

CRYSTAL BALL® AND DESIGN FOR SIX SIGMA

Lawrence I. Goldman
Crystal Campbell

Decisioneering, Inc.
1515 Arapahoe Street
Suite 1311
Denver, CO 80202, U.S.A.

ABSTRACT

In today's competitive market, businesses are adopting new practices like Design for Six Sigma (DFSS), a customer driven, structured methodology for faster-to-market, higher quality, and less costly new products and services. Monte Carlo simulation and stochastic optimization can help DFSS practitioners understand the variation inherent in a new technology, process, or product, and can be used to create and optimize potential designs. The benefits of understanding and controlling the sources of variability include reduced development costs, minimal defects, and sales driven through improved customer satisfaction. This tutorial uses Crystal Ball Professional Edition, a suite of easy-to-use Microsoft® Excel-based software, to demonstrate how stochastic simulation and optimization can be used in all five phases of DFSS to develop the design for a new compressor.

1 INTRODUCTION

An uncertain situation often involves risk, which is defined as the PROBABILITY of loss, damage, or any other undesirable event. Most people desire low risk, which translates to a higher probability of success, profit, or some form of gain. For example, if sales for next month are above a certain amount (a desirable event), then orders will reduce the inventory, and there will be a delay in shipping orders (an undesirable event). If a shipping delay means losing orders, then that possibility presents a risk.

Spreadsheets are a popular and accessible tool for modeling processes and plans. When using spreadsheets, analysts traditionally handle uncertainty by using average or best-guess values for uncertain inputs, primarily because Excel only lets them enter a single value in a cell. These deterministic models provide a single outcome upon which a business or technical decision is made. To capture uncertainty, analysts can perform simple "what-if analysis" or "scenario analysis" by manually changing model variables and analyzing their effect on the key outputs. This approach provides a

range of possible outcomes but does not impart an understanding of the likelihood or risk of any particular outcome.

1.1 Monte Carlo Simulation

One established solution to the limitations of spreadsheet risk analysis is Monte Carlo simulation. Since Excel alone does not have the ability to run and analyze simulations, modelers must rely on third-party programs such as Crystal Ball that add in and expand the features of Excel. Crystal Ball adds two techniques to Excel: the replacement of single values with probability distributions and the random simulation of a model. The result is a probability-enhanced spreadsheet with quantifiable outcomes, such as a 75% probability of making a profit on a new product or a 99.997% certainty of being within a specified tolerance requirement.

Monte Carlo simulation is a proven, efficient technique that only requires a random number table or a random number generator on a computer (Kelton and Law 1991). A random number is a mathematically selected value that is generated to conform to a probability distribution. For each simulation trial, Crystal Ball randomly selects a value from the distribution and enters that value into the Excel spreadsheet, which then recalculates. The multiple scenarios created through simulation can be analyzed to give more insight into the risks and mechanisms of the spreadsheet model. When used correctly, Monte Carlo simulation can provide valuable insights not available through deterministic models.

1.2 Design for Six Sigma Methodology

How can Monte Carlo simulation be applied to the creation of new products or services? For a growing number of companies, this question is being answered through Design for Six Sigma (DFSS). Six Sigma is a structured methodology for accelerated process improvement that aims for virtually error-free business processes (Pyzdek 2003). While based on many of the same principles as Six Sigma, DFSS differs from Six Sigma because (1) it tends to be more proactive

than reactive; (2) is applied to the new rather than existing products, services, and processes; and (3) is focused on marketing, R&D, and design rather than manufacturing or transactional processes (Creveling, Slutsky, and Antis 2003).

DFSS is divided into five distinct phases: Define, Measure, Analyze, Design, and Verify (DMADV). Alternative methodologies include PIDOV (Plan, Identify, Design, Optimize, Verify) and CDOV (Concept, Design, Optimize, Verify), to name just two. No matter what variation of DFSS a company may choose to adopt, the methodology will rely on a disciplined set of tools, such as spreadsheets, through the designing process (Brue and Launsby 2003). The inherent flexibility of spreadsheets means that they can be widely applied and can account for quality and cost changes across the different phases of the design process.

Monte Carlo simulation and stochastic optimization can play a critical role throughout the Six Sigma process (Goldman, Evans-Hilton, and Emmett 2003) as well as in the DMADV process of DFSS. For DFSS, critical benefits of simulation and optimization are the ability to prototype new products or processes without an appreciable investment of time or money, minimal defects, and sales driven through improved customer satisfaction.

2 CRYSTAL BALL PROFESSIONAL EDITION

Crystal Ball Professional Edition is a general desktop software suite that features spreadsheet-based analysis tools for Monte Carlo simulation (Crystal Ball), time-series forecasting (CB Predictor), and optimization (OptQuest). The suite also includes Crystal Ball and CB Predictor Developer Kits for building custom interfaces and processes using Visual Basic for Applications (VBA).

2.1 Crystal Ball

Crystal Ball is an easy-to-use Excel add-in designed to help all levels of spreadsheet modelers perform Monte Carlo simulation. Crystal Ball lets users define probability distributions on uncertain model variables, and then uses simulation to generate random values from within the defined probability ranges. Spreadsheet modelers can create and analyze thousands of alternate scenarios and quantify their level of risk for any given scenario. You can apply Crystal Ball to spreadsheet models, and the enhancements provided by Crystal Ball do not alter the formulae or functions of the original spreadsheet. Crystal Ball also includes the CB tools, seven wizard-driven add-ins that can assist in setting up and analyzing models.

2.2 OptQuest

OptQuest is a global optimization add-in that enhances Crystal Ball by automatically searching for and finding optimal solutions to simulation models. OptQuest uses a mixture of

techniques including scatter search and advanced tabu search to find the right combination of decision variables for the best possible results. As the program runs, adaptive and neural network technologies help it learn from past optimizations so that it achieves better results in less time. A setup wizard helps users define constraints, objectives, and requirements, and even to set up an efficient frontier option.

2.3 CB Predictor

CB Predictor is a wizard-driven Excel add-in that guides users through the process of time-series forecasting and multiple linear regression. CB Predictor analyzes historic data series and uses the level, trend, seasonality, and error of the data to project the series into the future. CB Predictor applies eight separate seasonal and nonseasonal time-series forecasting methodologies to data and can apply multiple linear regression in cases where the primary data series is dependent on other independent data series. The output forecasts can be defined as probability distributions and used in Monte Carlo simulation. CB Predictor will not be applied in this paper.

2.4 Developer Kits

The Crystal Ball and CB Predictor Developer Kits allows users with programming skills to completely automate and control Crystal Ball simulations and CB Predictor forecasts from within a VBA program or any other language outside of Excel supporting OLE 2 automation. The kits consist of libraries of macro commands and functions that unlock the programmability of Crystal Ball and CB Predictor. The Developer Kits will not be applied in this paper.

3 DESIGN FOR SIX SIGMA EXAMPLE

In this example, you are a compressor manufacturer in the process of developing a new type of compressor. Your project team has been charged with developing the design for the compressor using the DFSS tools and techniques. As you work through the DMADV process, you will see how simulation and optimization can provide a project justification, lend insight into the critical drivers of quality, and help create a cost effective design that meets customer requirements.

3.1 The Define Phase

The Define phase is used to determine the project goals and the requirements of customers, both internal and external. Define is primarily focused on customer and employee interviews to understand performance requirements. The Define phase is one of the most important phases, because it sets the framework and goals for the project. Simulation can be used for projecting the potential financial impact of a new product. Your team has decided to simulate the financial risks in developing the compressor prior to any design development.

3.1.1 Applying Crystal Ball

Before you use Crystal Ball Professional Edition, you will need an Excel spreadsheet model that represents the process, system, or problem that you are attempting to solve. This paper will discuss only how the software is applied to the DFSS process, not the specifics of model building. Many authors have recently provided excellent texts on how to create valid and verifiable spreadsheet models (Evans and Olsen 2001; Powell and Baker 2003; Mun 2003).

You start your feasibility study by developing a simple spreadsheet model that takes into account the greatest risks in developing this new product (Figure 1). You consider three major areas of concern: technical feasibility, manufacturing capability, and market uptake. Once the model is complete, you can add Crystal Ball elements.




Figure 1: Compressor Feasibility Spreadsheet

When you launch Crystal Ball and Excel through the Windows Start menu, Crystal Ball adds a new tool bar and three new menus to Excel. The tool bar has been designed from left to right to follow the modeling process from setup to simulation to the analysis and presentation of results.

The three Crystal Ball menus contain all of the functions on the tool bar plus additional features. The Cell menu contains all of the model setup functions, and the Run menu lists all of the simulation and analysis functions. The final menu, CBTools, opens the CB tools and other products in the Professional Edition suite.

3.1.2 Defining Crystal Ball Assumptions

In Crystal Ball, probability distributions are referred to as assumptions and are used to define the uncertainty or variability in any model input. The best way to identify which variables to convert from single values to probability distributions is to decide which variables are known and which are variable and uncertain.

To define an assumption, select a spreadsheet cell that contains numeric data that you know is variable or uncertain (e.g., material costs, staffing, market demand) and click on the Define Assumption button on the tool bar. In the Distribution Gallery (Figure 2), Crystal Ball offers sixteen pre-defined distributions and one customizable distri-

bution. The four most commonly used distributions are the normal, triangular, uniform, and lognormal.




Figure 2: Distribution Gallery

For the Technical Feasibility, you know from your design team that your probability of success is most likely 89% but could be as little as 85% or as high as 95%. You use these parameters to define a triangular distribution that represents the uncertainty of the technology in the new design (Figure 3). You also define assumptions around other uncertainties such as product development costs and price per unit.




Figure 3: Assumption for Technical Feasibility

3.1.3 Defining a Crystal Ball Forecast

After defining all the assumptions in your model, you need to define the outputs, or forecasts, that you want to examine with Crystal Ball. A forecast is any formula cell that you assign Crystal Ball to track during a simulation. As with assumptions, there is no limit to the number of forecasts you can define, although for more forecasts, more memory is required to store the simulated values.

In this model your primary concern is net present value (NPV). To define a forecast for NPV, you select the cell, click on the Define Forecast button to open the Define Forecast dialog (Figure 4), enter a unique name, and click on OK.




Figure 4: Forecast Chart for Product NPV

3.1.4 Setting the Simulation Run Preferences

Crystal Ball gives you full access to the simulation controls. These settings include the number of trials to run, what sampling method to use (Monte Carlo or Latin Hypercube), where and when during the simulation cycle to run macros, ways to increase the speed of simulations, and Precision Control, a method of confidence testing that lets you achieve a desired level of precision in your forecasts based on simulation statistics.

3.1.5 Testing and Running the Simulation

Before running the simulation, you can check that the model is functioning properly by clicking on the Single Step button on the tool bar. Each single step is a random trial in Crystal Ball. One benefit of single stepping is that, when the model recalculates for alternate scenarios, you may discover calculation errors that were not apparent from your original set of values.

During a simulation, Crystal Ball runs as many trials as you request and saves the forecast values for later analysis. To run the simulation, click on the Start Simulation button on the tool bar. As the simulation runs, you can watch the spreadsheet values change as Crystal Ball inserts randomly generated values (from the probability distributions) into the assumption cells and the spreadsheet recalculates.

3.1.6 Analyzing the Forecast Chart

In this project, when Crystal Ball has completed 5,000 trials, you can analyze these results using the NPV forecast (or frequency) chart, an interactive histogram that contains all of the statistics for the 5000 trials. Using this chart, you predict a 71% overall probability of success (Figure 4).

While the success rate for the product is not exceptional, the potential rewards of this project are great, compared to the relatively minor risk. The percentile values of the NPV forecast shown in Figure 5 illustrate that the likelihood of losing more than \$1.8MM is only 10% while the

likelihood of creating an NPV of at least \$48.5MM is more than 70%. Based on this information, you decide to proceed with the project.




Figure 5: Percentile Values for the NPV Forecast

In the Design phase, your team also develops spreadsheet models to simulate and forecast the projected timeline and the required resources for the project.

3.2 The Measure Phase

During the Measure phase, the design team gathers data that assess customer needs and specifications for the new compressor. Some of this data may be new, and much of it has never been scrutinized before. Your Six Sigma team must establish valid and reliable metrics to help monitor their progress towards the project goals.

Your primary goal at this point in the process is to establish the Critical-to-Quality (CTQ) characteristics. Based on customer feedback, you know the top CTQ is flow rate through the compressor at a specified temperature, pressure, and speed. Your team builds a model to calculate the flow rate from the input parameters.

As with the previous model, you apply assumptions to each of the uncertain input parameters to describe their variability. The data for these compressor part parameters are available from similar compressor systems your company has already designed. You then run a simulation and view the Sensitivity Chart (Figure 6) for the Flow Rate, which you have defined as a forecast.

As the simulation runs, Crystal Ball's sensitivity analysis uses rank correlation to calculate the relationships between the assumptions and the forecast. The sensitivity chart displays these relationships as correlation coefficients or as percentage values. The top listed assumptions have the greatest effect on the forecast, and the direction of the bar indicates a direct or inverse relationship.

The variation in piston stroke length was by far the most important input variable. In order to produce compressors with the least variability in the flow rate (as desired by your customers), you will have to minimize the variability in the piston stroke length of the compressor.




Figure 6: Sensitivity Analysis of Flow Rate

3.3 The Analyze Phase

In the Analyze phase, the team takes the results from the Measure phase and works to determine your design options for the piston assembly (Figure 7).




Figure 7: Tolerance Study Spreadsheet Model

Your team must decide what are the nominal values of arm length and crank length, which together with the other pre-defined dimensions determine the piston stroke length. Both simulation and optimization play an important role in the design analysis, and you will need to run an optimization on these two elements to ensure a proper design.

3.3.1 Decision Variables

To run an optimization, your team must first identify the decision variables in your spreadsheet. Decision variables are model inputs over which you have control, in this case, the arm length and crank length. Your target stroke length is 30.80 inches.

To define a decision variable, select a spreadsheet cell and click on the Define Decision button on the tool bar. In the Define Decision Dialog, you select the upper and lower bounds for the decision and whether the range should be

continuous or discrete. If discrete, you must define the step size of the discrete values. For the Crank Length, for example, you define the minimum length as 6.125 inches and the maximum length as 8.125 inches, with a discrete step of 0.1 (Figure 8).




Figure 8: Define Decision Variable for Crank Length

Since the crank angle producing the maximum stroke length depends on the arm length and crank length, your team wrote a simple macro that invokes Solver prior to each simulation to find the crank angle based on the two decision variables. The macro is called before each simulation through a setting defined in the Crystal Ball run preferences. Now, you are ready to run an optimization using OptQuest.

3.3.2 Setting Up OptQuest

With your Excel model and Crystal Ball running, you can open OptQuest by selecting the program from the CBTools menu. Once OptQuest is open, the step-by-step wizard guides the process of defining constraints, specifying the objective of the outcome, and setting the requirements. In the first step, OptQuest shows you the model's two decision variables, which you will use in your optimization.

The final step in setting up OptQuest is to define the objective of the optimization. All Crystal Ball forecasts are displayed in the Forecast Selection window, and you select which forecast to minimize or maximize. For this project, you want the simulated stroke length to be as close as possible to the target stroke length, so your goal is to minimize the mean of the error between the two lengths, which you have defined as a forecast. Once you have defined this objective, you select a run time and click on the Run button.

3.3.3 Running OptQuest

When running, OptQuest selects a value for each decision variable, enters those values into your spreadsheet, runs the Solver macro, runs a Monte Carlo simulation on the spreadsheet, records the results, and repeats the process. You could manually perform this sort of analysis, but as

you increase the number of decision variables, the number of possible variable combinations becomes unwieldy.

On a more advanced level, OptQuest does a much better job at finding optimal solutions than is possible with manual calculations. OptQuest surpasses the limitations of genetic algorithm optimizers because it uses multiple, complementary search methodologies, including advanced tabu search and scatter search, to help find the best global solutions. While running solutions, OptQuest also checks for compliance with any constraints or requirements. Additionally, OptQuest applies adaptive and neural network technologies to help it learn from past optimizations so it can converge on better results in less time.

3.3.4 Analyzing the Optimization Results

After running OptQuest, you can use the Solution Analysis feature (Figure 9) to analyze the designs that resulted in a nominal stroke length very near the target of 30.80 inches.




Figure 9: Analysis of Near-Optimal Solutions in OptQuest

OptQuest quickly identified five solutions that most closely approximate the desired stroke length, and these are the solutions that can be compared in the Design phase.

3.4 The Design Phase

In this phase of development, your team must design the product so that it meets customer requirements. This means that the piston assembly variability must be minimized. With five possible piston assembly solutions from the Analyze phase, your goal is to choose the optimum quality design that is also cost efficient.

The variation (uncertainty) in this design model is represented by the error in each of the five piston dimensions used in the calculation of piston stroke length. You define these errors as normal assumptions. You then modify your compressor model to include cost functions and levels of machined quality for each of the five dimensions. Because

you have control over the desired level of quality, you can define these five quality factors as decision variables.

You also use these quality costs and the simulated DPMO to calculate the total cost to produce one million in specification parts. As you increase quality (and decrease variability), your machining costs will increase, so your optimization needs to determine the lowest possible cost you can expect for the required level of piston quality.

You re-run OptQuest, this time with the objective of minimizing the total cost to produce 1,000,000 in specification parts. OptQuest will minimize cost by adjusting the nominal crank and arm lengths as well as the quality level of each of the five parameters that determine piston stroke length. The results are displayed below in Figure 10. The Performance Graph, shown in the bottom of Figure 10, is useful for gauging OptQuest's progress. As the graph levels off and remains unchanged, you can decide whether or not you have converged on an optimal solution. OptQuest selected the third design solution with a total cost value of \$17 million, which meets cost expectations.




Figure 10: Optimization to Minimize Total Cost

3.5 The Verify Phase

But does this final design meet customer requirements? The Verify phase requires you to check that the customer needs are satisfied. OptQuest allows you to copy the optimum design parameters back into your Excel model. After running 10,000 simulation trials on these new parameters (Figure 11), you find that roughly 80% of your production would be within the specification limits.

While the economic analysis indicates that this is the most cost-effective solution, you were charged with designing a piston assembly that provides 95% certainty of a stroke length between 30.70 and 30.90. This problem can still be solved in OptQuest.

By setting a requirement on a forecast in OptQuest, you ensure that OptQuest will reject all solutions that do




Figure 11: Certainty of Achieving Specification Limits on Stroke Length

not meet your set criteria. With this in mind, you rerun the optimization, with the same objective but with an additional requirement of 95% certainty that the stroke length will be within the specification limits (Figure 12).




Figure 12: Minimization of Costs Subject to 95% in Specification Requirement

Your new optimization, which selected design solution 1, is successful. While the costs have risen to \$24 million, you have ensured that the piston assembly will be within tolerance, that your company will create a compressor with few defects, and that the customer requirements will be satisfied.

4 CONCLUSION

Monte Carlo simulation and optimization have a crucial role to play in all phases of a Design for Six Sigma project. Without the ability to run simulations, spreadsheets alone can provide only a limited understanding of the variability surrounding a process or project design. Crystal Ball Professional Edition overcomes the limitations of spreadsheets by enhancing Excel with user-friendly tools for Monte Carlo simulation, global optimization, and time-series

forecasting. By moving to a probabilistic methodology, DFSS practitioners can better quantify the effects of variability and implement robust product designs with greater insight and confidence.

APPENDIX: CRYSTAL BALL RESOURCES

A variety of Six Sigma models and white papers are available for free on the Decisioneering Web site <www.crystalball.com/sixsigma>. Other free resources on the site include tutorials, case studies, and Internet links. Free downloadable trial versions of the software and tutorials are available to all site visitors.

REFERENCES

- Brue, G., and R.G. Launsby. 2003. *Design for Six Sigma*. New York: McGraw Hill.
- Creveling, C. M., J. L. Slutsky, and D. Antis, Jr. 2003. *Design for Six Sigma*. Upper Saddle River, New Jersey: Prentice Hall PTR
- Evans, J. R., and D. L. Olson. 2001. *Introduction to Simulation and Risk Analysis, 2nd Edition*. Upper Saddle River, New Jersey: Prentice Hall Business Publishing.
- Goldman, L.I., E. Evans-Hilton, and H. Emmett. Crystal Ball for Six Sigma Tutorial. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, 293-300. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Kelton, W. D., and A. Law. 1991. *Simulation Modeling & Analysis*. New York: McGraw Hill, Inc.
- Mun, J. 2003. *Applied Risk Analysis: Moving Beyond Uncertainty*. Hoboken, New Jersey: John Wiley & Sons, Inc.
- Powell, S. G., and K.R. Baker. 2003. *The Art of Modeling with Spreadsheets: Management Science, Spreadsheet Engineering, and Modeling Craft*. Hoboken, New Jersey: John Wiley & Sons, Inc.
- Pyzdek, T. 2003. *The Six Sigma Handbook*. New York: McGraw Hill, Inc.

AUTHOR BIOGRAPHIES

LAWRENCE GOLDMAN is a product marketing manager at Decisioneering in Denver, CO. He joined the company in June 1997 as a program manager with the product development group and has held positions as webmaster, web producer, trainer, and newsletter editor. His current focus is educational and training materials and seminars for Crystal Ball. Prior to working for Decisioneering, Lawrence was employed as a trainer and technical support assistant for Techbase, a company specializing in mining and engineering database software. He has an M.S. in geology from the University of Cincinnati and a B.A. in geology from Cornell University. His email address and

Goldman and Campbell

Web address are <lgoldman@crystalball.com>
and <www.crystalball.com>.

CRYSTAL CAMPBELL is a trainer and product consultant. After completing a BS in Chemical Engineering at Kansas State University, she went to work for The Dow Chemical Company in a polyethylene manufacturing facility. While at Dow, Crystal became a certified Six Sigma practitioner by leading a Six Sigma project team and participating in multiple black belt Six Sigma projects. There she learned first hand the need to make effective business decisions quickly and began training others to use the data available through technology to drive decision making. She is excited to help individuals learn to make more effective decisions by incorporating risk analysis and optimization into their decision making processes. Her email address and Web address are <ccampbell@crystalball.com> and <www.crystalball.com>.