
794. GETTING STARTED WITH CRYSTAL BALL

Learning Objectives

- Develop simulation models in MS Excel with a Monte Carlo simulation add-in program

STARTING CRYSTAL BALL®

Do one of the following to launch Crystal Ball:

- Select Oracle Crystal Ball from the list of Programs in the Start menu or list of apps
or
- Double-click the Oracle Crystal Ball icon on the desktop (if one exists there)


A fresh instance of Excel will open first, and then Crystal Ball will be launched within Excel. You will see a blank Excel workbook, and a Crystal Ball ribbon tab. Click on that tab.

It is best to have only one model open at a time. Running a simulation runs simulations in all open models.

SPECIFYING AN INPUT DISTRIBUTION (ASSUMPTION)

Cells that contain individual distributions are called *assumption* cells in CB.

1. Move to an empty worksheet cell.
2. Enter a value in the cell, typically the mean or EV value for your distribution. If no value is entered, a zero will automatically be entered once you Define Assumption.
 - Entering a value identifies the cell as a value cell, as opposed to a label or formula cell where assumptions are not allowed.
 - Note that the displayed cell value is not updated for changes made in the Assumption definition in CB. You may overwrite the value without changing the model.
3. Click the upper part of the <Define Assumption> icon to bring up the Distribution Gallery, or click just below the icon to show a drop-down menu of available distributions.
4. In the Distribution Gallery, double click on the distribution type that you want, or select from the drop-down menu.



5. A preliminary distribution is displayed. Crystal Ball has built-in defaults for most distributions. Data boxes and handles allow you to specify and alter the distribution.
 - Enter an assumption name. Excel names do not apply, though it is convenient to make them the same or similar.
 - Enter values in the parameter input boxes.
 - Instead of values, cell references can be placed in the parameter boxes.
 - Click on the <Enter> button to recompute and display the distribution.
 - You can truncate the upper and lower tails, either by dragging the triangle-handles on the chart, or by entering truncating values directly in the shown boxes.
If you are not seeing the range endpoints, click the Show More (^) button at the upper-right.
6. Click on <OK> to accept and return to the worksheet.

It is nice to have the mean displayed on a distribution figure. Change the setting within the Define Assumption window by:

1. Click Preferences, Chart, Chart Type.
2. From the list, check Mean.
3. Click <Apply to>
4. Click radio buttons for “All tabs” and “All open and new workbooks.”
5. Click <OK>’s to save and close the settings.

SPECIFYING OUTPUT CELLS (FORECASTS)

In designing a Monte Carlo model, it is important to anticipate what data need to be captured for analysis. Predetermine what calculated (output) values are needed at each trial in the simulation run. The output cells are called *forecasts* cells by CB.



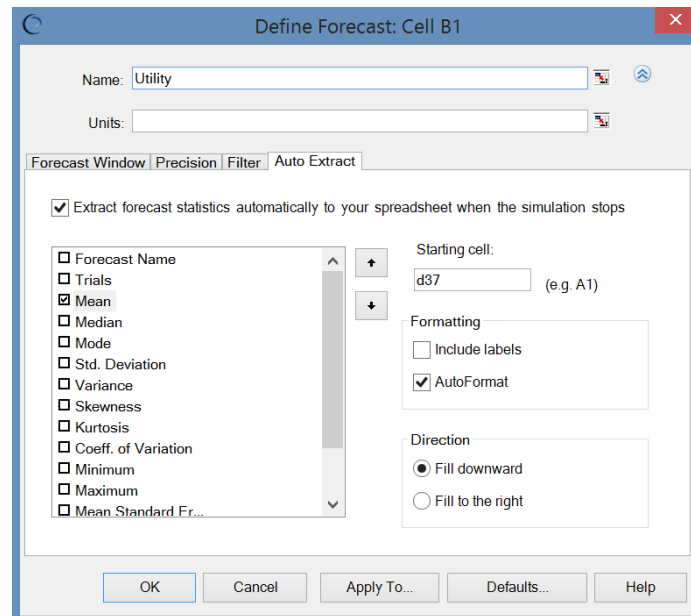
1. Position the cursor at an output formula cell, which must contain a formula.
2. Click the <Define Forecast> icon to bring up the Define Forecast popup window
3. Enter a variable name and units (for graphs and CB references). CB names are distinct from Excel named cells; it is convenient to make the names the same or similar.

Recommendations:

- Under the Forecast Window tab, you may want to uncheck the box next to <Show automatically> if you have many forecasts cells. This will help avoid a many open forecast windows when

you run a simulation. You can always see the results of a forecast cell by selecting the View Charts button.

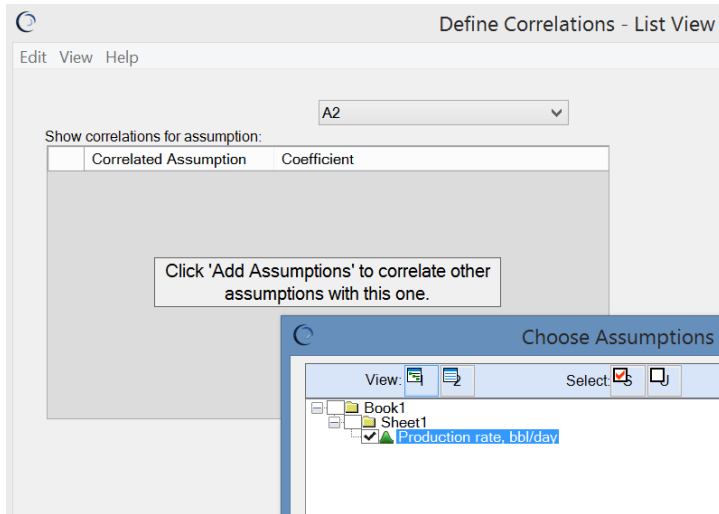
- For one key variable, such as *NPV*, leave its forecast window checked for Show automatically and for “While running simulation.” This is useful to monitor the simulation.
- It is helpful to export a forecast cell’s mean to the spreadsheet. As examples, the mean of the *NPV* cell is *EMV*, and the mean of the *Utility* cell is *EU*. Do this by
 1. Select the Auto Extract tab.
 2. Check the box next to Extract forecast statistic. Click <OK> to the warning.
 3. Select the statistic(s) you wish to extract.
 4. Type in the reference to the destination Starting Cell (unfortunately, selecting a cell in the worksheet doesn’t work).
 5. Uncheck Include Labels unless you want them.
 6. Click the button for whether you want the cells filled to the right or downward. Be careful that this export will not overwrite cells used in your model.



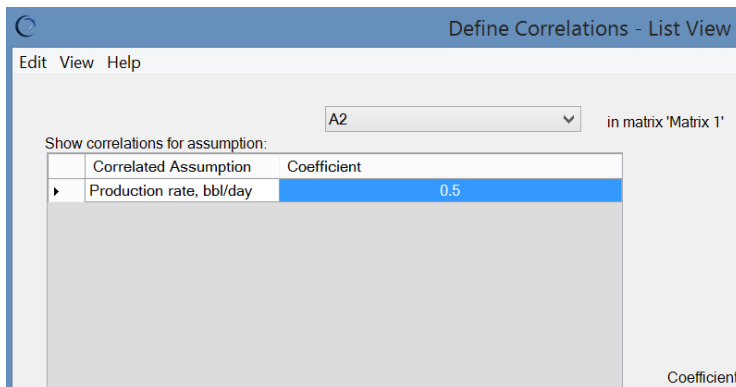
CORRELATED DISTRIBUTIONS

Often, a change in one variable is associated, or correlated, with a change in another variable. This occurs because (1) one variable is partially dependent upon another or (2) two variables each have a common partial dependency upon a third variable. The cause of the association, called *correlation*, is unknown but has been identified by data observations.

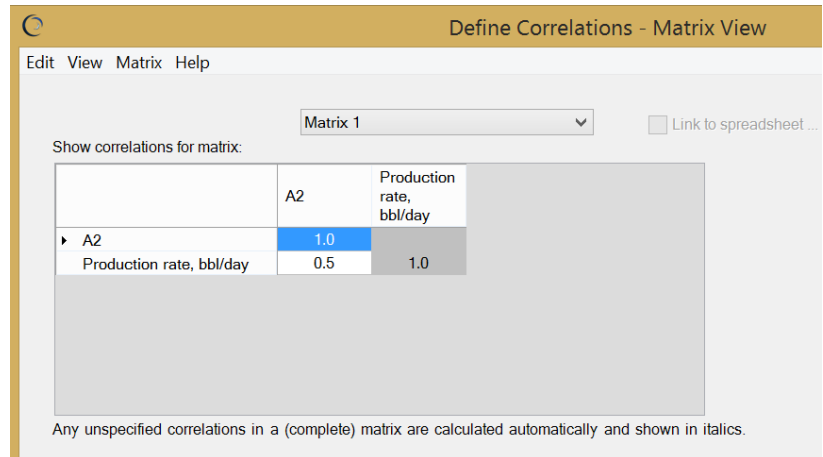
1. Define two or more variables as Assumptions (the order of each pair is not important, since the correlation coefficient is bi-directional).
2. While in the Define Assumption window, click the Correlate button. A Define Correlation window appears.
3. You may enter correlations in list view or matrix view. You can choose between the two views under the View menu. Here, we will start with the List View approach.
4. Click <Add Assumptions> to specify the other variable(s) cells for correlation. Select the assumption(s) and click <OK>.



5. Specify the Spearman rank correlation coefficient. This can range from -1 to $+1$. Either extreme represents perfect positive ($+1$) or negative (-1) rank correlation.



You can switch to matrix view under the View menu.



You may also want to copy the matrix view to your spreadsheet. You can do this within the Edit menu, and then link to the spreadsheet cells from within the correlation matrix.

6. You can return to Step 4 to specify a correlation to another variable.
7. Click on <OK> to complete the action.

EDITING ASSUMPTIONS AND FORECASTS

If you want to copy, paste or clear Assumptions or Forecasts, use Crystal Ball's <Copy>, <Paste> and <Clear> buttons. These work about the same as Edit Copy and Paste functions in Excel.

However, the destination range, if more than one cell, needs to be selected before pasting.

- Excel's Copy and Paste functions **will not work** with Crystal Ball definitions. The cells look okay but the cells do not carry or clear the CB attributes. Be sure to use the CB functions for CB cells.
- Cell or range dragging does work with CB and Excel cells without problem.

RUN PREFERENCES

Your simulation run is controlled by parameters set by clicking the <Run Preferences> button.

Number of trials. 1000 is typical. Other defaults (boxes checked) are okay.

Sampling. Check "Use same sequence ..." and enter a seed number. Documentation is unclear about allowable values other than to say the seed should be an integer. Suggestion: use an integer in the range 1 to 999. Specifying the seed allows better comparison of alternate cases.

You can choose Monte Carlo or Latin Hypercube sampling, and enter the number of bins for Latin Hypercube. The number of trials is best when an integer multiple of the number of “bins.”

Speed. Extreme speed is very fast. CB incorporates a vectorization package from Frontline Systems. The spreadsheet is compiled and solved outside Excel. If you are having any calculation issues, you may want to check “Normal speed” instead.

Options. Typically leave checked the first five check boxes.

Statistics. For E&P, we suggest selecting “Probability above a value” and “P10, P90 etc.” Thus for *NPV* or reserves, where higher values are better, the P90 is the 90% confidence value.

RUNNING A SIMULATION



Before launching a simulation run, we suggest that you try single-stepping through the model. Every time you click on Single Step button, a new random realization is generated.

Examining outlier outcomes is a good way to find errors in models.


When you’re satisfied that the model is working, press the Start button to finish the simulation for the number of trials you specified

OPTIMIZING WITH OPTQUEST

We’ll demonstrate using the exercise described in section 832.

Set the number of trials to 1000, check that Latin Hypercube Sampling is turned on.

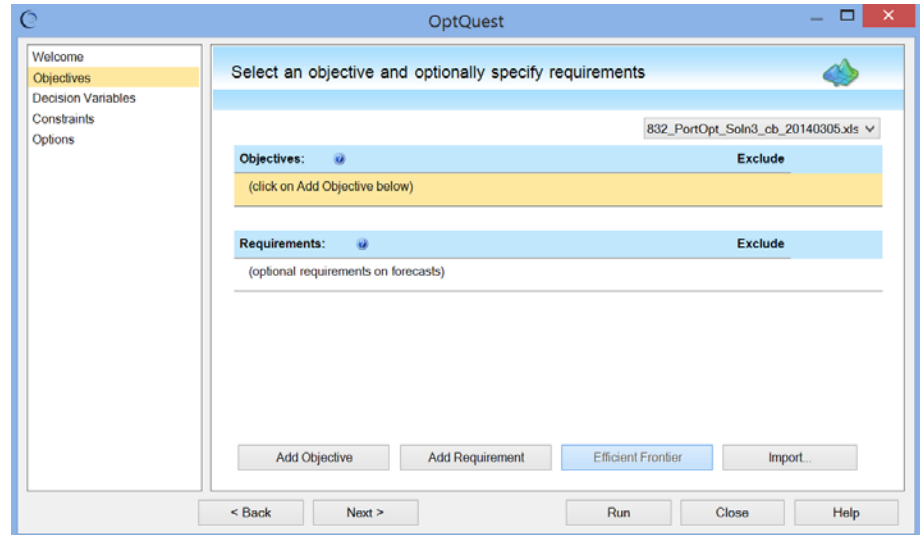
Set up the decision variables in Crystal Ball by selecting the working interests and pressing the Define Decision button on the Crystal Ball toolbar. Name the decision variables something like WI_1, WI_2, etc. The bounds will be 0 and 1, and this variable is continuous. You may want to put in reasonable starting values for these cells. Or, you may want OptQuest to work harder at optimizing and search broadly for candidate solutions.






You are now ready to use OptQuest. Open OptQuest by electing the OptQuest button from the Crystal Ball toolbar.

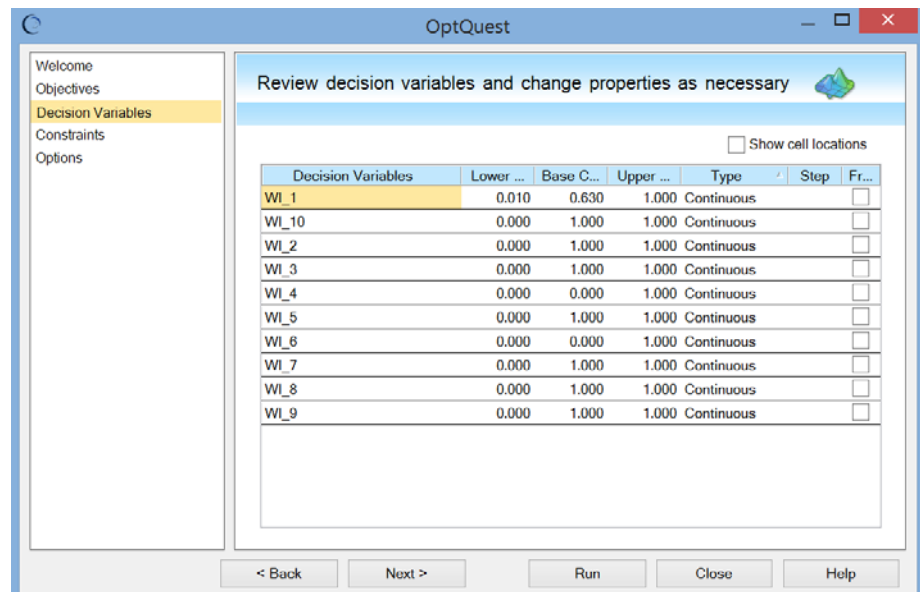
The OptQuest window appears. We want to add objectives first; click on the <Add Objectives> button at the bottom.



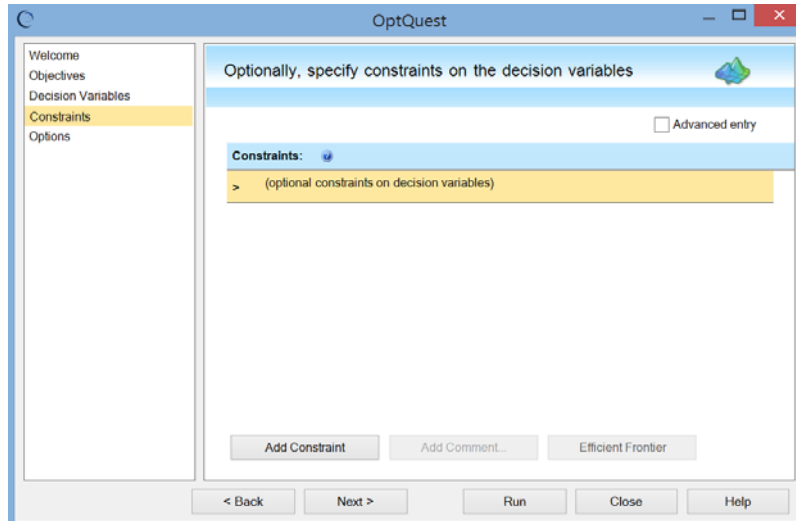
Choose to Maximize the Mean of NPV, in order to maximize EMV.




Requirements are restrictions on run statistics. For example, you may want to include solutions that replace reserves with a 90% confidence. Click the <Next> button at the bottom of the window, or the <Decision Variables> selection on the left side of the window. A list of the defined decision variables will appear.



Here, we can change properties of the decision variables if we desire. We will leave them as they are, so click the <Next> button at the bottom of the window, or the <Constraints> selection on the left side of the window. An opportunity to select constraints will appear.



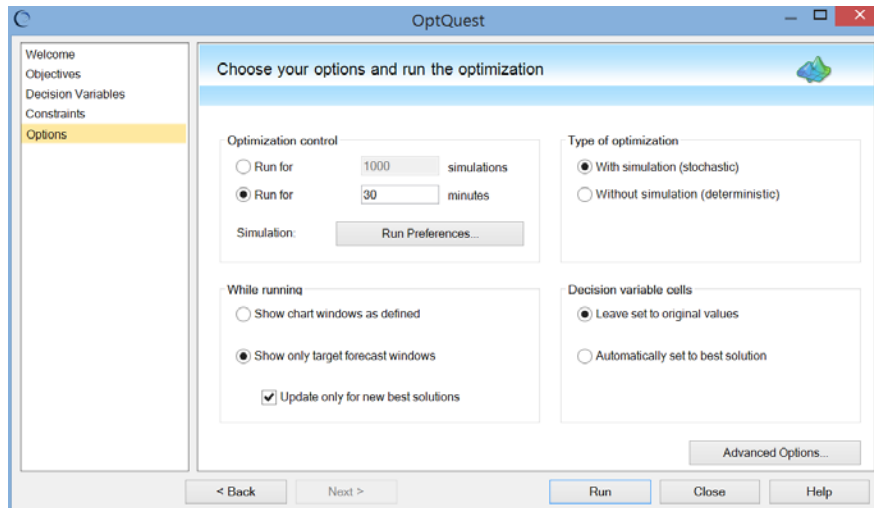
You may want to constrain decision variables. To add a constraint on the investment amount, click on the <Add Constraint> button at the bottom. Enter the cell reference for the Investment or Capital Budget (F31 in the provided file) in the space on the left, ensure that <= is selected, and enter the maximum budget in the space to the right.



You may want to set the constraint a little above the target budget to allow some flexibility.

Other types of constraints may also be assigned here. A possible use would be in allocating a budget among different investment types and requiring that the sum be total funds available.

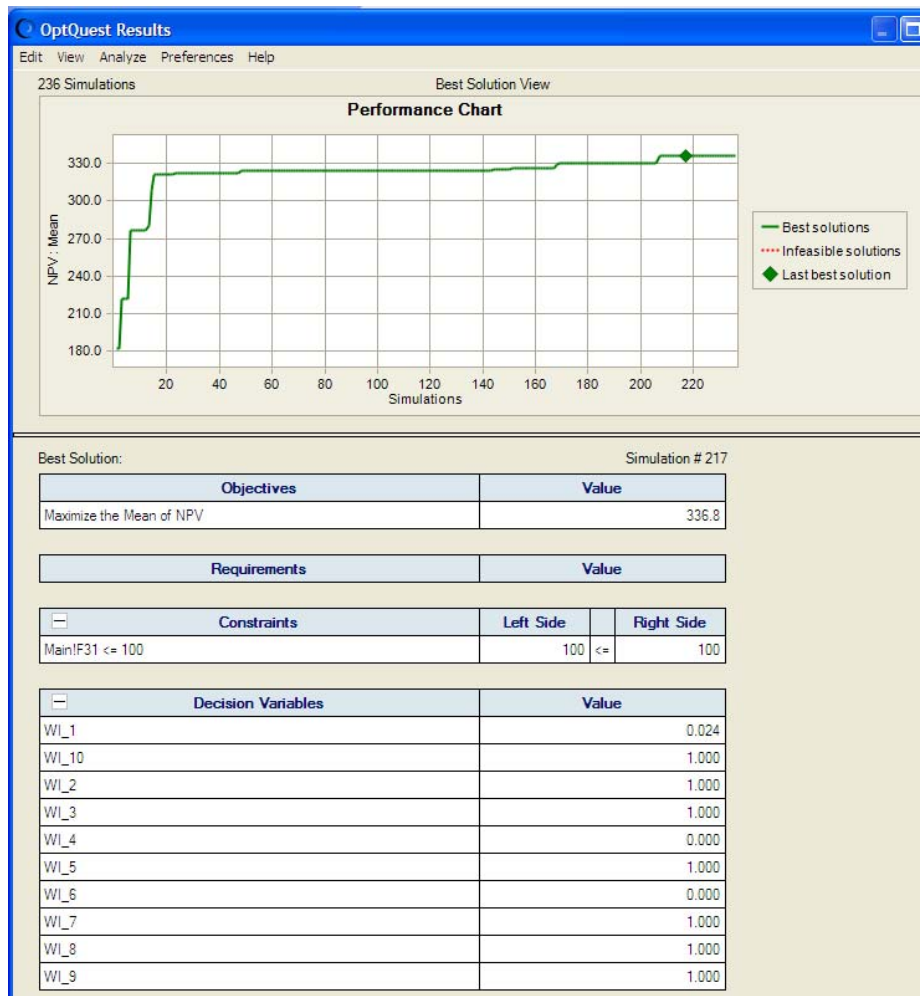
When finished, click the <Next> button at the bottom of the window, or the <Options> selection on the left side of the window. An opportunity to select optimization options will appear.



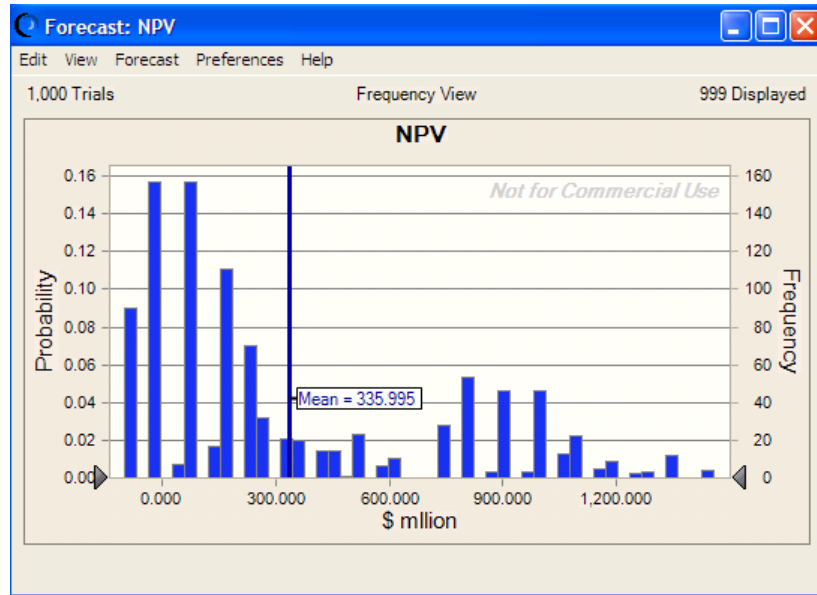
For single-point solutions, you can set a run time by specifying minutes. After that time is consumed, the run halts and you can inspect the performance graph to determine whether more time is needed to be confident of reasonable convergence. You can also set the run to stop after a set number of simulations.

Other options allow you to set simulation parameters by selecting the <Run Preferences> button, controlling the windows displayed while running, setting the type of optimization, and what to do with decision variable cells after optimization.

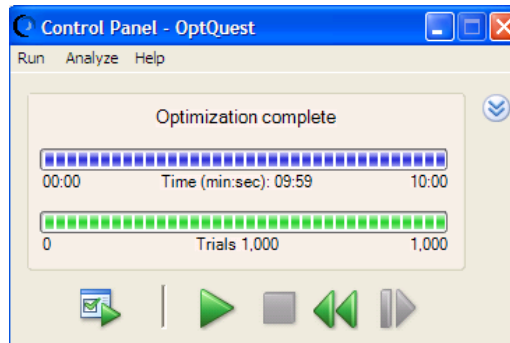
Once all parameters are set, click the <Run> button. The optimization will start, showing the progress as it runs. Once the optimization is finished, the window will show the best solution found to that point.



Once completed, you will also see some other results windows. The window below shows the distribution of NPV for the best solution.



To continue running the optimization, you can adjust the run preferences and then click the Start (play) button on the Control Panel.



Both OptQuest and RISKOptimizer will show somewhat different answers than a simulation run with the same decision values. This is influenced by choice of random number seed, number of trials (iterations), other program settings, and sampling order. As the number of trials gets larger, the answers from both tools should converge.

DETERMINING THE EFFICIENT FRONTIER

We want a portfolio that has both high value and low risk. Value is typically measured by *EMV* or the EV return¹ proxy.

Many measures have been used to represent risk, including but not limited to: variance, standard deviation, semi-variance, semi-standard deviation, test cost (e.g., dry hole cost), and maximum negative cashflow. In this exercise, we will use standard deviation, the most popular measure of risk for these analyses.

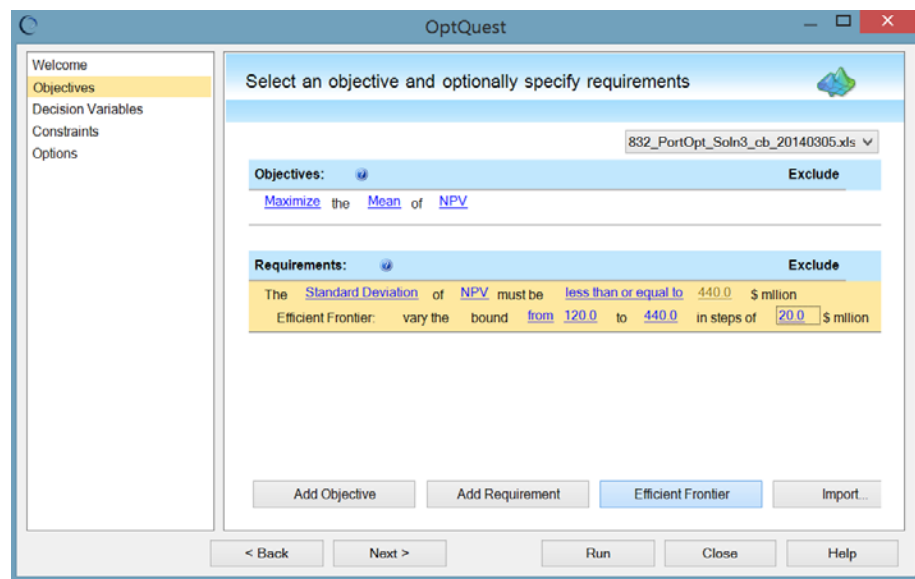
Because optimization calculations may require an hour or more, you may have time in class only to check or approximate the notebook solutions. With sufficient analysis and computer time, it is desirable to do most or all of the following:

1. Determine the portfolio and resulting *EMV* and *SD* if we maximize *EMV* without a capital constraint. This is a useful reference and what some companies call a “capability plan.”
2. Determine the portfolio and resulting *EMV* and *SD* if we maximize *EMV* with a \$100 million capital constraint. This is the upper-end of the efficient frontier with the candidate investments and this budget.
3. Determine the portfolio and resulting *EMV* and *SD* if we minimize *SD* (σ) with a \$100 million capital constraint (both ways; i.e., we must spend the \$100 million). This is the lower-, leftmost-end of the efficient frontier.
4. Maximize *EMV* versus *SD* by stepping through a range of *SD*s bounded by the answers in (2) and (3). This determines the upper side of the envelope of feasible portfolio solutions. This envelope is the *efficient frontier* and approximates the function of *EMV* vs. *SD*. OptQuest has an Efficient Frontier button to automate this in one calculation run.
5. Determine a graphic solution to approximate the optimal portfolio on the efficient frontier. Calculate and draw several iso-*CE* curves on the chart. Ideally we will have one iso-*CE* curve that touches the efficient frontier tangentially at one point. This is the optimal project mix for that given risk attitude.

¹ Most often, ‘return’ means internal rate of return (*IRR*). Using this as a proxy measure for value is especially suited to financial investment portfolios. This is okay for investments because the candidates are typically very liquid and unbounded as to investment amount. The value is the period’s profit, which equals the investment times the percentage return. For projects, especially those spanning partial or multiple years, *EMV* provides a better value measure.

To run the Efficient Frontier assessment in OptQuest:

1. Set any Objectives, as described in the previous section on running an optimization. For this example, set an Objective to maximize *EMV* (the Mean of *NPV*).
2. Now we need to set bounded increments of Standard Deviation so that the optimization can find the best solution within each bound for creating the Efficient Frontier plot. First, click the <Add Requirement> button at the bottom and select Standard Deviation of NPV must be less than 440.
3. Click the <Efficient Frontier> button at the bottom.
4. Set Efficient Frontier to vary the upper bound in steps; here it is set to vary from 120 to 440 in steps of 20 \$million.



Once all parameters are set, click the <Run> button. The optimization will start, showing the progress as it runs.

