Chapter 1

Crystal Ball Overview

In this chapter

- Model building and risk analysis overview
- Steps for using Crystal Ball
- Starting and closing Crystal Ball
- The Crystal Ball menus and toolbar

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

Chapters 2 through 7 of this *Crystal Ball User Manual* describe how to build, run, analyze, and present your own Crystal Ball model simulations.

Model building and risk analysis overview

Glossary Term:
simulation—
Any analytical method that is meant to imitate a real-life system, especially when other analyses are too mathematically complex or too difficult to reproduce.

Glossary Term:
spreadsheet model—
Any spreadsheet that represents an actual or hypothetical system or set of relationships.

Glossary Term:
forecast—
A statistical summary of the simulation results in a spreadsheet model, displayed graphically or numerically.

Glossary Term:
risk—
The possibility of loss, damage, or other undesirable event and the severity associated with the event.

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

The basic process for using Crystal Ball, then, is to:

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

This User Manual is structured to match those main tasks. The guidelines on page 13 fill in some details and indicate where each task is discussed.

If you are new to Crystal Ball and risk analysis tools, you might not be familiar with or know what is meant by models or risk analysis. Or, if you do know, you might want a better understanding of how Crystal Ball performs a risk analysis.

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

Risk and risk analysis

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

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

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

Almost any change, good or bad, poses some risk. Your own analysis will usually reveal numerous potential risk areas: overtime costs, inventory shortages, future sales, geological survey results, personnel fluctuations, unpredictable demand, changing labor costs, government approvals, potential mergers, pending legislation.

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

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

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

What is a model?

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

Data vs. analysis

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

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

Traditional spreadsheet analysis

So now you have a model, or you have created your first model. For each variable in your model, ask yourself, "How certain am I of its value? Will it vary? Is this a best estimate or a known fact?" You might notice that your model has some variables in it that aren't definitely certain. Perhaps you don't have the actual data yet (this month's sales figures) or the variable behaves unpredictably (individual item cost). This lack of knowledge about particular values or how some variables behave contributes to the model's *uncertainty*.

Traditional spreadsheet analysis tries to capture this uncertainty in one of three ways:

- Point estimates
- Range estimates
- What-if scenarios

Point estimates

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

Range estimates

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

What-if scenarios

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

As you can see, this is extremely time consuming, and results in lots of data, but still doesn't give you the probability of achieving different outcomes.

You are still faced with these two fundamental limitations of ordinary spreadsheets:

- You can change only one spreadsheet cell at a time. As a result, exploring the entire range of possible outcomes is next to impossible; you cannot realistically determine the amount of risk that is impacting your bottom line.
- "What-if" analysis always results in single-point estimates which do not indicate the likelihood of achieving any particular outcome. While single-point estimates might tell you what is *possible*, they do not tell you what is *probable*.

This is where simulation with Crystal Ball comes in.

Monte Carlo simulation and Crystal Ball

Glossary Term:
spreadsheet model— any spreadsheet that represents an actual or hypothetical system or set of relationships.

Glossary Term:
simulation— Any analytical method that is meant to imitate a real-life system, especially when other analyses are too mathematically complex or too difficult to reproduce.

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

History

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

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

Probability distributions and assumptions

Glossary Term:
probability distribution—
A set of all possible events and their associated probabilities.

For each uncertain variable in a simulation, you define the possible values with a ***probability distribution***. A simulation calculates numerous scenarios of a model by repeatedly picking values from the probability distribution for the uncertain variables and using those values for the cell. Commonly, a Crystal Ball simulation calculates hundreds or thousands of scenarios in just a few seconds.

Glossary Term:
assumption—
An estimated value or input to a spreadsheet model.

In Crystal Ball, distributions and associated scenario input values are called ***assumptions***. They are entered and stored in assumption cells. For more information on assumptions and probability distributions, see "About assumptions and probability distributions" beginning on page 18.

Forecasts

Glossary Term:
forecast—
A statistical summary of the simulation results in a spreadsheet model, displayed graphically or numerically.

Since all those scenarios produce associated results, Crystal Ball also keeps track of the ***forecasts*** for each scenario. These are important outputs of the model, such as totals, net profit, or gross expenses. They are defined with formulas in spreadsheet forecast cells.

For each forecast, Crystal Ball remembers the cell value for all the trials (scenarios). If you run a simulation at Demo speed, you can watch histograms of the results calculated for each forecast cell and can see how the results stabilize toward a smooth frequency distribution as the simulation progresses. After hundreds or thousands of trials, you can view sets of values, the statistics of the results (such as the mean forecast value), and the certainty of any particular value. Chapter 5 gives more information about charts of forecast results and how to interpret them.

Certainty

Glossary Term:
certainty—
The percent chance that a particular forecast value will fall within a specified range.

The forecast results show you not only the different result values for each forecast, but also the probability of obtaining any value. Crystal Ball normalizes these probabilities to calculate another important number: the ***certainty***.

The chance of any forecast value falling between –Infinity and +Infinity is always 100%. However, the chance — or certainty — of that same forecast being at least zero (which you might want to calculate to make sure that you make a profit) might be only 45%.

For any range you define, Crystal Ball calculates the resulting certainty. This way, not only do you know that your company has a chance to make a profit, but you can also quantify that chance by saying that the company has a 45% chance of making a profit on a venture (a venture you might, therefore, decide to skip).

Benefits of Monte Carlo analysis

Crystal Ball uses Monte Carlo simulation to overcome both of the spreadsheet limitations listed earlier:

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

In addition, Crystal Ball keeps track of the results of each scenario for you.

Steps for using Crystal Ball

Glossary Term:
assumption cell—
A value cell in a spreadsheet model that has been defined as a probability distribution.

Glossary Term:
forecast cell—
Cells that contain formulas that refer to one or more assumption and decision variable cells and combine the values in the assumption, decision variable, and other cells to calculate a result.

Glossary Term:
decision variable cell—
Cells that contain the values or variables that are within your control to change. The decision variable cells must contain simple numeric values, not formulas or text.

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

1. **Create a spreadsheet model in Microsoft Excel format with data and formula cells that represent the situation you want to analyze.**

   "What is a model?" on page 9 discusses spreadsheet models. Also see the references in the "Bibliography" beginning on page 335.

2. **Start Crystal Ball.**

   If you haven't set up Crystal Ball to load automatically with Microsoft Excel, start Crystal Ball as described on page 14.

3. **Load your spreadsheet model.**
4. **Using Crystal Ball, define assumption cells and forecast cells. If appropriate for your situation, you can also define decision variable cells.**

   For more information, see "Entering an assumption" beginning on page 19 and continue on with Chapter 3.

5. **Set run preferences for your simulation, as described beginning on page 80.**

6. **Run the simulation, following the instructions beginning on page 92.**

7. **Analyze your results. See "Understanding and using forecast charts" beginning on page 106 for suggestions.**

8. **If you have the Professional edition of Crystal Ball, consider using CB Predictor or OptQuest for further analysis.**

9. **Take advantage of the many resources Decisioneering offers to help you get the most out of Crystal Ball.**

   If you are new to Crystal Ball, the *Crystal Ball Getting Started Guide* offers tutorials to quickly introduce Crystal Ball's features and workflow. Consider completing the tutorials before you continue on with the more detailed chapters that follow in this User Manual.

   You can choose Start > Programs > Crystal Ball 7 > Crystal Ball Tutorial to run through a brief online tutorial that teaches Crystal Ball basics. Or, choose Start > Programs > Crystal Ball 7 > Training CD Demo to learn about a more extensive online tutorial available for purchase from Decisioneering, Inc.

   For a list of support, training, and referral services, see "Additional resources" on page 5. Papers, user group information, conference schedules, newsletter subscriptions, and more are available on our Web Site:

   http://www.crystalball.com

Starting and closing Crystal Ball

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

Starting Crystal Ball manually

To start Crystal Ball manually:

1. **In Windows, choose Start > Programs > Crystal Ball 7 > Crystal Ball.**

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

Starting Crystal Ball automatically

To set Crystal Ball to start automatically each time you start Excel:

1. **In Windows, choose Start > Programs > Crystal Ball 7 > Application Manager.**
2. **Check Automatically Launch Crystal Ball 7 When Excel Starts.**
3. **Click OK.**

Closing Crystal Ball

To close Crystal Ball, either:

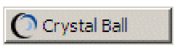

- Right-click the Crystal Ball icon in the Windows taskbar and choose Close, or
- Close Excel.

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

The Crystal Ball menus and toolbar

The Crystal Ball menus

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

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

The following chapters of this book explain how to use the various commands. For specific information about commands, see the Index at the end of this book as well as the Crystal Ball online help (see page 5 for more information on help).

The Crystal Ball toolbar

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

Figure 1.1 The Crystal Ball toolbar

The tools in the first three groups are from the Define menu. The tools from the next two groups are from the Run menu. The tools from the following two groups are from the Analyze menu, and the tool in the last group displays Crystal Ball online help.

To hide or display the Crystal Ball toolbar for the current session, choose View > Toolbars > Crystal Ball 7.

Chapter 2

Defining Model Assumptions

In this chapter

- Overview
- Defining assumptions
- Additional assumption features
- Additional Distribution Gallery features

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

If you are a new user, you should start by working through the tutorials in the *Crystal Ball Getting Started Guide*, and then read this chapter. After you complete this chapter and Chapter 3 of this User Manual, read Chapter 4 for information on setting preferences and running simulations.

Overview

Crystal Ball lets you define three types of cells:

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

Crystal Ball Note: *For previous versions of Crystal Ball, it might have been necessary to define forecasts in the same cells as assumptions or decision variables to capture that data for later extraction. Now, assumption and decision variable data can be extracted as well as forecast data. For this reason, Crystal Ball 7 no longer supports two types of cell definition in the same cell. If an assumption or a decision variable is defined in the same cell as a forecast in a Crystal Ball 4.x or 5.x (2000.x) workbook, the forecast will be deleted when the workbook is converted to Crystal Ball 7 format.*

Defining assumptions

About assumptions and probability distributions

For each uncertain variable in a simulation, or assumption, you define the possible values with a probability distribution. The type of distribution you select depends on the conditions surrounding the variable. For example, some common distribution types are shown in Figure 2.1.

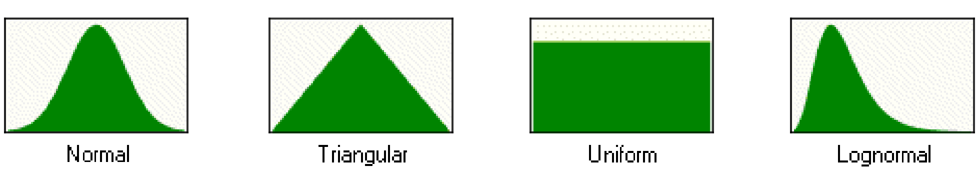

Figure 2.1 Common distribution types

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

Because distributions for independent variables are so important to simulations, selecting and applying the appropriate distribution is the main part of defining an assumption cell. For more information on probability distributions, see "Understanding probability distributions" beginning on page 242

Defining an assumption

To define an assumption, you must:

1. **Identify a distribution type as described in "Selecting a probability distribution" beginning on page 246.**

   Crystal Ball uses probability distributions to describe the uncertainty in your assumption cells. From a gallery of distribution types, you choose the ones that best describe the uncertain variables in the problem you are trying to solve. Appendix A describes each distribution type in detail.

   You can also select a distribution type by fitting a distribution to data. For details, see "Fitting distributions to data" beginning on page 27.

2. **Enter the assumption as described in the next section, "Entering an assumption."**

Entering an assumption

To enter an assumption:

1. **Select a cell or a range of cells.**

   Select value cells or blank cells only. Assumptions cannot be defined for formula or non-numeric cells.

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

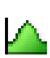

2. **Choose Define > Define Assumption.**

For each selected cell or cells in the selected range, Crystal Ball displays the Distribution Gallery dialog.

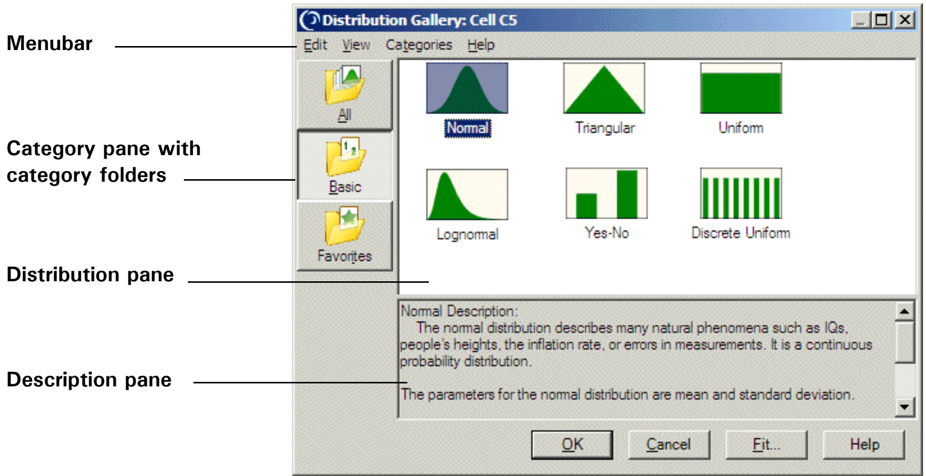

Figure 2.2 The Distribution Gallery with Basic category selected

By default, the Basic category appears when the Distribution Gallery opens. Only the most common probability distributions appear in the window.

Crystal Ball Note: *The All category contains all distributions originally shipped with Crystal Ball. If you modify and save one of these original distributions, it appears in the Favorites category unless you create and specify another category.*

You can perform many tasks in the Distribution Gallery:

- To see more distributions, click the All folder in the category pane. You can use the upper scroll bars, resize the window, or change the View menu settings to view all the distributions in that category.
- To view the description of a distribution, click it; the description appears in the lower pane.
- To have Crystal Ball select a distribution for you based on your data, click the Fit button. See "Fitting distributions to data" on page 27 for details.

You can also add categories, add distributions to categories, customize distributions, share categories with other Crystal Ball users, and more. For details, see "Additional Distribution Gallery features" beginning on page 41.

3. **Select a category from the folders in the category pane, and then double-click the distribution you want to use.**

   A dialog appears, showing the distribution type you chose for the selected cell (or for the first cell in a range of cells). Figure 2.3 shows an example of the normal distribution.

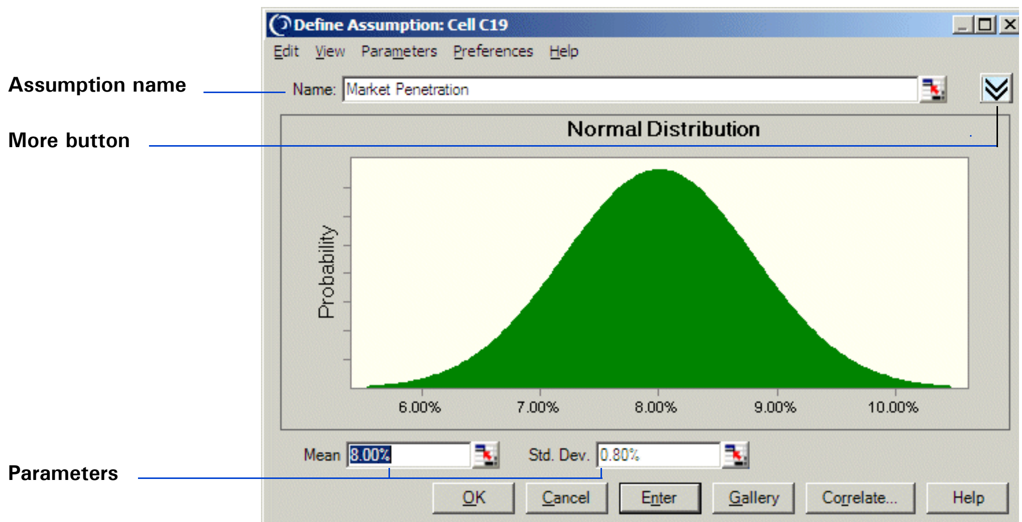

Figure 2.3 Normal distribution

Crystal Ball Note: *If you want to change the distribution type, click Gallery to return to the Distribution Gallery and then select another distribution.*

4. **In the dialog, type a name for the assumption (optional).**

   If the assumption cell already has a name next to it or above it on the spreadsheet (or has been named in Excel), the name appears in the dialog. You can use that name or type a new name. You can also use cell referencing to name the assumption.

5. **Type the parameters for the distribution.**

   Default values appear for the distribution parameters. You can type new values or cell references and formulas in any field. For more about using cell references and formulas, see "Entering cell references and formulas" on page 24.

6. **To see more information, click the More button near the Name field.**

   More information appears in the Define Assumption dialog as shown in Figure 2.4.

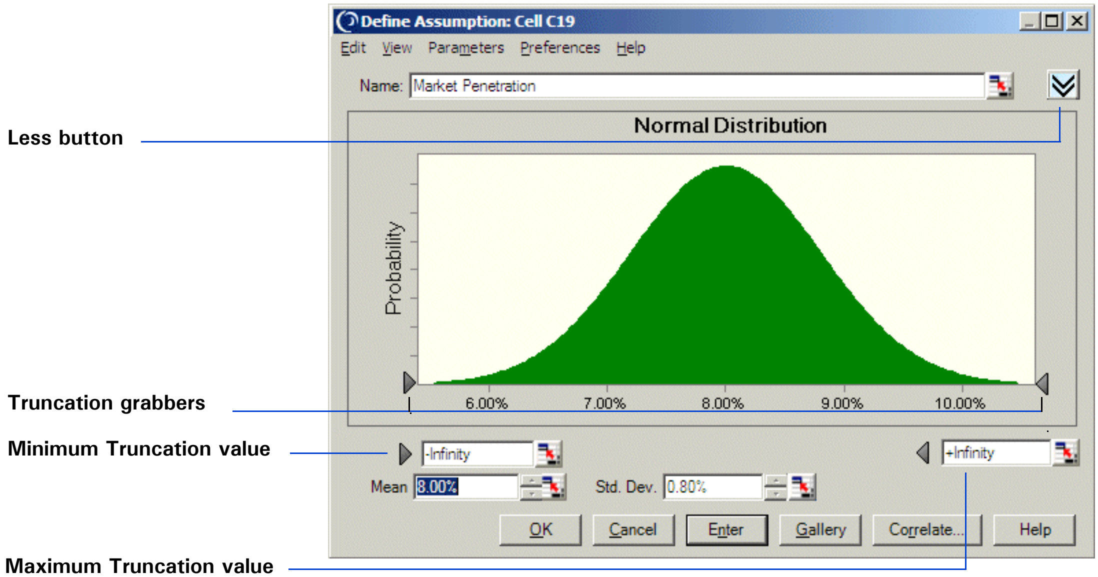

Figure 2.4 The Define Assumption dialog, expanded

In the expanded Define Assumption dialog, you can:

- Enter truncation minimum and maximum values in the fields just below the distribution.
- Use the truncation grabbers to truncate the value range.
- Use numeric spinners (arrows to the right of the field) to adjust parameter settings.
- Click the Less button to hide the minimum and maximum value fields and truncation grabbers.

Crystal Ball Note: *For more information about truncating distributions, see "Truncating distributions" on page 308.*

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

- Click the Gallery button to display the Distribution Gallery window and choose another distribution.
- Click the Correlate button to define correlations as described on page 33.

- Choose Edit > Add in the menubar to add the currently defined assumption distribution to the Favorites category or a user-defined category in the Distribution Gallery.
- Use other menu commands to copy the chart, paste it into Excel or another application, print data, change the view, use alternate parameters, set assumption and chart preferences, and display help as described in "Additional assumption features" beginning on page 23.

7. **When you have finished entering parameters to define the assumption, click Enter.**

   The distribution changes to reflect the values you entered.

Crystal Ball Note: *If you click OK instead of Enter, Crystal Ball accepts the parameters and closes the dialog.*

8. **Click OK.**

   If you selected a range of cells, repeat steps 3-8 to define the assumption for each cell.

Additional assumption features

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

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

- "Entering cell references and formulas," below
- "Alternate parameter sets" on page 25
- "Setting assumption preferences" on page 38
- "Fitting distributions to data" on page 27
- "Specifying correlations between assumptions" on page 33
- "Freezing Crystal Ball data cells" on page 88

Entering cell references and formulas

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

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

Dynamic vs. static cell references

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

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

Crystal Ball Note: *In previous versions of Crystal Ball, you could choose whether to use static or dynamic cell referencing in parameters. With static referencing, all cell references are resolved at the start of a simulation and then frozen while a simulation is running. If you open a model from a previous version, any static references are converted into dynamic references. If you don't want parameter values to change when a simulation is running, be sure cell references in parameters do not reference Crystal Ball data cells (assumptions, decision variables, and forecasts) directly or indirectly through formulas. If a dynamic cell reference is circular, that is, it references itself directly or indirectly, the assumption's value will treated as static. Otherwise, the model will run as expected.*

Relative references

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

Absolute references

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

Range names

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

Formulas

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

Alternate parameter sets

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

For example, if you are defining a triangular distribution, but are unsure of the absolute minimum and maximum values of the variable, you could instead define the distribution using the 10th and 90th percentiles along with the likeliest value. This gives you a distribution that has 80%, or four-fifths of the values, occurring between the two specified percentiles, as in Figure 2.5.

To change the parameter sets for the continuous distributions, use the Parameters menu in the menubar of the Define Assumption dialog. The currently selected parameter set has a check mark next to it, as shown on the menu in Figure 2.5.

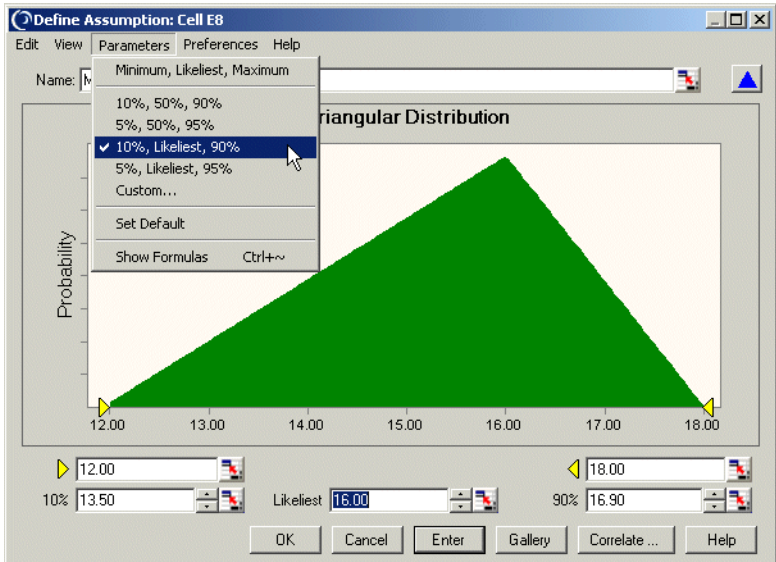

Figure 2.5 10th and 90th percentiles with Likeliest parameter

In addition to the standard parameter set, each continuous distribution's Parameters menu has additional pre-defined parameter sets that include various combinations of the standard parameters and percentiles. There is also a Custom command that lets you define your own parameter set.

If you choose Custom in the Parameters menu, you can replace any or all of the standard parameters with any percentile. You will always have the same number of parameters, either standard or alternate, for any given distribution. For example, even if you choose to use custom alternate parameters for a triangular distribution, you will always have three parameters, either minimum, likeliest, and maximum, or, for example, 10th percentile, likeliest, and 99th percentile.

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

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

Fitting distributions to data

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

How distribution fitting works

Glossary Term: **goodness-of-fit**— A set of mathematical tests performed to find the best fit between a standard probability distribution and a data set's distribution.

In distribution fitting, Crystal Ball automatically matches your data against each continuous probability distribution. A mathematical fit is performed to determine the set of parameters for each distribution that best describe the characteristics of your data. The quality or *goodness* of each fit is judged using one of several standard ***goodness-of-fit*** tests. The distribution with the highest ranking fit is chosen to represent your data.

You can review the distributions sorted in order of their fit tests using the comparison chart. This chart shows the fitted distributions superimposed over your data so you can visually check the quality of the fits. Several chart preferences make it easier to pinpoint discrepancies in the fits. If desired, you can override the highest-ranking probability distribution with another one of your choice.

Crystal Ball Note: *Only continuous distributions are considered for distribution fitting. Continuous and discrete distributions are defined on page 245. Distribution fitting can also be used to check the characteristics of a forecast chart. See "Fitting a distribution to a forecast" on page 120 for more information.*

Use the Fit Distribution dialog to specify the source of your data, the distributions to be fitted, and the goodness-of-fit test to use. Each goodness-of-fit test is calculated for every distribution, but only the selected test determines how the distributions are ranked.

Using distribution fitting

To use distribution fitting:

1. **Select the cell where you want to create an assumption.**

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

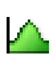

2. **Choose Define > Define Assumption.**

   The Distribution Gallery appears.

3. Click Fit to select the source of the fitted data.

The Fit Distribution dialog appears, as shown in Figure 2.6.

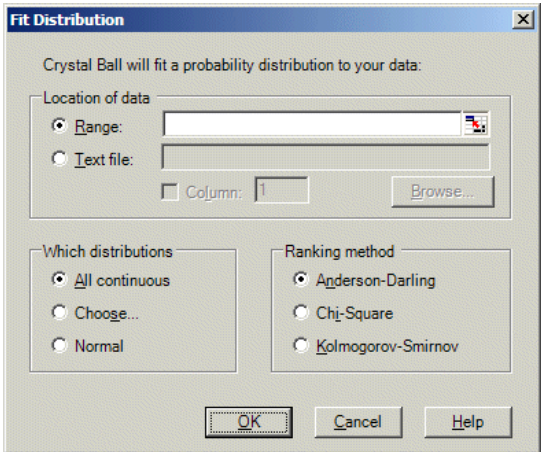

Figure 2.6 Fit Distribution dialog

4. Choose one of the following two options.

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

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

5. Specify which distributions are to be fitted:

- All Continuous fits the data to all of the built-in continuous distributions (these distributions appear as solid shapes on the Distribution Gallery).
- Choose displays another dialog from which you can select a subset of the continuous distributions to include in the fitting.

- The third option selects the continuous distribution that was highlighted on the Distribution Gallery when you clicked the Fit button.

Crystal Ball Note: *If you try to fit negative data to a distribution that can only accept positive data, that distribution will not be fitted to the data.*

6. **Specify how the distributions should be ranked.**

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

   - **Anderson-Darling**. This method closely resembles the Kolmogorov-Smirnov method, except that it weights the differences between the two distributions at their tails greater than at their mid-ranges. This weighting of the tails helps to correct the Kolmogorov-Smirnov method's tendency to over-emphasize discrepancies in the central region.
   - **Chi-Square**. This test is the oldest and most common of the goodness-of-fit tests. It gauges the general accuracy of the fit. The test breaks down the distribution into areas of equal probability and compares the data points within each area to the number of expected data points. The chi-square test in Crystal Ball does not use the associated *p-value* the way other statistical tests (e.g., *t* or *F*) do.
   - **Kolmogorov-Smirnov**. The result of this test is essentially the largest vertical distance between the two cumulative distributions.

7. **Click OK to fit the distributions to your data.**

   Crystal Ball successively fits the selected distributions to your data. The fitted distributions appear in the Comparison Chart dialog, starting with the highest-ranked distribution down through to the lowest.

   You can use the Next and Previous buttons to scroll through the fitted probability distributions. Each probability distribution is shown superimposed over the data, as shown in Figure 2.7.

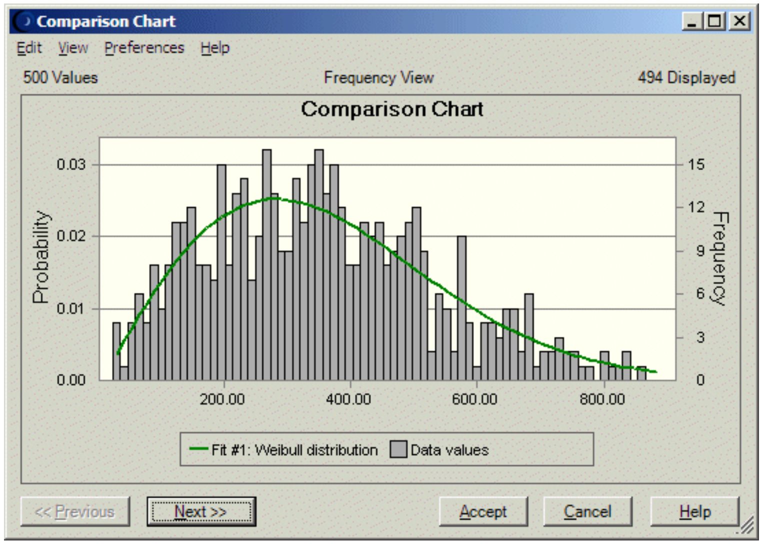

Figure 2.7 Comparison Chart dialog

8. **Use the Comparison Chart dialog to visually compare the quality of the fits or to view the goodness-of-fit statistics.**

Use the Comparison Chart features as described below:

- Click the Next and Previous buttons to scroll through the fitted distributions. You can view the quality of each fit graphically and statistically in decreasing order.
- Use the View menu to change the chart view. Choose Goodness of Fit to display results of the different goodness-of-fit tests for each distribution type.

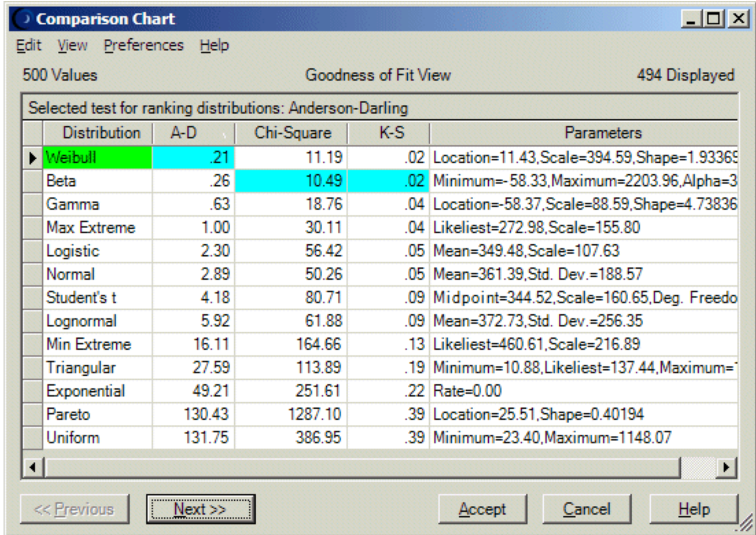

Figure 2.8 The Comparison Chart Goodness of Fit view

- Choose Preferences > Chart to change chart features so that similarities or differences are more clearly accentuated.
- Click Cancel to return to the Fit Distribution dialog.

9. **To use the currently selected distribution, either the best fit or another of your choice, click Accept.**

   The Assumption dialog appears with the best-fitting parameters taken from the chosen distribution as shown in Figure 2.9.

   You can change the distribution parameters before you click OK.

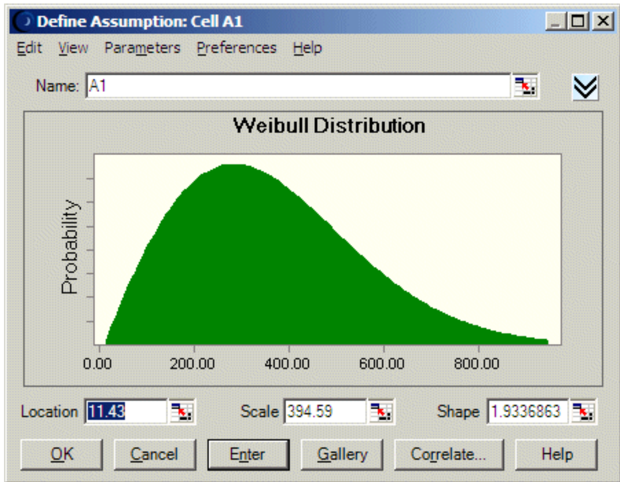

Figure 2.9 The best-fitting distribution after acceptance

A distribution fitting example

The following steps offer a specific example of fitting a distribution to data in a file.

1. **Create a new spreadsheet and select a cell.**

2. **Choose Define > Define Assumption.**

   The Distribution Gallery appears.

3. **Click Fit to select the source of the fitted data.**

   The Fit Distribution dialog appears as shown in Figure 2.6 on page 28.

4. **Choose the location of historical data.**

   For this example, click Text File and Browse and locate the TESTDATA.txt file in the Examples folder under the Crystal Ball installation folder, by default C:\Program Files\Decisioneering\Crystal Ball 7\Examples.

5. **Choose the distribution fitting characteristics.**

   For this example, use these settings:

   - For Which Distributions, choose All Continuous.
   - For Ranking Method, choose Anderson-Darling.

6. **Click OK to display the Comparison Chart dialog.**

For this Test Data example, the Weibull distribution had the best fit of any distribution using the Anderson-Darling fit test. The parameters that were calculated for the Weibull distribution are displayed in Figure 2.8 on page 31.

7. **Use the Next and Previous buttons to view comparison charts for the other distributions.**
8. **Return to Weibull and click Accept.**

   The Define Assumption dialog appears with the accepted Weibull distribution, as shown in Figure 2.9 on page 32.

Specifying correlations between assumptions

In Crystal Ball, assumption values are usually calculated independently of each other. Crystal Ball generates random numbers for each assumption without regard to how random numbers are generated for other assumptions.

However, dependencies often do exist between variables in a system being modeled. The Correlated Assumptions feature in Crystal Ball lets you enter correlation coefficients to describe the dependencies between assumptions. Correlation coefficients and correlation in general are explained in the "Statistical Definitions" chapter of the online *Crystal Ball Reference Manual*.

Crystal Ball Note: *Crystal Ball uses Spearman rank correlation to calculate correlation coefficients. For more information on how Crystal Ball calculates Spearman rank correlation coefficients, see "Rank correlation" in the "Statistical Definitions" chapter of the online Crystal Ball Reference Manual.*

When defining correlations, the more correlations you define, the greater the possibility that some correlations might be in conflict with each other, preventing Crystal Ball from running a simulation. Conflicts can arise when a group of assumptions are improperly related to each other by large positive and/or large negative correlation coefficients. When this condition occurs, the correlations are said to be "inconsistent." For more information, see "Correlated assumptions" on page 319.

With Crystal Ball, you use the Define Correlation dialog to specify a correlation coefficient for any pair of assumptions in the same workbook.

To define a correlation coefficient:

1. **Select the cell of one of the assumptions you want to correlate — for example, a cell describing the inflation rate.**

Statistical Note: *Which of the pair you select first is not important, since the correlation coefficient is bidirectional.*

2. Choose Define > Define Assumption.

The distribution previously defined for this assumption appears, as in Figure 2.10.

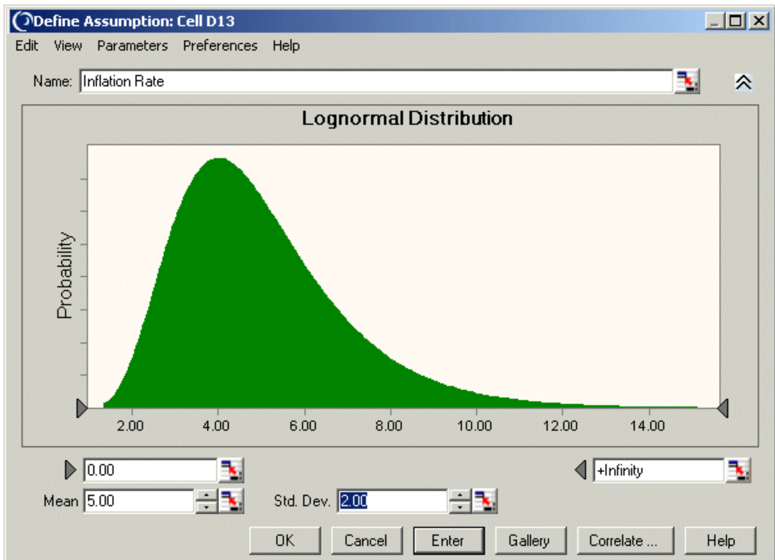

Figure 2.10 Lognormal distribution

3. Click Correlate.

The Define Correlation dialog appears, as in Figure 2.11.

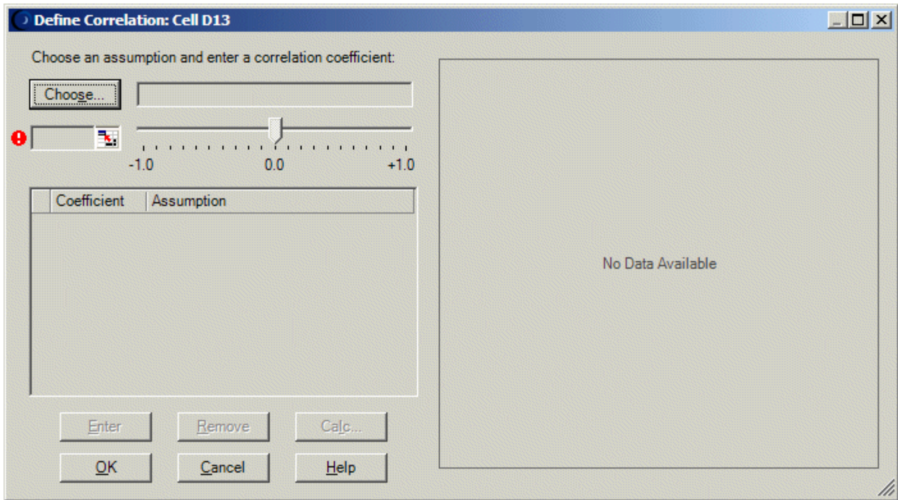

Figure 2.11 Define Correlation dialog

4. **Click Choose to select the second assumption from the Choose Assumptions dialog.**

Crystal Ball Note: *You can only correlate assumptions in the same workbook.*

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

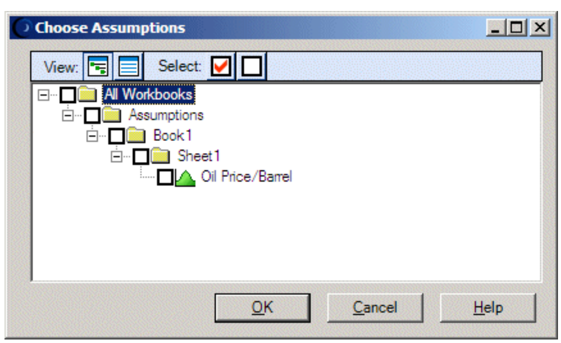

Figure 2.12 The Choose Assumptions dialog, Tree view

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

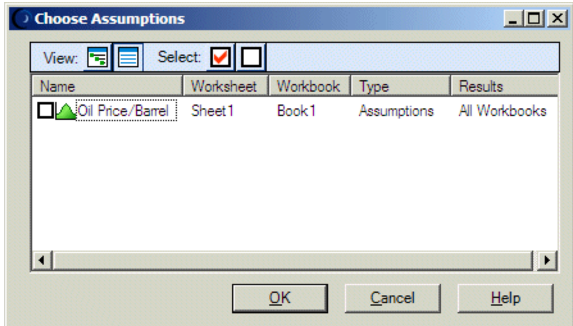

Figure 2.13 The Choose Assumptions dialog, List view

5. **Choose one or more of the assumption names on the list and click OK.**

   The cell reference or name of the assumption appears in the list in the left pane of the dialog, as shown in Figure 2.14.

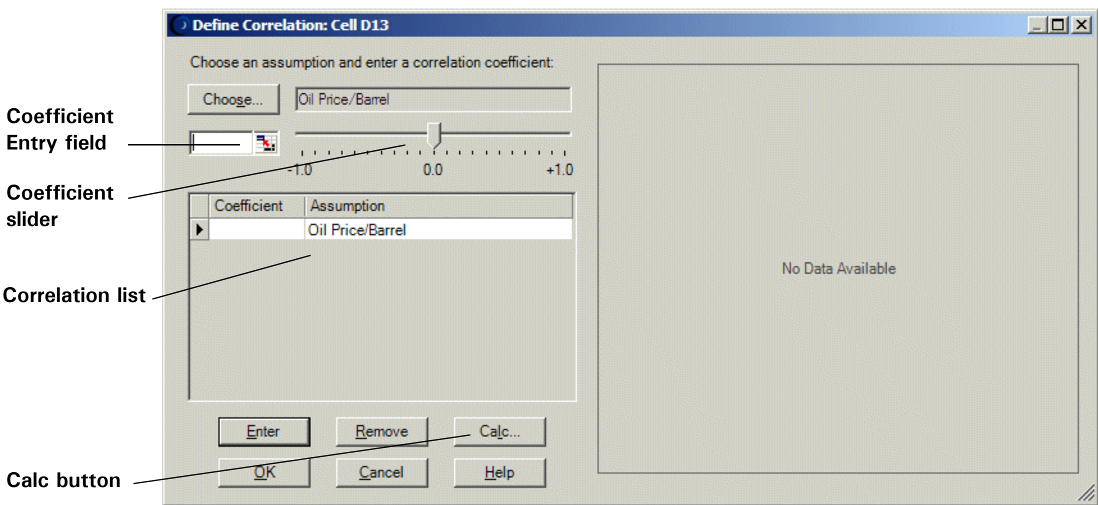

Figure 2.14 Define Correlation dialog with second assumption

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

The chosen assumptions appear in the list of correlations. The currently selected assumption also appears immediately next to the Choose button.

6. **Enter a correlation coefficient using *one* of the following methods:**

   - Enter a value between -1 and 1(inclusive) in the Coefficient Entry field.

     Type the number that you want to use in the field to the left of the slider control. After you type the number, the slider control on the correlation coefficient scale moves to the selected value.

   - Choose a cell that contains the correlation coefficient.

   Crystal Ball Note: *If you choose a cell with values that change during the simulation, it is the initial value of the cell that is used for the coefficient.*

   - Drag the slider control along the correlation coefficient scale.

     The value you select appears in the field to the left of the scale.

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

   - Click Calc.

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

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

     Crystal Ball calculates the correlation coefficient, enters it in the field to the left of the correlation coefficient scale, and moves the slider control to the correct position.

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

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

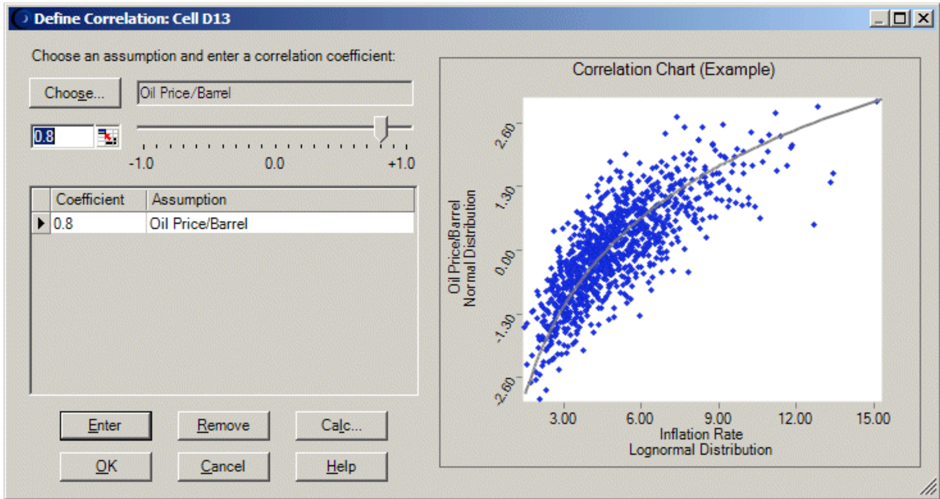

Figure 2.15 Correlation chart

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

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

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

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

Setting assumption preferences

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

Table 2.1 Preferences menu, Define Assumptions dialog

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

The Chart Preferences settings are discussed in "Setting chart preferences" beginning on page 125.

If you choose Assumption Preferences, the Assumption Preferences dialog appears as shown in Figure 2.16.

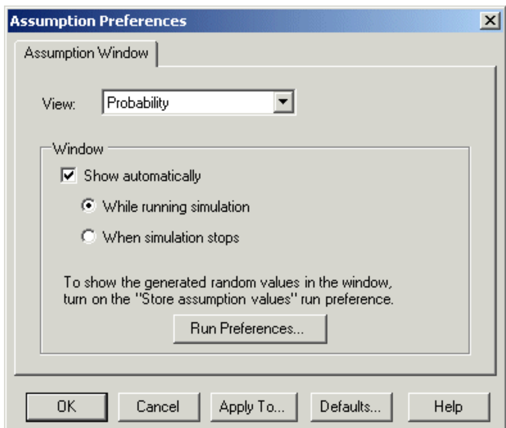

Figure 2.16 The Assumption Preferences dialog

This dialog lets you:

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

- Percentiles – shows a table of percentiles and their associated values for the assumption variable.

  Note: *For examples of each view, see "Changing the distribution view and interpreting statistics" beginning on page 115.*

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

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

For more information about the Assumption Preferences settings, click the Help button in the Assumption Preferences dialog.

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

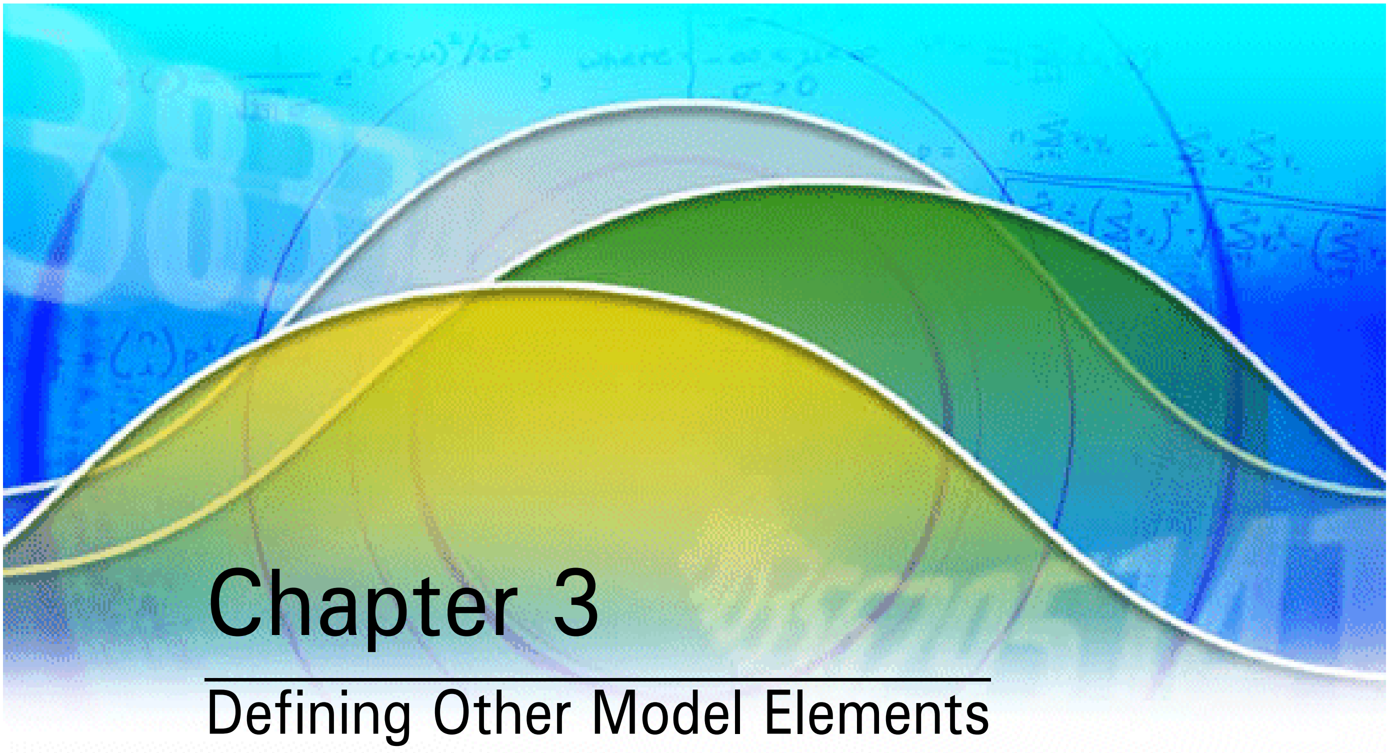

Chapter 3

Defining Other Model Elements

In this chapter

- Defining decision variable cells
- Defining forecasts
- Working with Crystal Ball data
- Setting cell preferences
- Saving and restoring your models

Chapter 2 describes how to start building a spreadsheet model by defining assumption cells. This chapter provides step-by-step instructions for completing models in Crystal Ball so you can run simulations against them. As you work through these instructions, you will learn to define decision variable cells and forecasts and how to cut, copy, and paste Crystal Ball data.

After you complete this chapter, read Chapter 4 for information on setting run preferences and running simulations.

Defining decision variable cells and forecast cells

Chapter 2 describes how to define assumption cells and use the Distribution Gallery. This chapter describes how to define decision variable cells and forecast cells and perform other tasks needed to complete a model definition.

Defining decision variable cells

Decision variables are not required for simulation models, but they can be useful when comparing and optimizing alternate scenarios. Several of the Crystal Ball tools discussed in Chapter 8, “Crystal Ball Tools,” use and benefit from decision variables. Decision variables are the variables that you can control, such as rent to charge or the amount of money to invest.

OptQuest Note: *You also use decision variables with OptQuest, available with Crystal Ball Professional and Premium Editions.*

To define a decision variable cell:

1. **Select a cell or range of cells.**

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

2. **Select Define > Define Decision.**

   The Define Decision Variable dialog appears.

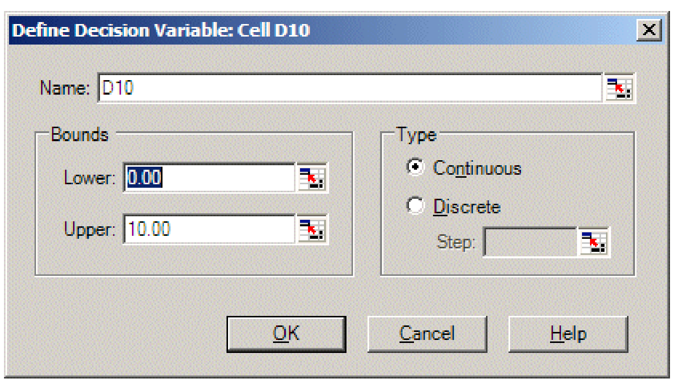

Figure 3.1 Define Decision Variable dialog

3. Complete the Define Decision Variable dialog:

- Name is the name of the decision variable.
- Bounds are the upper and lower limits for the decision variable range.
- Type defines whether the variable is continuous or discrete.
- Step specifies the interval between values for discrete variables. For example Step = 1 could be used to specify whole dollars, while Step = .5 could specify 50-cent increments.

For details about the fields and options in this dialog, click the Help button in the dialog.

Crystal Ball Note: *You can use cell referencing to name a decision variable, define the lower and upper limits, and set the step size. For more information, see "Entering cell references and formulas" on page 24.*

4. Click OK.

5. Repeat steps 1–4 for each decision variable in your model.

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

Defining forecasts

After you define the assumption cells and decision variables, you are ready to select and define forecast cells. Forecast cells usually contain formulas that refer to one or more assumption and decision variable cells. The forecast cells combine cells in your model to produce the results you need.

When you define a forecast cell, you:

- Name the forecast
- Specify the units for the forecast
- Indicate whether you want to automatically display the forecast window during the simulation
- Set precision control settings for the forecast
- Set filtering settings for the forecast
- Indicate whether to automatically extract statistics

To define forecast cells:

1. **Select a formula cell or a range of formula cells.**
2. **Select Define > Define Forecast.**

   The Define Forecast dialog appears.

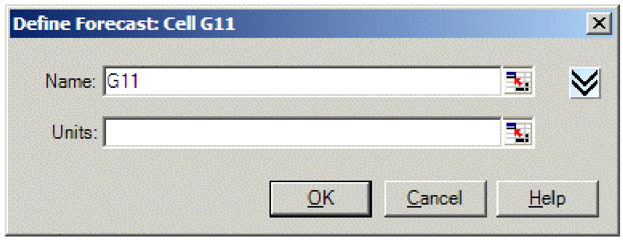

Figure 3.2 Define Forecast dialog

3. **Complete the Define Forecast dialog:**
   - Name is the name of the forecast.
   - Units is the name of the units that will appear at the bottom of the forecast chart, such as Millions.

   For details about the fields and options in this dialog, click the Help button in the dialog.
4. **To set additional forecast preferences, click the More button to expand the Define Forecast dialog.**

   The expanded Define Forecast dialog has four tabs of additional options and fields, listed in “Setting forecast preferences” on page 61.
5. **Click OK.**
6. **Repeat steps 1–4 for each forecast in your model.**

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

The Expanded Define Forecast dialog also has the Defaults button. Clicking on Defaults restores the original default settings in place of any new settings you have made. You can also click Apply To to use the settings in other charts and worksheets. For more information, click the Help button in the Apply To dialog.

Chapter 4

Running Simulations

In this chapter

- About Crystal Ball simulations
- Setting run preferences
- Running simulations
- Managing chart windows
- Saving and restoring simulation results
- Running user-defined macros

Chapters 2 and 3 describe how to define a spreadsheet model in Crystal Ball. This chapter provides step-by-step instructions for setting run preferences and running a simulation in Crystal Ball. After you complete this chapter, read Chapter 5 for information about analyzing simulation results.

About Crystal Ball simulations

Glossary Term:
spreadsheet model—
Any spreadsheet that represents an actual or hypothetical system or set of relationships.

Glossary Term:
simulation—
Any analytical method that is meant to imitate a real-life system, especially when other analyses are too mathematically complex or too difficult to reproduce.

Spreadsheet risk analysis uses both a ***spreadsheet model*** and ***simulation*** to analyze the effect of varying inputs on outputs of the modeled system.

After you define assumption, forecast, and decision variable cells in your spreadsheet model, you are ready to run the simulation. Crystal Ball uses a technique called Monte Carlo simulation to forecast the entire range of results most likely to occur in the situation you have defined in your spreadsheet model. This technique involves generating random numbers for the assumption variables. For more information on Monte Carlo simulation, see page 11.

While the simulation is running, Crystal Ball displays these assumption results in a forecast chart that shows the entire range of possible outcomes. Crystal Ball provides descriptive statistics for any forecast, which summarize the results numerically. The "Statistical Definitions" chapter of the online *Crystal Ball Reference Manual* discusses descriptive statistics.

How Crystal Ball uses Monte Carlo simulation

Glossary Term:
random number—
A mathematically selected value which is generated (by a formula or selected from a table) to conform to a probability distribution.

Glossary Term:
random number generator—
A method implemented in a computer program that is capable of producing a series of independent, random numbers.

Most real-world problems involving elements of uncertainty are too complex to solve analytically. There are simply too many combinations of input values to calculate every possible result. Monte Carlo simulation is an efficient technique that requires only a ***random number*** table or a ***random number generator*** on a computer.

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

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

This is an iterative process that continues until either:

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

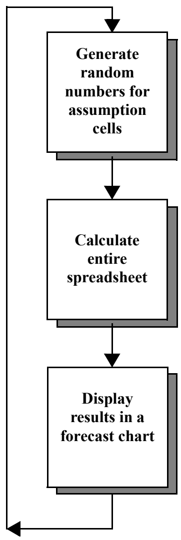

The final forecast chart reflects the combined uncertainty of the assumption cells on the model's output. Keep in mind that Monte Carlo simulation can only approximate a real-world situation. When you build and simulate your own spreadsheet models, you need to carefully examine the nature of the problem and continually refine the models until they approximate your situation as closely as possible. For more information about Monte Carlo simulations and Crystal Ball, see page 11.

Crystal Ball provides statistics that describe the forecast results. For more information on forecast results and statistics, see Chapter 5, "Analyzing Forecast Charts," and the "Statistical Definitions" chapter of the online *Crystal Ball Reference Manual*

Steps for running simulations

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

1. **Define assumptions (page 19), forecasts (page 59), and decision variable cells if appropriate (page 58).**
2. **If you want, customize the appearance of each cell as described in "Setting cell preferences" beginning on page 73.**
3. **Set run preferences as described in the next section, beginning on page 80.**
4. **Run the simulation as described in "Running simulations" beginning on page 90.**

The following sections explain each step.

Setting run preferences

You can change a number of run preferences, factors that determine how Crystal Ball runs a simulation.

Crystal Ball Note: *You need to reset the simulation before changing run preferences.*

To change any run preferences:

1. Select Run > Run Preferences to display the Run Preferences dialog.

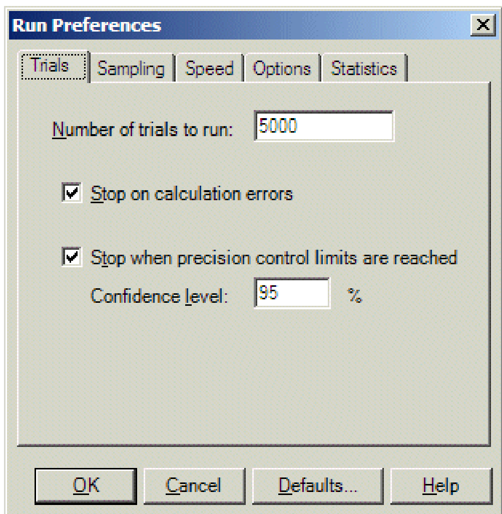

Figure 4.1 Run Preferences dialog, Trials tab

2. Click the tab with the preferences you want to change, described in the following sections:

- "Trials preferences" on page 81 — Specify when to stop a simulation, namely number of trials, calculation errors, and precision control.
- "Sampling preferences" on page 83 — Set the sampling seed value, method, and sample size.
- "Speed preferences" on page 84 — Determine whether a simulation runs in Normal, Demo, or Extreme speed (if available) and set additional speed control options.
- "Options preferences" on page 86 — Set a number of run preferences, including whether sensitivity data and assumption values are stored, whether assumption correlations are activated, whether

user macros are run, and whether the Crystal Ball Control Panel appears.

- "Statistics preferences" on page 87 — Determine how Crystal Ball displays percentiles.

For details on each tab, see the referenced pages or click the Help button in the Run Preferences dialog.

3. **Change any preferences on any tab.**
4. **Click OK.**
5. **To reset settings on the active tab to the original defaults, click Defaults.**

Trials preferences

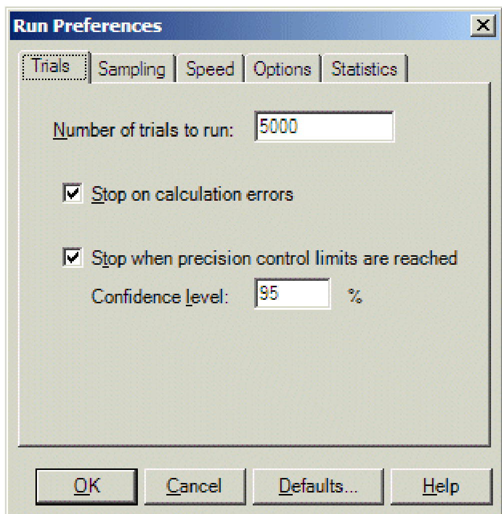

Figure 4.2 Trials tab, Run Preferences dialog

The Trials tab sets preferences that stop a simulation: number of trials, calculation errors, and precision control. For general instructions, see "Setting run preferences" beginning on page 80. Table 4.1 describes the settings available on this tab.

Crystal Ball Note: *The current simulation must be reset before precision control settings will take effect.*

Table 4.1 Trials tab, Run Preferences dialog

<table>
<tr><th>Setting</th><th colspan="2">Effect</th></tr>
<tr><td>Number Of Trials To Run</td><td colspan="2">Defines the maximum number of trials that Crystal Ball runs before it stops the simulation. If you check either of the checkboxes on this dialog, Crystal Ball only uses the maximum number of trials if forecast results do not meet the other stop criteria first.</td></tr>
<tr><td>Stop On Calculation Errors</td><td colspan="2">When checked, stops the simulation when a mathematical error (such as division by zero) occurs in any forecast cell. If a calculation error occurs, to let you find the error, Crystal Ball doesn’t restore the cell values.
If no calculation errors occur, the simulation continues until it reaches the Number Of Trials To Run or (if set) when the specified precision is reached.</td></tr>
<tr><td rowspan="2">Stop When Precision Control Limits Are Reached</td><td colspan="2">When checked stops the simulation when certain statistics reach a specified level of precision. You choose the statistics and define the precision that triggers this option in each Define Forecast dialog. For instructions, see “Precision preferences” on page 64.
Any forecasts set to use precision control must all reach their specified precision within the confidence level to stop the simulation.
If all the forecasts set to use precision control don’t meet the specified precision, the simulation stops when it reaches the Number Of Trials To Run.
By default, precision control is on.</td></tr>
<tr><td>Confidence Level</td><td>Sets the precision level (confidence level) that indicates when to stop a simulation.</td></tr>
</table>

Sampling preferences

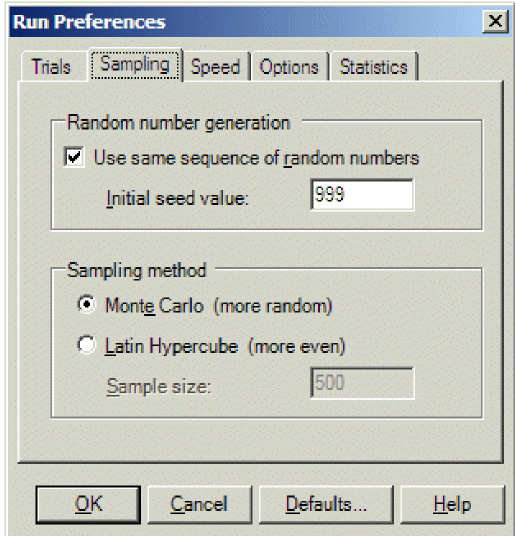

Figure 4.3 Sampling tab, Run Preferences dialog

The Sampling tab sets the sampling seed value, sampling method, and sample size. For general instructions, see "Setting run preferences" beginning on page 80. Table 4.2 describes the settings available on this tab.

Table 4.2 Sampling tab, Run Preferences dialog

Setting	***Effect***
Use Same Sequence Of Random Numbers	Sets the random number generator to generate the same set of random numbers for assumptions, letting you repeat simulation results. When you select this option, enter an integer seed value in the Initial Seed Value field.
Initial Seed Value	Determines the first number in the sequence of random numbers generated for the assumption cells (integer). ***Crystal Ball Note:*** *To reproduce the sample results shown in this manual, check Use Same Sequence... and use a seed value of 999.*
Sampling Method	Indicates whether to use Monte Carlo or Latin hypercube simulation. Latin hypercube sampling generates values more evenly and consistently across the distribution, but requires more memory. For more information, see page 317.

Table 4.2 Sampling tab, Run Preferences dialog (Continued)

Setting	Effect	
	Sample Size	For Latin hypercube sampling, divides each distribution into the specified number of intervals (bins). A higher number increases the evenness of the sampling method, while reducing the randomness. For more information, see "Simulation accuracy" on page 314.

Speed preferences

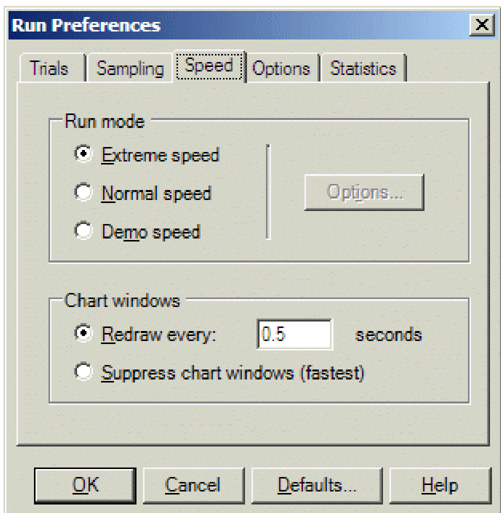

Figure 4.4 Speed tab, Run Preferences dialog

The Speed tab sets run mode and adjusts how fast a simulation runs. Extreme speed is the default simulation speed for Crystal Ball Professional Edition and Crystal Ball Premium Edition. Crystal Ball Standard Edition can only run in Normal or Demo speed. If you choose Normal or Demo speed, the Options button is active and you can make additional settings, described in Table 4.3.

Crystal Ball Note: *If you are using Crystal Ball Professional Edition or Crystal Ball Premium Edition, be sure to read Appendix C, "Using the Extreme Speed Feature," for important information about Extreme speed.*

For general instructions, see "Setting run preferences" beginning on page 80. Table 4.3 describes the settings available on this tab.

Table 4.3 Speed tab, Run Preferences dialog

Group	Setting		Effect
Run Mode	Determines overall simulation speed.		
	Extreme Speed		Available only in Crystal Ball Professional Edition and Crystal Ball Premium Edition. This option runs simulations up to 100 times faster than Normal mode but is not suitable for some models. For details, see Appendix C, “Using the Extreme Speed Feature.”
	Normal Speed		The standard simulation option for general model processing.
	Demo Speed		Runs simulations in "slow-motion" to make it easier to watch values change in spreadsheet cells and charts.
Options	Sets update rules for the active worksheet (Normal and Demo speeds only):		
	For Normal speed:	Update Every Trial	Updates Crystal Ball data in Excel after each simulation trial. Dynamic references are still updated internally if another setting is chosen.
		Update Every _ Seconds	Defines the update rate in terms of time. The default value is 1.
		Minimize Workbooks (Fastest)	Minimizes the Excel window. This option produces the fastest simulations.
	For Demo speed:	Maximum Number of Trials/Second	With optimal processing, what is the greatest number of trials to run each second. The default value is 10.
Chart Windows	Sets the redraw rate for any charts open during a simulation:		
	Redraw Every _ Seconds		Defines the redraw rate in terms of time. The default value is 0.5.
	Suppress Chart Windows (Fastest)		Closes all charts during simulation. Selecting this option overrides the Show Window preferences set for any charts. This option produces the fastest simulations.

Options preferences

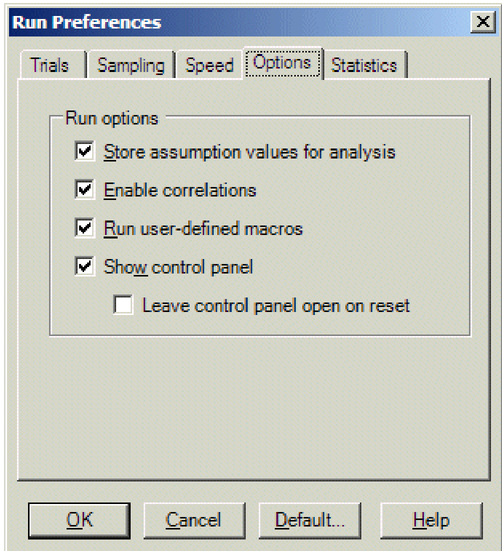

Figure 4.5 Options tab, Run Preferences dialog

The Options tab sets a number of run preferences. For general instructions, see "Setting run preferences" beginning on page 80. Table 4.4 describes the settings available on this tab.

Table 4.4 Options tab, Run Preferences dialog

Setting	*Effect*
Store Assumption Values For Analysis	Stores the randomly generated values used during the simulation for display while running. To see those values during a simulation, you must select this option and the appropriate Window preference for each assumption you want displayed. Values can also be exported to a spreadsheet after the simulation using the Extract Data command. This setting also allows Crystal Ball to generate sensitivity data during simulations. This information appears in the sensitivity chart to display the influence each assumption has on a particular forecast. For more information on the sensitivity chart, see "Understanding and using sensitivity charts" on page 161. The sensitivity chart is not available unless you select this option before you run a simulation.
Enable Correlations	Activates any defined correlations between assumptions.
Run User-Defined Macros	Runs any user-defined macros as part of the simulation process. For details, see "Running user-defined macros" beginning on page 99.

Table 4.4 Options tab, Run Preferences dialog (Continued)

Setting	Effect	
Show Control Panel	When checked, activates the Crystal Ball Control Panel.For more information, see “The Crystal Ball Control Panel” on page 93.	
	Leave Control Panel Open On Reset	When checked, continues to display the Control Panel after a simulation is reset.

Statistics preferences

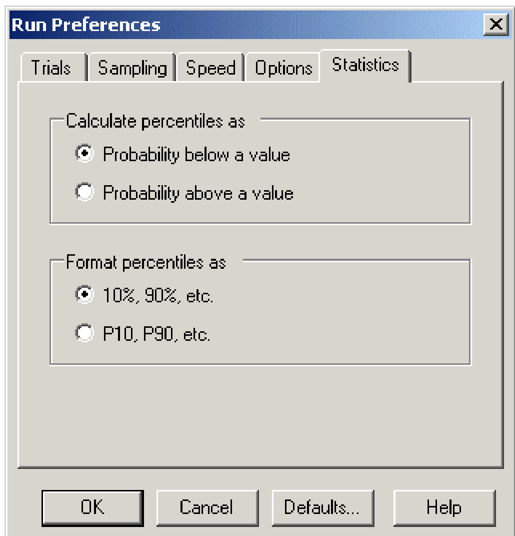

Figure 4.6 Statistics tab, Run Preferences dialog

The Statistics tab settings determine how Crystal Ball displays percentiles. For general instructions, see “Setting run preferences” beginning on page 80. Table 4.5 describes the settings available on this tab.

Table 4.5 Statistics tab, Run Preferences dialog

Setting	Effect	
Calculate Percentiles As	Sets how Crystal Ball defines percentiles. Selecting either of these options also affects the percentiles used for the assumption alternate parameters.	
	Probability Below A Value	Defines percentiles as the percent chance (probability) that the associated variable value is at or below a particular value. This setting is the default.
	Probability Above A Value	Defines percentiles as the percent chance (probability) that the associated variable value is at or above a particular value.
Format Percentiles As	Sets how Crystal Ball displays percentiles in charts and reports, with a percent sign or the percentile preceded by "P."	

Freezing Crystal Ball data cells

The Freeze command excludes or "freezes" certain assumptions, decision variables, or forecasts from a simulation. The Freeze command lets you:

- Temporarily disable cells that have been defined, but are not wanted for the current simulation.
- See the effect of changing certain assumptions or decision variables while holding other cells to their spreadsheet values.

Crystal Ball Note: *The Freeze command can be especially useful when you have multiple workbooks open and don't want to include all of their data cells in a simulation. You can freeze any unwanted cells instead of closing the workbooks that contain them.*

To freeze cells in a Crystal Ball model:

1. **Choose Run > Freeze.**

   The Freeze dialog appears as shown in Figure 4.7.

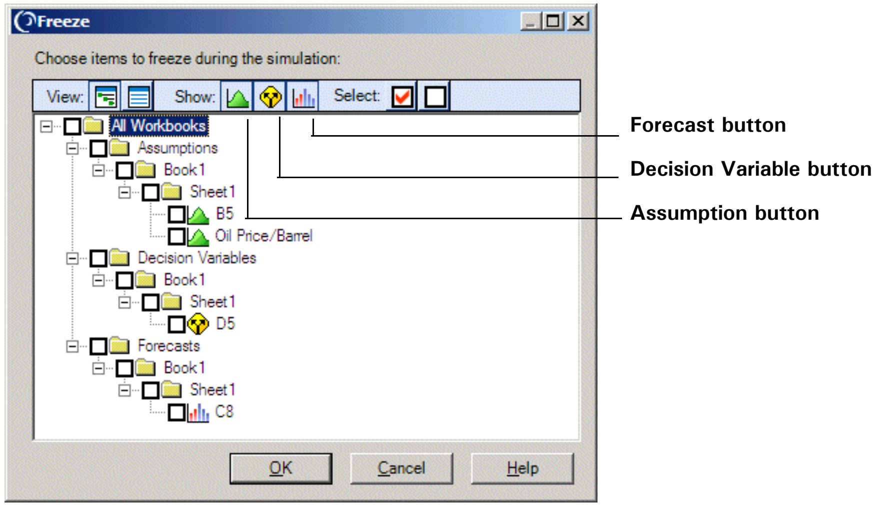

Figure 4.7 Freeze dialog, Tree view

All Crystal Ball data cells defined for the current simulation appear in the dialog.

By default, the dialog appears in a hierarchical Tree view. You can:

- Click Show: Assumptions, Decision Variables, or Forecasts to hide all assumptions, decision variables, or forecasts, respectively. Each button is a toggle; when clicked again, it shows that type of data cell.

- Click Select: All to select or check all cells currently showing in the tree.

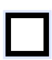

- Click Select: None to select no cells, that is, to uncheck all data cells currently showing in the tree.

You can use these buttons together. For example, to select all assumptions, you can click Show: Decision Variables and Forecasts to hide all decision variables and forecasts. Then, click Select: All. Since only assumptions are visible, only assumptions are selected for freezing.

Click on a workbook or worksheet box to select or deselect all data cells in that folder.

You can also click View: List to change from Tree view to List view, shown in Figure 4.8.

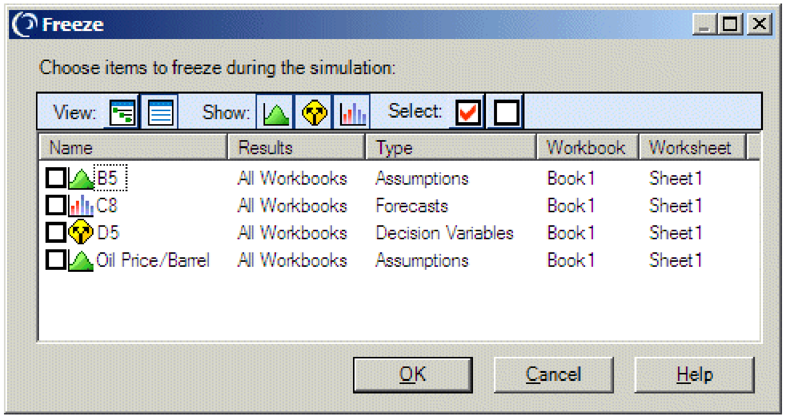

Figure 4.8 Freeze dialog, List view

2. **Check the boxes in front of cells you want to freeze.**
3. **When all data cells to be frozen are selected, click OK.**

Running simulations

About running simulations

When running a simulation in Crystal Ball, you can stop, continue, or reset the simulation at any time. You also have the option of observing each individual trial to see how the values change. You can display a chart for each assumption or forecast cell or run the simulation with all chart windows closed. You can select windows, cascade windows, open various windows at the same time, or bring the spreadsheet or a chart window to the front of the other windows.

While the simulation is running, Crystal Ball creates a forecast chart for each forecast cell using frequency distributions. A frequency distribution shows the number or frequency of values occurring in a given group interval (bin).

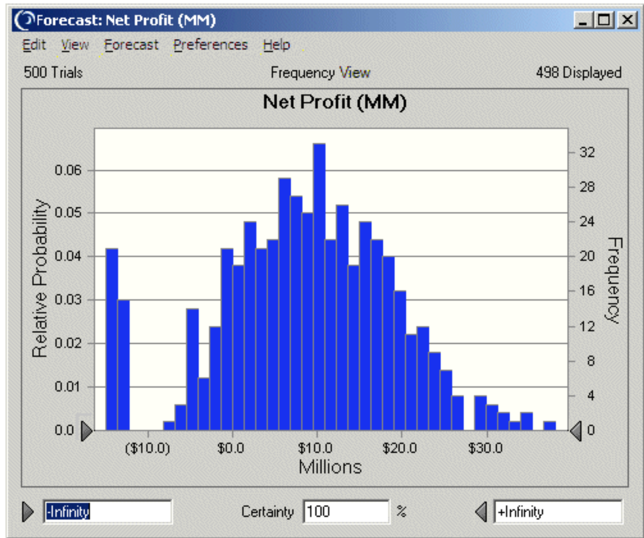

Figure 4.9 Net Profit Forecast dialog

In Chapter 5, you can find a more detailed explanation on how to interpret the forecast chart. For now, remember that the chart displays the forecast results, showing you how the forecast values are distributed and the likelihood of achieving a given result. As you run the simulation, observe the shape of the distribution to get a feel for the problem you are trying to solve. If you are new to Crystal Ball, try running simulations in Demo speed to watch the forecast values change from trial to trial.

In Figure 4.9 (frequency distribution for Net Profit), for example, you see that most of the values are clustered around central group intervals. You also see the variation around these group intervals and a clump of values at the left end. Chapter 5 also explains how to focus on a particular range of forecast values.

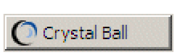

Crystal Ball Note: *Clicking on the Excel menu bar brings Excel forward and makes the forecast windows disappear. If this happens, you can quickly bring the forecast windows back to the front by clicking the Crystal Ball and Excel icons in the Windows task bar.*

During the simulation, Crystal Ball saves the forecast values in a list that grows as the simulation progresses. Beginning on page 186, you will learn how to export these forecast values to other programs, such as a statistical analysis program.

Running a simulation

To run a simulation:

1. **Select Run > Start Simulation or click the Start Simulation button in the Crystal Ball toolbar.**

   If you set the forecast to appear in the Define Forecast dialog, a forecast window appears. As the simulation proceeds, the forecast chart updates to reflect the changing values in the forecast cell.

Crystal Ball Note: *When you start Crystal Ball, the Crystal Ball Control Panel appears as shown in Figure 4.10 on page 94. You can use the Crystal Ball Control Panel to perform many of the procedures described in this section. For details, see "The Crystal Ball Control Panel" on page 93.*

Stopping a simulation

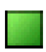

To stop a simulation:

1. **Select Run > Stop Simulation or click the Stop button on the Crystal Ball toolbar or Control Panel.**

Continuing a simulation

To continue a simulation:

1. **Select Run > Continue Simulation or click the Continue button on the Crystal Ball toolbar or Control Panel.**

Resetting and rerunning a simulation

To reset the simulation and start over again:

1. **Select Run > Reset Simulation or click the Reset button on the Crystal Ball toolbar or Control Panel.**

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

2. **Click OK.**

   Crystal Ball resets the number of trials to 0 and clears the list of values and statistics for each assumption and forecast. However, the assumption and forecast definitions remain.

3. **Change any assumptions or forecasts as needed.**

4. **Select Run > Start Simulation or click the Run button on the Crystal Ball toolbar.**

   The simulation starts over again.

Single-stepping

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

To observe an individual trial:

1. **Select Run > Reset Simulation or click the Reset button on the Crystal Ball toolbar or Control Panel.**

2. **Select Run > Single Step or click the Single-Step button to run one trial of the simulation. Click the button again to run another.**

The Crystal Ball Control Panel

You can use the Crystal Ball Control Panel to perform many simulation and analysis commands. When you start the first run as described in “Running a simulation” on page 92, the Crystal Ball Control Panel appears as shown in Figure 4.10 on page 94.

Once it appears, you will find its buttons and menus convenient to use for stopping, continuing, resetting, single-stepping, and analyzing the results.

Run preference settings determine whether the Control Panel appears or is hidden when a simulation runs. To hide the Control Panel, uncheck Show Control Panel on the Options tab of the Run Preferences dialog (click the Run Preferences tool or choose Run > Run Preferences).

Crystal Ball Note: *By default, the Control Panel stays open following a reset. To close it, uncheck Leave Open On Reset on the Options tab of the Run Preferences dialog. For more information on the Run Preferences Options tab, see “Options preferences” on page 86.*