt-a, q
Analyze	Cascade	Alt-a, c
Analyze	Close All	Alt-a, l
Analyze	Create Report	Alt-a, r
Analyze	Extract Data	Alt-a, d
Help	Control Panel Help	Alt-h, c



References and Bibliography

In This Appendix

Introduction.....	129
References.....	129
Spreadsheet Design	130
Optimization Topics	130
Financial Applications	130
Quality and Six Sigma Applications	131
Petrochemical Engineering Applications.....	131
Inventory System Applications.....	131

Introduction

This appendix provides a list of references on advanced topics suggested in this Guide. It is intended for advanced users who want more detail on topics such as metaheuristic methods and how optimizations work.

This appendix also includes bibliography entries by subject.

References

These references provide further detail on metaheuristic methods, comparisons of optimization methods, and optimization of complex systems. See the following references on our Web site. Also see “[Optimization Topics](#)” on page 130.

- Glover, F., J. P. Kelly, and M. Laguna. "The OptQuest Approach to Crystal Ball Simulation Optimization." Graduate School of Business, University of Colorado (1998).
- M. Laguna. "Metaheuristic Optimization with Evolver, Genocop and OptQuest." Graduate School of Business, University of Colorado, 1997.
- M. Laguna. "Optimization of Complex Systems with OptQuest." Graduate School of Business, University of Colorado, 1997.

A variety of white papers concerning optimization are available on the Web site of OptTek Systems, Inc., the company that developed the OptQuest calculation engine. For a list of papers with abstracts, see <http://www.opttek.com/News/WhitePapers.html>

Spreadsheet Design

Powell, S.G., and K.R. Baker. *The Art of Modeling with Spreadsheets: Management Science, Spreadsheet Engineering, and Modeling Craft*. Hoboken, NJ: John Wiley, 2003.

Ragsdale, C.T. *Spreadsheet Modeling and Decision Analysis: A Practical Introduction to Management Science*. 5th Ed. Mason, OH: South-Western College Publishing, 2007.

Thommes, M.C. *Proper Spreadsheet Design*. Boston: Boyd and Fraser Publishing Co., 1992.

Optimization Topics

Metaheuristics

Glover, F., J.P. Kelly, and M. Laguna. "New Advances and Applications of Combining Simulation and Optimization." Proceedings of the 1996 Winter Simulation Conference. Edited by J.M. Charnes, D.J. Morrice, D.T. Brunner, and J.J. Swain, 1996: 144-152.

Glover, F., and M. Laguna. *Tabu Search*. Boston: Kluwer Academic Publishers, 1997.

Laguna, M. "Scatter Search." *Handbook of Applied Optimization*. P.M. Pardalos and M.G.C. Resende (Eds.), Oxford Academic Press, 1999.

Stochastic (Probabilistic) Optimization Theory

Infanger, G. "Planning Under Uncertainty." Boston: Boyd & Fraser Publishing, 1994.

Kall, P., and S.W. Wallace. *Stochastic Programming*. New York: John Wiley and Sons, 1994.

Multiobjective Optimization

Chankong, V., and Y.Y. Haimes. *Multiobjective Decision Making: Theory and Methodology*. New York: North-Holland, 1983.

Hwang, C., and A. S. M. Masud. *Multiple Objective Decision Making -Methods and Applications*. Berlin: Springer-Verlag, 1979.

Keeney, R., and Raiffa, H. *Decisions with Multiple Objectives*. New York: John Wiley, 1976.

Financial Applications

Brealey, R., and S. Myers. *Principles of Corporate Finance*. 4th ed. New York: McGraw-Hill, Inc., 1991.

Chen, N., R. Roll, and S. Ross. "Economic Forces in the Stock Market." *Journal of Business*, 59 (July 1986): 383-403.

Markowitz, H.M. *Portfolio Selection*. 2nd ed. Cambridge, MA: Blackwell Publishers Ltd., 1991.

Quality and Six Sigma Applications

Creveling, C. *Tolerance Design: A Handbook for Developing Optimal Specifications*. Reading, MA: Addison-Wesley, 1997.

Pyzdek, T. *The Six Sigma Handbook, Revised and Expanded : The Complete Guide for Greenbelts, Blackbelts, and Managers at All Levels*. 2nd Ed. New York: McGraw-Hill, 2003.

Sleeper, A. *Design for Six Sigma Statistics*. New York: McGraw-Hill Professional, 2005.

Sleeper, A. *Six Sigma Distribution Modeling*. New York: McGraw-Hill Professional, 2006.

Petrochemical Engineering Applications

Humphreys, K.K. *Jelen's Cost and Optimization Engineering*. 3rd ed. New York: McGraw-Hill, 1991, 257-262.

Inventory System Applications

Evans, J.R., and D.L. Olsen. *Introduction to Simulation and Risk Analysis*. New York: Prentice-Hall, 1998.

Glossary

APT Arbitrage Pricing Theory.

assumption An estimated value or input to a spreadsheet model. Assumptions capture the uncertainty of model data using probability distributions.

bound A maximum or minimum limit you set for each decision variable.

certainty The percentage of simulation results that fall within a range.

coefficient of variability A measure of relative variation that compares the standard deviation to the mean. Results can be represented in percentages for comparison purposes.

constraint A limitation that restricts the possible solutions to a model. You must define constraints in terms of decision variables.

continuous A variable that can be fractional (that is, it can take on any value between the lower and upper bounds). No step size is required and any given range contains an infinite number of possible values. Continuous also describes an optimization model that contains only continuous variables.

correlation 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.

decision variable A variable in your model that you can control.

deterministic A model or system with no random variables that yields single-valued results.

discrete variable A variable that can only assume values equal to its lower bound plus a multiple of its step size; a step size is any number greater than zero, but less than the variable's range. Discrete also describes an optimization model that contains only discrete variables.

distribution See probability distribution.

efficient frontier The curve that plots an objective value against changes to a requirement or constraint. A typical use is for comparing portfolio returns against different risk levels.

efficient portfolio Combinations of assets for which it is impossible to obtain higher returns without generating higher risk or lower risk without generating lower returns. An efficient portfolio lies directly on the efficient frontier.

EOQ Economic Order Quantity.

feasible solution A solution that satisfies any constraints imposed on the decision variables, as well as any requirements imposed on forecast statistics.

final value The last value that is calculated for a forecast during a simulation. The final value is useful for when a forecast contains a function that accumulates values across the trials of a simulation, or is a function that calculates the statistic of another forecast.

forecast A statistical summary of the mathematical combination of the assumptions in a spreadsheet model, output graphically or numerically. Forecasts are frequency distributions of possible results for the model.

forecast objective One forecast from a model that OptQuest uses as the primary goal of the optimization. OptQuest maximizes or minimizes a statistic of the forecast's distribution.

forecast statistic Summary values of a forecast distribution, such as the mean, standard deviation, or variance. You control the optimization by maximizing or minimizing forecast statistics or setting them to a target value.

frequency distribution A chart that graphically summarizes a list of values by sub-dividing them into groups and displaying their frequency counts.

heuristic An approximate and self-educating technique for improving solutions.

inventory Any resource set aside for future use, such as raw materials, semifinished products, and finished products. Inventory also includes human, financial, and other resources.

inventory level The amount of inventory on hand, not counting ordered quantities not received.

inventory position The amount of inventory on hand plus any amount on order but not received, less any back orders.

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 A sampling method that divides an assumption's probability distribution into intervals of equal probability. The number of intervals corresponds to the Sample Size option available in the Crystal Ball 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.

linear A mathematical relationship where all terms in the formulas can only contain a single variable multiplied by a constant. For example, $3x - 1.2y$ is a linear relationship since both the first and second term involve only a constant multiplied by a variable.

maximum The largest value in a dataset.

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.

metaheuristic A family of optimization approaches that includes genetic algorithms, simulated annealing, tabu search, scatter search, and their hybrids.

minimum The smallest value in a dataset.

mixed A type of optimization model that has both discrete and continuous decision variables.

mode The value that, if it exists, occurs most often in a data set.

model A representation of a problem or system in a spreadsheet application such as Excel.

multiobjective optimization A technique that combines multiple, often conflicting objectives, such as maximizing returns and minimizing risks, into one objective.

nonlinear A mathematical relationship where one or more terms in the formulas are nonlinear. Terms such as x^2 , xy , $1/x$, or $3.1x$ make nonlinear relationships. See linear.

NPV Net Present Value. The NPV equals the present value minus the initial investment.

objective A forecast formula in terms of decision variables that gives a mathematical representation of the model's goal.

optimal solution The set of decision variable values that achieves the best outcome.

optimization A process that finds the optimal solution to a model.

optimization model A model that seeks to maximize or minimize some quantity (the objective), such as profit or risk.

order quantity The standard amount of product you reorder when inventory reaches the reorder point.

percentile A number on a scale of zero to one hundred that indicates the percent of a probability distribution that is equal to or below a value (default definition).

performance For an optimization program, the ability to find high-quality solutions as fast as possible.

probability The likelihood of an event.

probability 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 A method whereby Crystal Ball replaces assumption values with their ranking from lowest value to highest value (1 to N) prior to computing the correlation coefficient. With this method, you can ignore the distribution types when correlating assumptions.

RAROC A multiobjective function that calculates the Risk-adjusted Return On Capital.

reorder point The inventory position when you reorder.

requirement A restriction on a forecast statistic that requires the statistic to fall between specified lower and upper limits for a solution to be considered feasible.

risk The uncertainty or variability in the outcome of some event or decision.

risk factor A number representing the riskiness of an investment relative to a standard, such as U.S. Treasury bonds, used especially in APT.

safety stock The additional quantity kept in inventory above planned usage rates.

seed value The first number in a sequence of random numbers. A given seed value produces the same sequence of random numbers for assumption values 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 or decision variable cell.

sensitivity analysis The computation of a forecast cell's sensitivity with respect to the assumption or decision variable cells.

simulation A set of Crystal Ball trials. OptQuest finds optimal solutions by running multiple simulations for different sets of decision variable values.

skewed An asymmetrical distribution.

skewness The measure of the degree of deviation of a curve from the norm of an asymmetric distribution. The greater the degree of skewness, the more points of the curve lie on one side of the peak of the curve as compared to the other side. A normal distribution curve, having no skewness, is symmetrical.

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, that is, the dispersion of values around the mean.

step size Defines the difference between successive values of a discrete decision variable in the defined range. For example, a discrete decision variable with a range of 1 to 5 and a step size of 1 can take on only the values 1, 2, 3, 4, or 5; a discrete decision variable with a range of 0 to 17 with a step size of 5 can take on only the values 0, 5, 10, and 15.

stochastic A model or system with one or more random variables.

STOIIP Stock Tank Oil Initially In Place. STOIIP is the estimated reserves of an oil field in millions of barrels (mmbbls).

trial A three-step process in which Crystal Ball generates random numbers for assumption cells, recalculates the spreadsheet models, and displays the results in a forecast chart. A Crystal Ball simulation is made up of multiple trials.

variable A quantity that might assume any one of a set of values and is usually referenced by a formula.

variance The square of the standard deviation, where standard deviation is approximately the average of the sum of the squares of the deviations of a number of observations (n) from their mean value (except the sum is divided by $n-1$ instead of n , which would yield a true average).

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 larger.

wizard A feature that leads you through the steps to create and run an optimization model. This wizard presents panels for you to complete in the proper order.

Index

Symbols

.opt files, 46

A

analysis

 solution, 38

 using tornado chart, 119

apartment tutorial, 49

assumptions

 defining, 25

B

bibliography

 by subject, 129

 financial applications, 130

 inventory systems, 131

 optimization topics, 130

 petrochemical engineering, 131

 spreadsheet design, 130

 tolerance design, 131

bounds

 affecting performance, 117

 defined for decision variables, 17

budget-constrained project selection example, 75

C

capability metrics, 21

charts

 viewing, 42

charts, viewing, 42

cleanup, groundwater example, 79

closing OptQuest, 45

commands

 start/pause/stop, 36

commands, keyboard

 OptQuest, 126

complexity of objective, 118

constraint editor syntax, 32

constraint feasibility defined, 19

constraints

 affecting performance, 117

 defined, 18, 60

 defining, 28

 editor, 31

 recommended number, 116

 simple entry, 29

 syntax, 32

continuous

 decision variables, 17

 models, 113

creating reports, 42

Crystal Ball

 models, creating, 24

Crystal Ball charts, 42

Crystal Ball web page, 129

D

data

 extracting, 44

 extracting to spreadsheets, 44

 including in reports, 42

decision variables

 bounds defined, 17

 defined, 13, 17

 in constraints, 31

 number affecting performance, 116

 recommended number, 116

 selecting to optimize, 27

 step size, 18

 types, 17

Design for Six Sigma, 21

deterministic

 model illustrated, 112

discrete

- decision variables, 17
- models, 113
- variable step size, 18

drill bit replacement example, 102

E

efficient portfolios, 20

engineering, petrochemical references, 131

examples

- drill bit replacement, 102
- groundwater cleanup, 79
- hotel design and pricing, 72
- inventory system, 96
- oil field development, 83
- overview, 67
- portfolio revisited, 86
- product mix, 69
- project selection, 75
- requirements, 17
- tolerance analysis, 92

exiting OptQuest, 45

extracting data, 44

extracting OptQuest data, 44

F

feasibility

- constraint, defined, 19
- requirement, 19

feasible solutions, 53

feature changes, 24

files

- optimization, name, 36

financial applications, references, 130

flow chart, OptQuest, 15

forecast statistics

- defined, 16
- maximizing or minimizing, 16

forecasts

- cells as objectives, 16
- defining, 25
- restricting statistics, 17
- selecting objective, 26
- statistics, defined, 113

Futura Apartments tutorial, 49

G

getting started, 49

groundwater cleanup example, 79

guidelines for using OptQuest, 55

H

help, 10

hotel design and pricing example, 72

how OptQuest works, 14

how this manual is organized, 9

I

icons

- OptQuest, 126

importing settings from .opt files, 46

initial values, affecting performance, 116

inventory system example, 96

inventory systems, references, 131

K

keyboard commands

- OptQuest, 126

L

Lean principles, 21

limits

- OptQuest elements, 116

linear models, 114

M

mathematical operations in constraints, 32

maximizing forecast statistic, 16

menus

- OptQuest, 125

metaheuristics

- references, 130

minimizing forecast statistic, 16

mixed models, 113

models, 15

- creating, 24
- deterministic illustrated, 112
- linear and nonlinear, 114
- optimization, defined, 15
- setting up, 23

N

new features, 24
 nonlinear models, 114

O

objectives

- complex, 118
- defined, 58
- selecting forecast, 26
- using forecasts as, 16

 oil field development example, 83
 online help, 10
 operations, mathematical in constraints, 32
 optimization performance, 114
 optimization process, overview, 23
 optimization topics, references, 130
 optimizations

- deterministic model, illustrated, 112
- model, defined, 15
- running, 35
- starting and stopping, 36

 options

- selecting, 34

 OptQuest

- flow, 15
- how it works, 14
- keyboard commands and icons, 126
- options, 34
- steps to use, 23
- what it does, 13

 OptQuest element limits, 116
 OptQuest menus, 125
 OptQuest results

- printing, 42

 OptQuest, guidelines for using, 55
 OptQuest, starting, 26
 organization, manual, 9

P

pause command, 36
 performance factors

- bounds and constraints, 117
- complex objectives, 118
- initial values, 116
- number of decision variables, 116
- requirements, 117

 simulation speed, 118
 performance, affecting factors, 114
 performance, optimization, 114
 petrochemical engineering, references, 131
 Portfolio Allocation tutorial, 54
 portfolio revisited example, 86
 portfolios, efficient, 20
 preferences, suggested run, 25
 printing OptQuest results, 42
 printing results, 42
 process capability, 21
 process, optimization, 23
 product mix example, 69
 project selection example, 75

Q

quality programs, 21

R

ranges

- decision variable, 17

 references

- financial applications, 130
- inventory systems, 131
- metaheuristics, 130
- on the web, 129
- optimization topics, 130
- petrochemical engineering, 131
- spreadsheet design, 130
- tolerance design, 131

 remaining time, viewing, 36
 reporting on OptQuest data, 42
 reports

- creating, 42
- using OptQuest data, 42

 requirements

- affecting performance, 117
- defined, 17, 58, 113
- defining, 26
- examples, 17
- feasibility, 19
- recommended number, 116
- variable, 46

 results

- analyzing in Crystal Ball, 42
- interpreting in OptQuest, 40

printing, [42](#)
 run preferences, suggested, [25](#)

S

saving OptQuest settings, [45](#)
 screen capture notes, [10](#)
 sensitivity analysis, using tornado chart, [119](#)
 Simple Entry mode, [29](#)
 simulations
 accuracy of, [115](#)
 current number, [36](#)
 running longer, [42](#)
 speed of, [118](#)
 Six Sigma, [21](#)
 solution analysis of results, [40](#)
 Solution Analysis window, [38](#)
 solutions
 feasible, defined, [53](#)
 viewing, [36](#)
 speed of simulations, [118](#)
 spreadsheet design, references, [130](#)
 spreadsheet models, creating, [24](#)
 start command, [36](#)
 starting OptQuest, [26](#)
 statistics
 forecast, defined, [16](#), [113](#)
 forecast, optimizing, [16](#)
 restricting forecasts, [17](#)
 selecting forecast, [26](#)
 Status And Solutions window, [36](#)
 step sizes, for decision variables, [18](#)
 steps for using OptQuest, [23](#)
 stop command, [36](#)
 syntax, constraint, [32](#)

T

time remaining, viewing, [36](#)
 tolerance analysis example, [92](#)
 tolerance design, references, [131](#)
 tornado chart, [119](#)
 transferring settings from .opt files, [46](#)
 tutorials
 Futura Apartments, [49](#)
 Portfolio Allocation, [54](#)
 types
 decision variable, [17](#)

U

using OptQuest, [55](#)

V

variable requirements, [46](#)
 variables
 decision, defined, [17](#)
 decision, range, [17](#)
 decision, step size, [18](#)
 decision, types, [17](#)
 variables, decision
 defined, [13](#)
 in constraints, [31](#)
 number affecting performance, [116](#)
 selecting to optimize, [27](#)
 viewing charts, [42](#)

W

web pages
 references, [129](#)
 what OptQuest does, [13](#)
 what's new, [24](#)
 windows
 Solution Analysis, [38](#)
 Status And Solutions, [36](#)