UVA-QA-0593

CORRELATIONS IN CRYSTAL BALL

Crystal Ball has features that allow users to correlate the outputs from different assumptions. By using correlations,[1] the output of a specific assumption variable becomes dependent upon another assumption variable. For example, a model might contain two assumption cells: the value of the S&P 500 and the value of the NASDAQ. If these assumptions were treated as independent (i.e., no correlation), the model would be missing the connection between the movements of the markets. If one were simulating the risk in a portfolio that has components tied to each market, the amount of portfolio risk would be wrong unless correlation were considered. Correlation would be useful in the study of diversification, hedging, and other risk-management activities.

There are two ways to introduce correlations between assumptions: (1) using the assumptions dialog box and (2) using the correlation matrix tool.

1. *Correlations in the assumptions dialog box:* The main dialog box that opens when you define a Crystal Ball assumption has a button called "correlate." When you click the "correlate" button, a dialog appears in which all the assumptions of the model appear. When you select any of the other assumptions, you will be able to fill the correlation coefficient in the specified box. Notice that you can introduce the correlation coefficient directly, use a slide bar, or include a reference to a cell that contains the correlation. After setting the correlation coefficient (or referencing a cell that has the correlation number), push the "enter" button. This will instruct Crystal Ball to generate correlated trials for the assumptions when a simulation is run.

2. *Correlation matrix tool:* The previous method works when you are creating the correlation between a small number of assumptions, perhaps two or three. When the

[1]The correlation between two variables measures the extent to which the uncertain outcome of one tends to move together with the uncertain outcome of the other. Correlation can take values between –1 and 1. The sign of the correlation coefficient shows whether both variables move in the same direction (e.g., height and weight) or in the opposite direction (e.g., GDP per capita and percentage of people under the poverty income level). A correlation of 1 means that the variables move entirely together, and a correlation of 0 means that there is no relationship between the two variables—they are independent.

This technical note was prepared by Miguel Palacios under the supervision of Samuel E. Bodily, John Tyler Professor of Business Administration.

number of assumptions increases, manually inputting the correlation coefficients can become tedious. For such cases, Crystal Ball has a tool called "correlation matrix," which you can find in the CBtools menu at the top of your screen (see Figure 1). When you use the correlation matrix tool, Crystal Ball automatically refers the correlation coefficients to a matrix that it creates in your spreadsheet. The numbers that appear in the matrix are actually linked to the assumptions.

You may calculate correlations from data in Excel using the correlate function.

There are two caveats for using correlations in Crystal Ball:

1. You might input a set of correlation coefficients that is inconsistent. For example, an inconsistency would be to input that the S&P 500 has a .8 correlation with the NASDAQ, and that the NASDAQ has a .6 correlation with the Dow Jones Industrial Average (DJIA), and that there is 0 correlation between the S&P 500 and the DJIA. In that case, Crystal Ball will give you the option of adjusting the coefficients to consistent values or to cancel the correlation. If you choose to let Crystal Ball adjust the coefficients, it will ask you to determine whether you want to change the coefficients permanently or only for the current simulation. If you wish to continue using the matrix, choose to change the values *only for the current simulation*.

2. You cannot see the effect of the correlation using the single-step button. The correlation between the variables works correctly only for an entire simulation run. Thus, if you input a correlation coefficient of 1, in which both variables would move in lock-step fashion, you will notice that this does not happen on each single step. The correlation resulting from a simulation run, however, will show a correlation of 1 or very close to it.

Figure 1: Correlation Matrix Example

	Variable 1	Variable 2	Variable 3	Variable A	Variable B	Variable C
Variable 1	1.000	0.900	0.950	0.900	0.850	0.800
Variable 2		1.000	0.900	0.950	0.800	0.850
Variable 3			1.000	0.900	0.850	0.800
Variable A				1.000	0.800	0.850
Variable B					1.000	0.900
Variable C						1.000