

DOE Wizard – Mixture Designs

Summary

When the experimental factors to be studied are ingredients or components of a mixture, the response function typically depends upon the relative proportions of each component, not the absolute amount. Since the proportions must sum to a fixed amount, usually 100%, the factors cannot be varied independently of one another. Consequently, the designs normally used for screening and optimization cannot be applied directly, nor is the rectangular coordinate space the most desirable representation of the experimental region.

For example, consider a blending problem in which q components may be mixed together in different proportions. Let

X_j = amount of component j in the mixture,

usually represented as a proportion or percentage, although the factors may be expressed in any convenient units. The sum of the q components is constrained to equal a fixed value T , i.e.,

$$X_1 + X_2 + \dots + X_q = T,$$

where T often equals 100%. In addition, each component is subject to lower and upper bounds:

$$L_j \leq X_j \leq U_j,$$

which may be as simple as $L_j = 0\%$ and $U_j = 100\%$ or more restrictive. The only legitimate experimental runs are those that satisfy all of the above constraints.

STATGRAPHICS contains several different types of designs intended to study the effect of up to 12 components on one or more responses. These include:

1. simplex-lattice designs
2. simplex-centroid designs
3. extreme vertices designs

In addition, one or more process variables may be added to a mixture experiment when using the DOE Wizard.

Example

A typical example of a mixture problem, discussed by Myers and Montgomery (2002), involves the formulation of a rocket propellant. The propellant is a mixture of three components: a fuel, an oxidizer, and a binder. The researcher wished to find a combination of these three components which achieved a satisfactory burn rate. Since an inert component made up 10% of the propellant, the factors were constrained by:

$$fuel + oxidizer + binder = 90\%. \quad (1)$$

In addition, there were lower bounds for each of the three components:

$$\begin{aligned} 30\% &\leq fuel \\ 20\% &\leq oxidizer \\ 20\% &\leq binder \end{aligned} \quad (2)$$

Given the constraints, there is a remaining 20% of the mixture that can be any combination of *fuel*, *oxidizer*, and/or *binder*.

Sample StatFolio: doewiz mixture.sgp

Design Creation

To begin the design creation process, start with an empty StatFolio. Select *DOE – Experimental Design Wizard* to load the DOE Wizard's main window. Then push each button in sequence to create the design.

Step #1 – Define Responses

The first step of the design creation process displays a dialog box used to specify the response variables. For the rocket propellant example, there is a single response variable:

Design of Experiments Wizard - Define Responses

Design file: <untitled>

Comment: Rocket propellant example

Number of responses: 1

Response	Name	Units	Analyze	Goal	Target	Impact (1-5)	Sensitivity	Minimum	Maximum
1	burn rate	cm per second	Mean	Hit target	85	3.0	Medium		
2	Var_2		Mean	Maximize	0.5	3.0	Medium		
3	Var_3		Mean	Maximize	0.5	3.0	Medium		
4	Var_4		Mean	Maximize	0.5	3.0	Medium		
5	Var_5		Mean	Maximize	0.5	3.0	Medium		
6	Var_6		Mean	Maximize	0.5	3.0	Medium		
7	Var_7		Mean	Maximize	0.5	3.0	Medium		
8	Var_8		Mean	Maximize	0.5	3.0	Medium		
9	Var_9		Mean	Maximize	0.5	3.0	Medium		
10	Var_10		Mean	Maximize	0.5	3.0	Medium		
11	Var_11		Mean	Maximize	0.5	3.0	Medium		
12	Var_12		Mean	Maximize	0.5	3.0	Medium		
13	Var_13		Mean	Maximize	0.5	3.0	Medium		
14	Var_14		Mean	Maximize	0.5	3.0	Medium		
15	Var_15		Mean	Maximize	0.5	3.0	Medium		
16	Var_16		Mean	Maximize	0.5	3.0	Medium		

OK Cancel Help

- **Name:** The name for the variable is *burn rate*.
- **Units:** Burn rate is measured in centimeters per second.
- **Analyze:** The parameter of interest is the *mean* burn rate.
- **Goal:** The goal of the experiment is to achieve a target value for the mean burn rate.
- **Target:** The desired mean burn rate is 85 centimeters per second.

Step #2 – Define Experimental Factors

The second step displays a dialog box on which to specify the components that will be varied. In the rocket propellant example, there are 3 components:

Design of Experiments Wizard - Define Factors

Design file: <untitled>

Comment: Rocket propellant example

Number of controllable process factors: 0 Number of controllable mixture components: 3 Number of noise factors: 0

Factor	Name	Units	Type	Role	Low	High	Levels
A	fuel	percent	Mixture	Controllable	30	50	1,2,3,4
B	oxidizer	percent	Mixture	Controllable	20	40	1,2,3,4
C	binder	percent	Mixture	Controllable	20	40	1,2,3,4
D	Factor_D		Continuous		-1.0	1.0	1,2,3,4
E	Factor_E		Continuous		-1.0	1.0	1,2,3,4
F	Factor_F		Continuous		-1.0	1.0	1,2,3,4
G	Factor_G		Continuous		-1.0	1.0	1,2,3,4
H	Factor_H		Continuous		-1.0	1.0	1,2,3,4
I	Factor_I		Continuous		-1.0	1.0	1,2,3,4
J	Factor_J		Continuous		-1.0	1.0	1,2,3,4
K	Factor_K		Continuous		-1.0	1.0	1,2,3,4
L	Factor_L		Continuous		-1.0	1.0	1,2,3,4
M	Factor_M		Continuous		-1.0	1.0	1,2,3,4

Total for controllable mixture components: 90

Factors A-M Factors N-Z

OK Back Cancel Help

- **Name** – Each component must be assigned a unique name.
- **Units** – In this case, the value of each component will be its percent of the mixture.
- **Type** – Each component is a *Mixture* component.
- **Role** – Each component is *Controllable*.
- **Low** - the lower bound L_j for the component.
- **High** - the upper bound U_j for the component.
- **Total for controllable mixture components** - the total or sum T of the q components. The component levels in all runs will sum to this total.

When the OK button is pressed, one of several events may occur:

1. If the constraints pose no problems, the dialog box disappears and control returns to the main DOE Wizard window.
2. If any upper bound is less than any lower bound, an error message is displayed.
3. If the sum of the low fields is greater than T or the sum of the high fields is less than T , an error message is displayed.

- If the constraints are inconsistent, i.e., a high value can never be achieved, an error message is displayed.

Note: the specification of components may be done in any convenient unit, such as percentages, proportions, weights, etc. For purposes of analysis, some of the results (and the mathematical discussion below) will be expressed in terms of *pseudo-components*. A pseudo-component, which will be denoted by a lowercase x , rescales a component so that it equals 0 at its minimum permissible value and 1 at its maximum permissible value, i.e.,

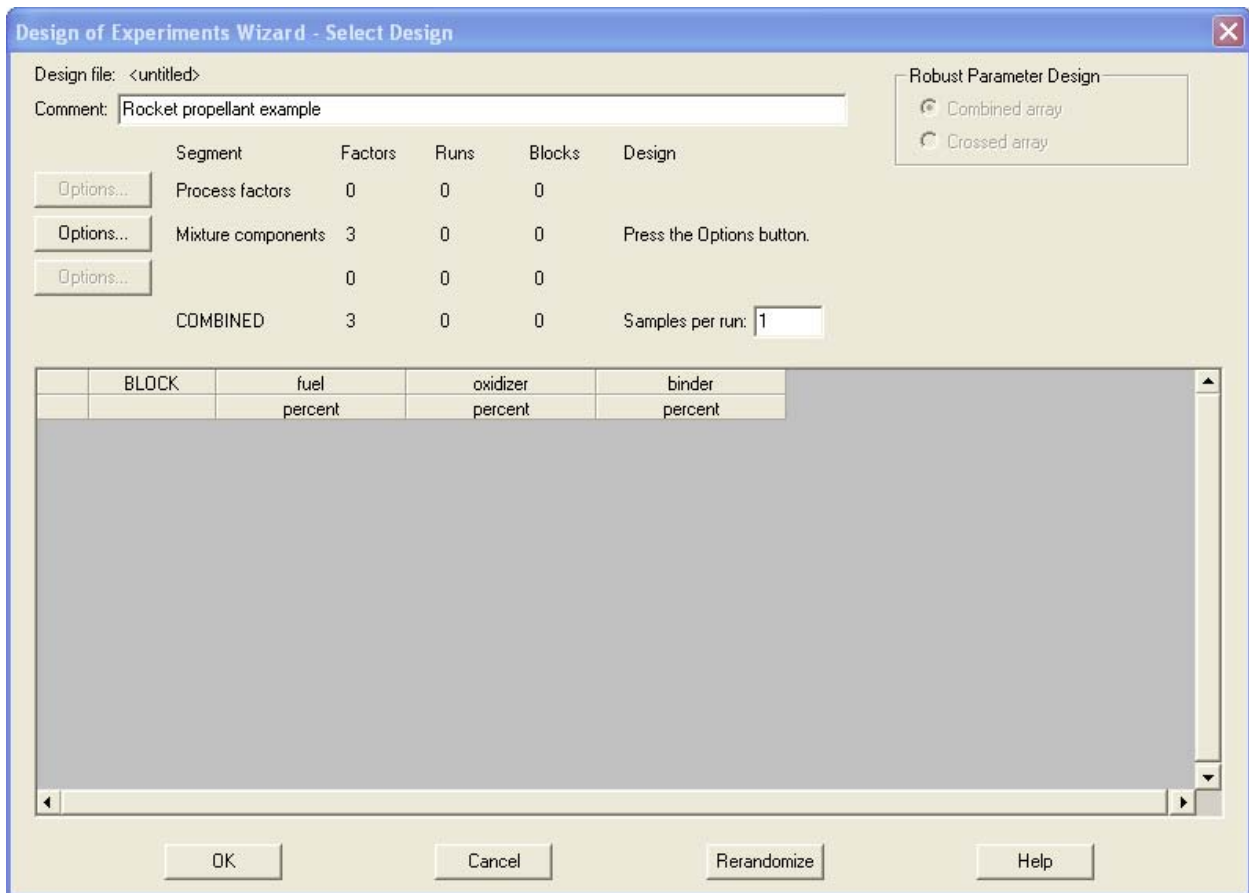
$$x_j = \frac{X_j - L_j}{T - \sum_{j=1}^q L_j} \tag{3}$$

The pseudo-components must sum to 1 for each experimental run:

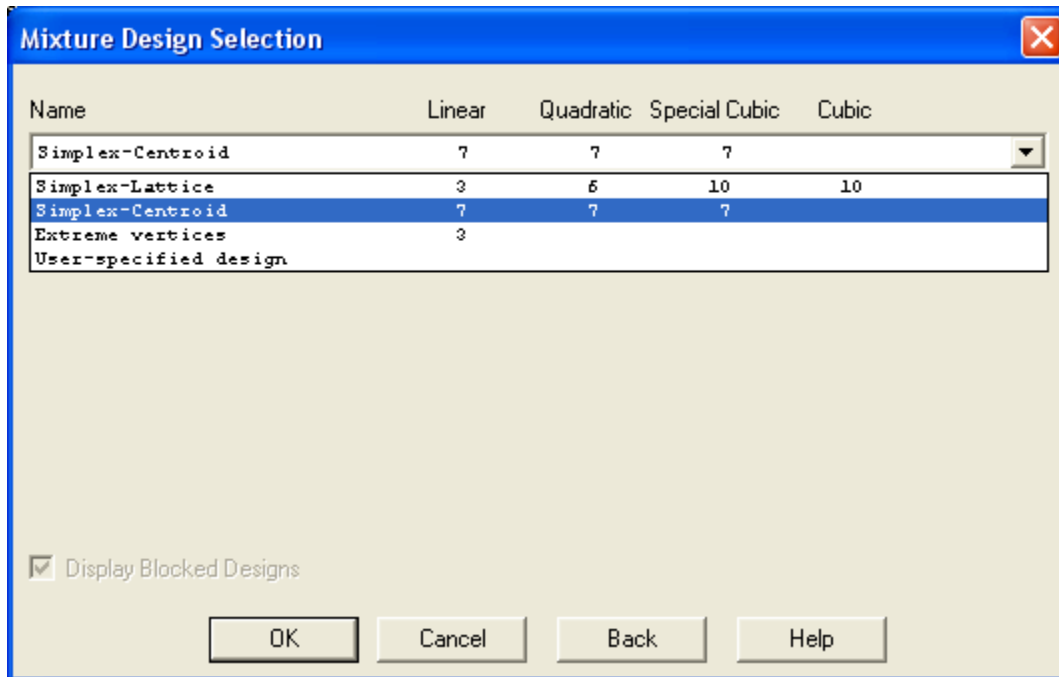
$$x_1 + x_2 + \dots + x_q = 1. \tag{4}$$

Step #3 – Select Design

The third step begins by displaying the dialog box shown below:



Since all of the factors are controllable mixture components, only one *Options* button is enabled. Pressing that button displays a list of all mixture designs in the STATGRAPHICS catalog for the selected number of components:



- **Name** - the design name. For mixture designs, the following types may appear in the list, depending upon the number of components and the nature of any constraints imposed:
 1. *Simplex-Lattice* - a uniformly spaced set of points on a simplex, where a simplex is a regularly sided figure with q vertices in $q-1$ dimensions (e.g., a triangle in a 2-dimensional space). Depending on the constraints, these designs may not be available.
 2. *Simplex-Centroid* – a design consisting of $2^q - 1$ runs at all primary blends, binary blends, tertiary blends, etc. up to the design centroid. Depending on the constraints, these designs may not be available.
 3. *Extreme vertices* – a design consisting of a run at each vertex of the feasible region. This design is always available.
 4. *User-Specified* – columns are created in the datasheet for each component and response variable, but the user must define the runs to be performed.
- **Linear** - the number of runs in the base design if the most complicated model to be fit is a linear model. A linear model has the form

$$Y = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 \quad (5)$$

- **Quadratic** - the number of runs in the base design if the most complicated model to be fit is a quadratic model. A quadratic model has the form

$$Y = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_{12} x_1 x_2 + \beta_{13} x_1 x_3 + \beta_{23} x_2 x_3 \quad (6)$$

- **Special cubic** - the number of runs in the base design if the most complicated model to be fit is a special cubic model. A special cubic model has the form

$$Y = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_{12} x_1 x_2 + \beta_{13} x_1 x_3 + \beta_{23} x_2 x_3 + \beta_{123} x_1 x_2 x_3 \quad (7)$$

- **Cubic** - the number of runs in the base design if the most complicated model to be fit is a full cubic model. A full cubic model has the form

$$Y = \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_{12} x_1 x_2 + \beta_{13} x_1 x_3 + \beta_{23} x_2 x_3 + \beta_{123} x_1 x_2 x_3 + \delta_{12} x_1 x_2 (x_1 - x_2) + \delta_{13} x_1 x_3 (x_1 - x_3) + \delta_{23} x_2 x_3 (x_2 - x_3) \quad (8)$$

Notice that terms such as a constant and factors raised to the second power do not appear. Because of the constraint on the sum of the components, the full polynomial model would be over-parameterized. The special cubic model is popular and can represent many types of response surfaces.

Simplex-Lattice Designs

A simplex-lattice design consists of a uniformly spaced set of points on a simplex. As will be discussed more fully below, the simplex is the natural representation of the experimental region for a mixture experiment. The simplex-lattice design uses $m+1$ equally spaced values between 0 and 1, defined by

$$x_j = 0, \frac{1}{m}, \frac{2}{m}, \dots, 1 \quad (9)$$

where m is the order of the model to be fit. To fit a linear model, only pure blends are used, such as

$$\begin{aligned} (x_1=1, x_2=0, x_3=0) \\ (x_1=0, x_2=1, x_3=0) \\ (x_1=0, x_2=0, x_3=1) \end{aligned}$$

To fit a quadratic model, binary blends such as

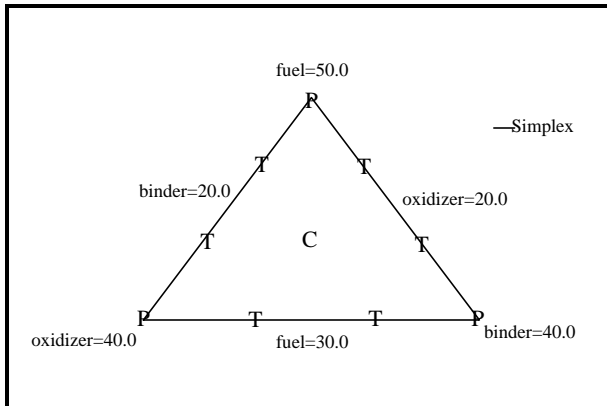
$$\begin{aligned} (x_1=.5, x_2=.5, x_3=.0) \\ (x_1=.5, x_2=.0, x_3=.5) \\ (x_1=.0, x_2=.5, x_3=.5) \end{aligned}$$

are added. A cubic model adds the following tertiary blends instead of the binary blends:

$$\begin{aligned} (x_1=.667, x_2=.333, x_3=.000) \\ (x_1=.667, x_2=.000, x_3=.333) \\ (x_1=.000, x_2=.667, x_3=.333) \\ (x_1=.333, x_2=.667, x_3=.000) \\ (x_1=.333, x_2=.000, x_3=.667) \\ (x_1=.000, x_2=.333, x_3=.667) \end{aligned}$$

$$(x_1=.333, x_2=.333, x_3=.333)$$

The design for a cubic or special cubic model is shown below:



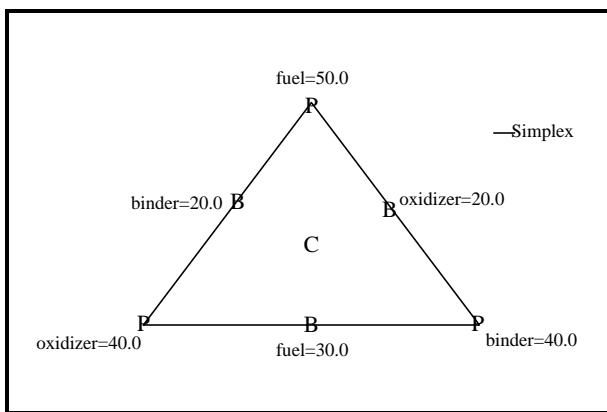
The primary blends are marked by a P, the tertiary blends by a T, and the centroid by a C.

Simplex-Centroid Designs

The simplex-centroid design consists of $2^q - 1$ runs at all primary blends, binary blends, tertiary blends, etc. up to the design centroid given by

$$\left(x_1 = \frac{1}{q}, x_2 = \frac{1}{q}, \dots, x_q = \frac{1}{q} \right) \tag{10}$$

This design tends to place more points at interior locations of the simplex, rather than on the design boundaries as does the simplex-lattice design. The design for a special cubic model is shown below:



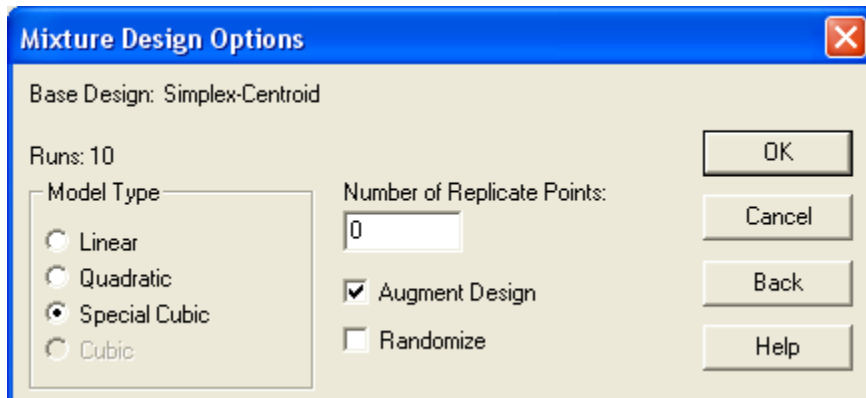
Extreme Vertices Designs

If restrictive lower and upper bounds exist on the components, the design space may be constrained such that it is not a regularly shaped simplex. In such cases, the simplex-lattice and simplex-centroid designs are not available. The only design that will be available in such cases consists of one that places runs at each vertex of the design region, regardless of its shape. As discussed more fully in Myers and Montgomery (1995), there may be quite a few such vertices. STATGRAPHICS automatically finds all vertices when this type of design is selected. The

Optimize Design procedure on the DOE menu can then be used if necessary to reduce the number of design points.

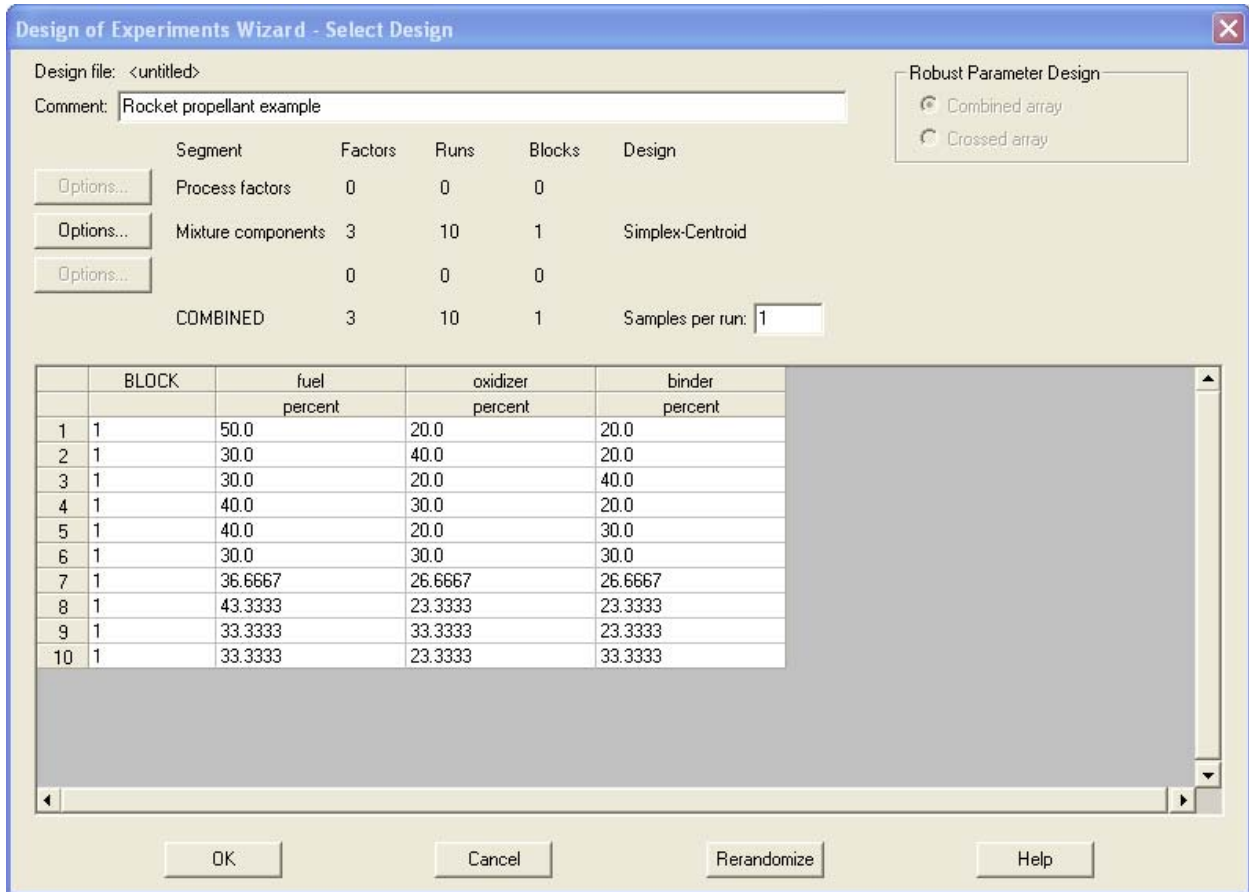
For the current example, a simplex-centroid design will be selected. This design contains 7 runs and is capable of estimating all of the models except the full cubic.

After the basic design is selected, a final dialog box allows the analyst to add additional runs to the design and to specify the order in which the runs will be performed:

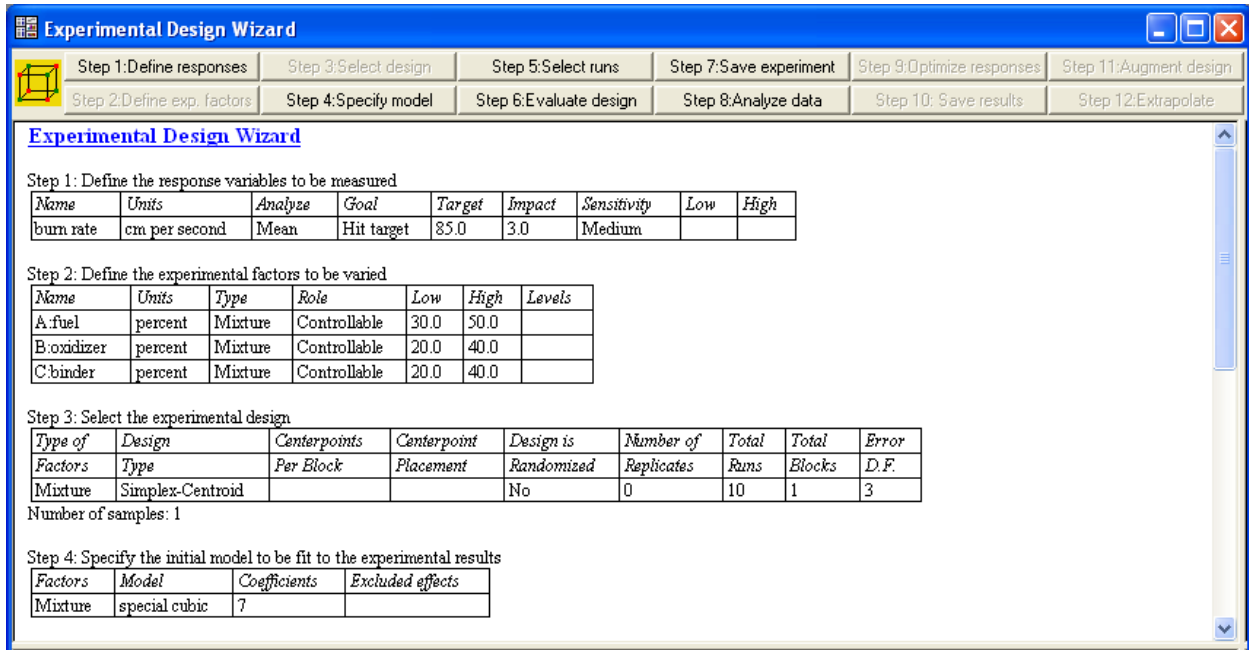


- **Model Type** - the most complicated model form that is to be fit. Enough runs will be added to fit that type of model. When analyzing the results, the model can of course be simplified.
- **Number of Replicate Points** - if not zero, the specified number of points will be selected at random from the original design points and replicated. This can be used to provide additional degrees of freedom to estimate the experimental error. Alternatively, the analyst can manually choose runs to replicate rather than letting the program do so simply by adding runs to the bottom of the datasheet after the design is created.
- **Augment Design** - if selected, additional runs will be placed at special checkblends, i.e., combinations of the components that can be used to check the fit of the selected model. For the simplex-lattice and simplex-centroid designs, these runs are located along axial lines running from the pure blends through the centroid. For the extreme vertices designs, the program adds points on axial lines running from the centroid to the vertices, at the centers of the edges connecting the vertices, at the constraint plane centroids, and at the overall centroid. Note: for extreme vertices designs with simple constraints, it may be necessary to select this option in order to have enough runs to estimate all but the simplest type of model.
- **Randomize** - if selected, the runs will be arranged in random order.

In the current example, the experimenters choose a 10-run design consisting of a 7-run simplex-centroid with 3 additional checkblends. The tentatively selected design is displayed in the *Select Design* dialog box:

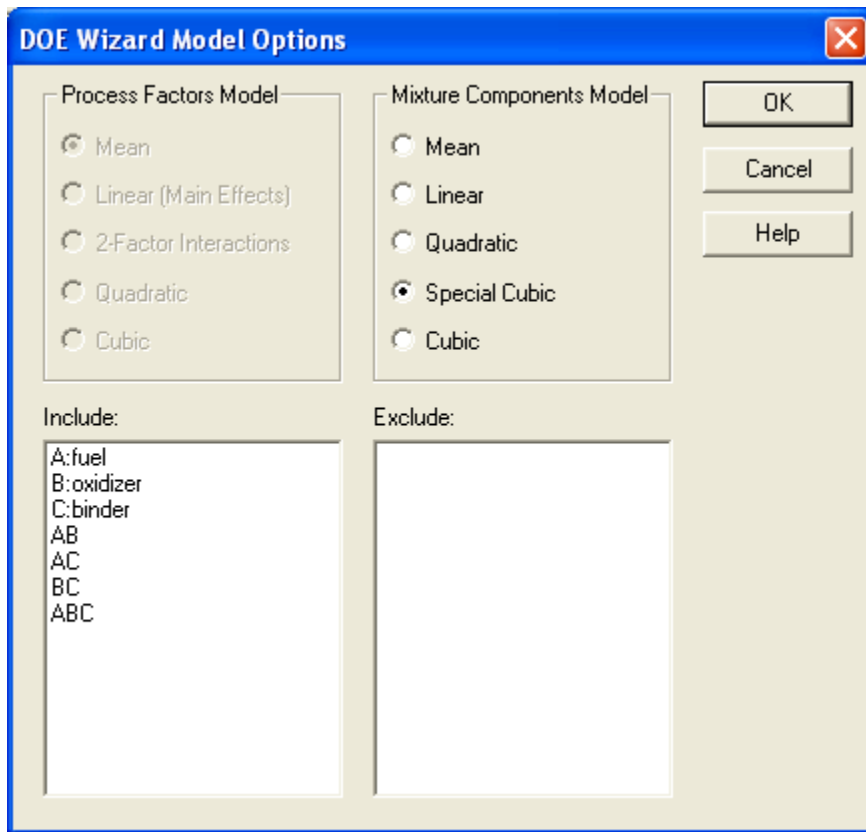


If the design is acceptable, press OK to save it to the STATGRAPHICS DataBook and return to the DOE Wizard’s main window, which should now contain a summary of the design:



Step #4: Specify Model

Before evaluating the properties of the design, a tentative model must be specified. Pressing the fourth button on the DOE Wizard's toolbar displays a dialog box to make that choice:



In this case, we will use the *special cubic* model. Selected terms could be excluded by double-clicking on them with the left mouse button.

Step #5: Select Runs

Since we intend to run all of the runs in the base design, this step can be omitted.

Design Properties

Step #6: Evaluate Design

Several of the selections presented when pressing button #6 are helpful in evaluating the selected design:

Design Worksheet

The design worksheet shows the 10 runs that have been created, in the order they are to be run:

Worksheet for rocket.sgx - Rocket propellant example				
<i>run</i>	<i>fuel</i>	<i>oxidizer</i>	<i>binder</i>	<i>burn rate</i>
	<i>percent</i>	<i>percent</i>	<i>percent</i>	<i>cm per second</i>
1	50.0	20.0	20.0	
2	30.0	40.0	20.0	
3	30.0	20.0	40.0	
4	40.0	30.0	20.0	
5	40.0	20.0	30.0	
6	30.0	30.0	30.0	
7	36.6667	26.6667	26.6667	
8	43.3333	23.3333	23.3333	
9	33.3333	33.3333	23.3333	
10	33.3333	23.3333	33.3333	

ANOVA Table

The ANOVA table shows the breakdown of the degrees of freedom in the design:


ANOVA Table	
Source	D.F.
Model	7
Total Error	3
Lack-of-fit	3
Pure error	0
Total (corr.)	10

7 of the 10 total degrees of freedom are used to estimate terms in the special cubic model. 3 degrees of freedom are left to estimate the experimental error. The 0 degrees of freedom for *pure error* is due to the lack of any replication in the design.

Design Points

This plot displays the experimental region and the points at which data will be collect

Rocket propellant example



All points within the triangular region represent feasible combinations of the 3 components. Each vertex of the triangle represents a point at which one of the components is set at its maximum level, while the other components are set at their minimum levels. In the current example, there are 10 unique combinations of the components.

Adding Additional Runs

According to Myers and Montgomery, the rocket propellant researchers decided to add 5 additional runs to the base design, each of which would be a replicate of a run already performed. Additional runs may be easily added by hand to the design by going to the DataBook and typing or pasting in the desired information, as shown below:

	BLOCK	fuel percent	oxidizer percent	binder percent	burn rate cm per second	Col
1	1	50.0	20.0	20.0		
2	1	30.0	40.0	20.0		
3	1	30.0	20.0	40.0		
4	1	40.0	30.0	20.0		
5	1	40.0	20.0	30.0		
6	1	30.0	30.0	30.0		
7	1	36.6667	26.6667	26.6667		
8	1	43.3333	23.3333	23.3333		
9	1	33.3333	33.3333	23.3333		
10	1	33.3333	23.3333	33.3333		
11	1	50.0	20.0	20.0		
12	1	30.0	40.0	20.0		
13	1	30.0	20.0	40.0		
14	1	36.6667	26.6667	26.6667		
15	1	36.6667	26.6667	26.6667		
16						
17						

In this case, rows 11 through 13 are replicate runs at the 3 vertices, while runs 14 and 15 are additional runs at the centerpoint. Be sure to add the block numbers as well.

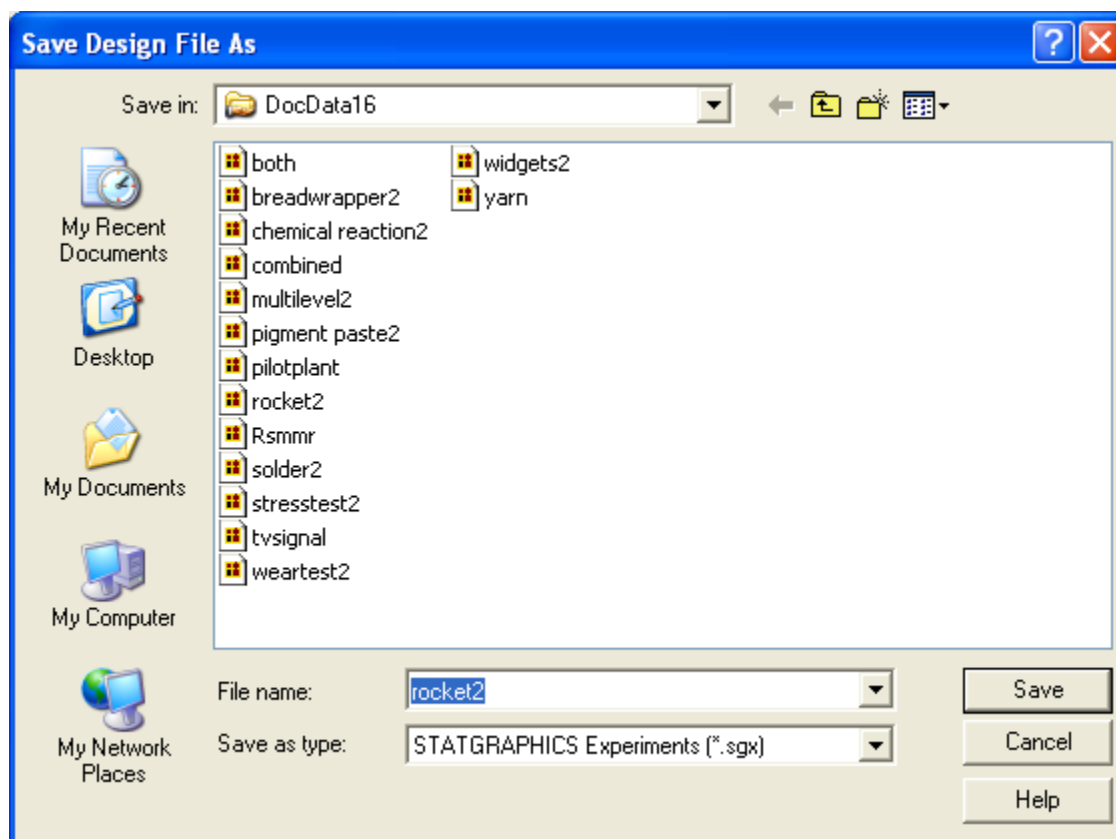
Returning to the ANOVA table shows that the additional 5 runs provide degrees of freedom for estimating pure error:

Source	D.F.
Model	7
Total Error	8
Lack-of-fit	3
Pure error	5
Total (corr.)	15

Saving the Design File

Step #7: Save experiment

Once the experiment has been created and any additional runs entered, it must be saved on disk. Press the button labeled Step 7 and select a name for the experiment file:



Design files are extended data files and have the extension *.sgx*. They include the data together with other information that was entered on the input dialog boxes.

To reopen an experiment file, select *Open Data File* from the *File* menu. The data will be loaded into the datasheet, and the *Experimental Design Wizard* window will be displayed.

Analyzing the Results

After the design file has been created and saved, the experiment would be performed. At a later date, once the results have been collected, the experimenter would return to STATGRAPHICS and reopen the saved design file using the *Open Data Source* selection on the main *File* menu. The results can then be typed into the response columns. The results for the example are displayed below:

<i>run</i>	<i>fuel</i> <i>(percent)</i>	<i>oxidizer</i> <i>(percent)</i>	<i>binder</i> <i>(percent)</i>	<i>burn rate</i> <i>(cm per second)</i>
1	50.0	20.0	20.0	32.5
2	30.0	40.0	20.0	54.5
3	30.0	20.0	40.0	64.0
4	40.0	30.0	20.0	44.0
5	40.0	20.0	30.0	63.2
6	30.0	30.0	30.0	94.0
7	36.6667	26.6667	26.6667	112.5
8	43.3333	23.3333	23.3333	67.1
9	33.3333	33.3333	23.3333	73.0
10	33.3333	23.3333	33.3333	87.5
11	50.0	20.0	20.0	37.9
12	30.0	40.0	20.0	32.5
13	30.0	20.0	40.0	78.5
14	36.6667	26.6667	26.6667	98.5
15	36.6667	26.6667	26.6667	103.6

Important Notes:

1. If more than one sample was taken at each set of experimental conditions, the data values should be entered into data tables B through Z. The summary statistics in data table A will then be automatically calculated from the other tables.
2. If any experiments were not performed, leave the corresponding cell blank. The program will recognize the imbalance in the design and handle it.
3. If any experimental runs were done at conditions different than originally planned, change the entries in the experimental factor columns to correspond to the values that were actually used.
4. If additional runs were performed, you may add them to the bottom of the datasheet. They will be included in the fit.

Step #8: Analyze data

Once the data have been entered, press the button labeled *Step #8* on the Experiment Design Wizard toolbar. This will display a dialog box listing each of the response variables:

Analysis Summary

The *Analysis Summary* displays information about the estimated effects:

[Analyze Mixture - burn rate](#)
File name: rocket.sgx

Estimated Full Model Effects for burn rate (cm per second)

Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Mean	72565.0	1	72565.0		
Linear	2395.93	2	1197.96	2.01	0.1770
Quadratic	5486.89	3	1828.96	9.83	0.0034
Special Cubic	1063.59	1	1063.59	13.92	0.0058
Error	611.369	8	76.4211		
Total	82122.8	15			


Full Model Results

Model	SE	R-Squared	Adj. R-Squared
Linear	24.4299	25.07	12.58
Quadratic	13.6421	82.48	72.74
Special Cubic	8.74191	93.60	88.81

The table shows:

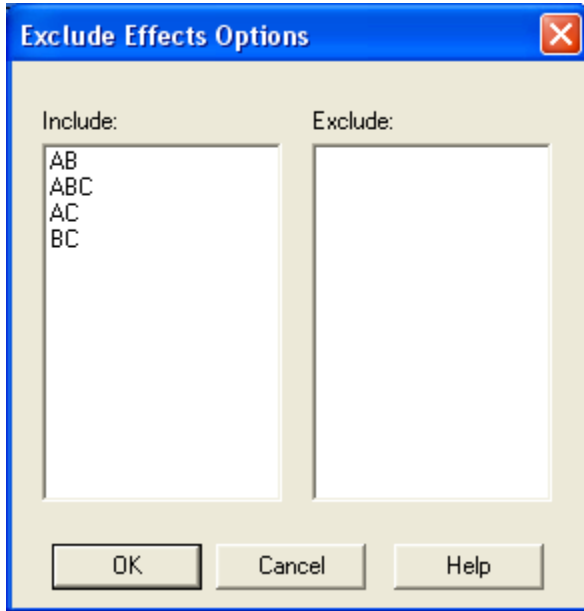
- Estimated Full Model Results** - an analysis of variance table for each model, showing the number of additional coefficients (df) needed by that model compared to the simpler model immediately above it. The *F-ratios* test the hypothesis that the additional terms in the more complicated model are not needed. Operating at a 5% significance level, the analyst would pick the most complicated model for which the P-value is less than 0.05. In the above case, that is the *special cubic* model. Had the P-value for the special cubic been greater than 0.05, it would not have been significantly better statistically than the quadratic model, which would then have been selected instead.
- Full Model Results** - shows the standard error of estimation, the R-squared statistic, and the adjusted R-squared. In general, the model with the smallest standard error and highest adjusted R-squared would be the model selected. In this case, the cubic model has an adjusted R-squared of 88.8%, which is considerably better than either of the other models.

Analysis Options



- Model to Fit:** model for which the results are displayed in the tables and graphs.

- **Exclude Button:** push to exclude terms from the fitted model. The following dialog box will be displayed:



Effects can be excluded from the model by double-clicking on them one at a time. Double-clicking on an effect in either of the two columns moves it to the other column.

ANOVA Table

A complete analysis of variance table is also available for the full model:

ANOVA for burn rate - Rocket Propellant Example					
Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Special Cubic Model	8946.42	6	1491.07	19.51	0.0002
Total error	611.36	8	76.42		
Total (corr.)	9557.78	14			

R-squared = 93.6035 percent
 R-squared (adjusted for d.f.) = 88.8062 percent
 Standard Error of Est. = 8.74185
 Mean absolute error = 5.22437
 Durbin-Watson statistic = 2.59519 (P=0.8685)
 Lag 1 residual autocorrelation = -0.306465

The ANOVA partitions the variance of the response into two components: one for the fitted model, and one for the experimental error. The ANOVA table shows:

- **Sum of Squares** - the sums of squares attributable to the model and to the error. The sum of squares for error is given by

$$S_{error} = \sum_{i=1}^n e_i^2 = \sum_{i=1}^n (y_i - \hat{y}_i)^2 \tag{11}$$

where e_i is the i -th residual, measuring the difference between the observed response for run i and the value predicted by the fitted model.

- **Df** - the degrees of freedom associated with each term.
- **Mean Square** - the mean square associated with each term, obtained by dividing the associated sum of squares by its degrees of freedom. The mean squared error (MSE) estimates the variance of the experimental error:

$$\hat{\sigma}^2 = MSE = \frac{S_{error}}{df_{error}} \quad (12)$$

- **F-Ratio** - an F ratio which divides the mean square of the model by the mean squared error:

$$F = \frac{MS_{MODEL}}{MSE} \quad (13)$$

- **P-Value** - the P-Value associated with testing the null hypothesis that the model does not explain any of the variability in the response. P-Values below a critical level (such as 0.05 if operating at the 5% significance level) indicate that the model is statistically significant at that significance level.
- **R-squared** - the percentage of the variability in the response variable that has been accounted for by the fitted model, calculated from

$$R^2 = 100 \left(1 - \frac{S_{error}}{S_{total}} \right) \% \quad (14)$$

R-squared ranges from 0% to 100% and measures how well the model fits the observed response data.

- **R-squared (adjusted for d.f.)** - the adjusted R-squared, which accounts for the number of degrees of freedom in the fitted model. In situations such as the current one where the number of coefficients in the fitted model is large relative to the total number of runs, the ordinary R-squared statistic may overstate the ability of the fitted model to predict the response. The adjusted R-squared compensates for this effect by

$$R^2_{adj} = 100 \left[1 - \left(\frac{n-1}{n-p} \right) \frac{S_{error}}{S_{total}} \right] \% \quad (15)$$

where p is the number of estimated coefficients in the fitted model.

- **Standard error of est.** - the estimated standard deviation of the experimental error, given by

$$\hat{\sigma} = \sqrt{MSE} \quad (16)$$

This value is used when constructing prediction intervals for the response.

- **Mean absolute error** - the average of the absolute values of the residuals, given by

$$MAE = \frac{\sum_{i=1}^n |e_i|}{n} \quad (17)$$

This value indicates the average error in predicting the observed response using the fitted model.

- **Durbin-Watson statistic** - a statistic calculated from the residuals according to


$$DW = \frac{\sum_{i=1}^{n-1} (e_{i+1} - e_i)^2}{\sum_{i=1}^n e_i^2} \quad (18)$$

The Durbin-Watson statistic measures serial correlation in the residuals to determine whether there is any dependence between successive observations. In this case, it could detect drifts over the course of the experiment. A small P-value would indicate that the analyst should take a close look at the residuals to look for any trends, which may be done using the *Diagnostic Plots* graph option.

- **Lag 1 residual autocorrelation** – a measure of the serial correlation in the residuals, based on the order that the runs appear in the datasheet. Values close to 0 correspond to random residuals.

In the example above, the P-Value for the model is very small, indicating that the special cubic model is a useful predictor of *burn rate*. The adjusted R-squared indicates that nearly 90% of the variability in *burn rate* has been explained by the model.

Pane Options



- **Include Lack-of-Fit Test:** If checked, a line will be added to the ANOVA table to determine whether the current model adequately represents the observed data. Note: this option has no effect unless there are replicate experimental runs at identical settings of the experimental factors.

The resulting ANOVA table is shown below:

ANOVA for burn rate - Rocket Propellant Example					
Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
Special Cubic Model	8946.42	6	1491.07	16.13	0.0039
Lack-of-fit	149.249	3	49.7495	0.54	0.6764
Pure error	462.112	5	92.4223		
Total (corr.)	9557.78	14			

R-squared = 93.6035 percent
 R-squared (adjusted for d.f.) = 88.8062 percent
 Standard Error of Est. = 9.61365
 Mean absolute error = 5.22437
 Durbin-Watson statistic = 2.59519 (P=0.8685)
 Lag 1 residual autocorrelation = -0.306465

Note the lines labeled *Lack-of-fit* and *Pure error*, which provide two separate estimates of the experimental error sigma:

1. *Pure error*: an estimate calculated by pooling the variance within sets of observations at identical levels of X. It is “pure” in the sense that it estimates the experimental error σ whether or not the proper model has been selected.
2. *Lack-of-fit*: an estimate calculated from the deviation between the average response for each group of replicate values and the values predicted by the fitted model. If the model is not correct, this estimates σ plus a positive quantity that measures the lack-of-fit of the selected model.

The P-Value in the lack-of-fit line may be used to test the hypothesis that the current model is adequate. A small P-Value would indicate an inadequate model. In the current example, the P-Value is well above 0.05, so the selected model appears to be adequate.

Model Results

The equation of the fitted model is obtained by selecting *Model Results*:

Special Cubic Model Fitting Results for burn rate				
Parameter	Estimate	Standard Error	T Statistic	P-Value
A:fuel	35.4946	6.07193		
B:oxidizer	42.7756	6.07193		
C:binder	70.3613	6.07193		
AB	16.0213	38.2911	0.418408	0.6867
AC	36.3356	38.2911	0.948931	0.3704
BC	136.821	38.2911	3.57319	0.0073
ABC	854.962	229.174	3.73063	0.0058

Interpretation
 This pane displays the equation of the fitted special cubic model. The equation of the fitted model is

$$\text{burn rate} = 35.4946*\text{fuel} + 42.7756*\text{oxidizer} + 70.3613*\text{binder} + 16.0213*\text{fuel}*\text{oxidizer} + 36.3356*\text{fuel}*\text{binder} + 136.821*\text{oxidizer}*\text{binder} + 854.962*\text{fuel}*\text{oxidizer}*\text{binder}$$
 where the values of the components are specified in pseudo-components. To have STATGRAPHICS evaluate this function, select Predictions from the list of Tabular Options. To plot the function, select Response Plots from the list of Graphical Options.

The table shows:

- **Estimate** - the estimate of each coefficient in the selected model, where the model is expressed in pseudo-components. If the *StatAdvisor* is requested, the model is written in full.
- **Standard error** - the estimated standard error of each coefficient.
- **T statistic** - the results of a t-test calculated by

$$t = \text{coefficient} / \text{standard error} \tag{19}$$

- **P-Value** - the statistical significance of each t-test. Small P-Values (less than 0.05) indicate coefficients that are significantly different from 0.

The results of this test could be used to simplify the model if desired. As mentioned above, individual terms may be excluded using the *Exclude* button on the *Analysis Options* dialog box. In this case, neither AB nor AC is significant. However, they are both involved in the significant ABC effect and would typically not be removed.

Correlation Matrix

The correlation matrix displays the estimated correlation between the coefficients in the fitted regression model:

	A:fuel	B:oxidizer	C:binder	AB	AC	BC	ABC
A:fuel	1.0000	0.0130	0.0130	-0.3510	-0.3510	0.0247	0.0799
B:oxidizer	0.0130	1.0000	0.0130	-0.3510	0.0247	-0.3510	0.0799
C:binder	0.0130	0.0130	1.0000	0.0247	-0.3510	-0.3510	0.0799
AB	-0.3510	-0.3510	0.0247	1.0000	0.1065	0.1065	-0.4750
AC	-0.3510	0.0247	-0.3510	0.1065	1.0000	0.1065	-0.4750
BC	0.0247	-0.3510	-0.3510	0.1065	0.1065	1.0000	-0.4750
ABC	0.0799	0.0799	0.0799	-0.4750	-0.4750	-0.4750	1.0000

The correlations are estimated from the variance-covariance matrix of the coefficients, given by

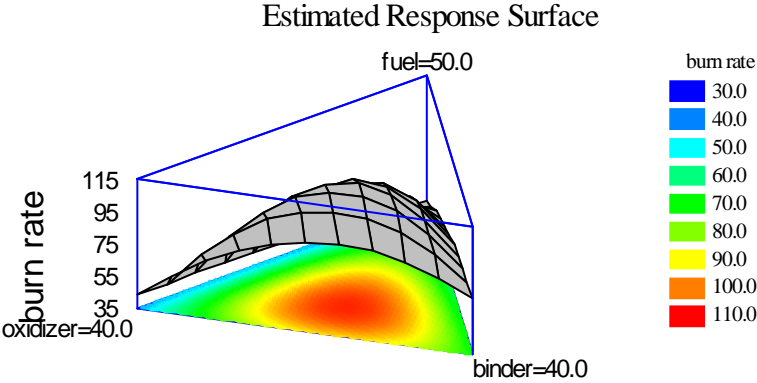
$$s^2(b) = \text{MSE}(X'X)^{-1} \tag{20}$$

In most mixture designs, there will be non-zero correlations between the coefficients.

Response Plots

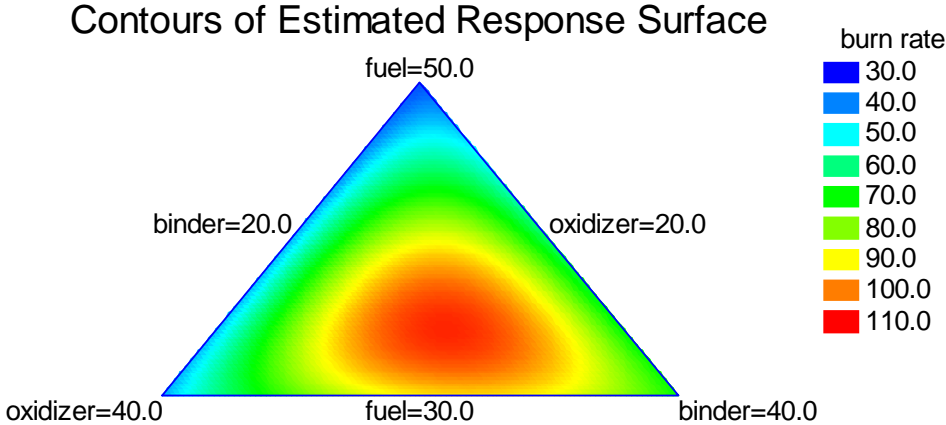
The list of graphs available in the *Analyze Data* procedure contains two selections labeled *Response Plots* that allow the predicted values of the response to be plotted. By default, the first selection creates a response surface plot and the second a contour plot. However, each is controlled by the same *Pane Options* dialog box, which allows the analyst to display either two type of plot.

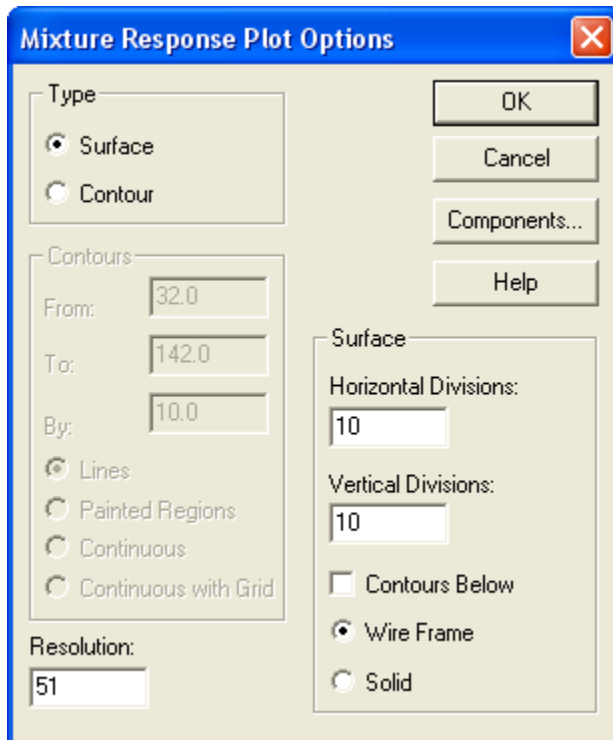
The surface plot displays a plot of the predicted response as a function of any three of the experimental components, with the other components held at selected values. For example, the plot below shows *burn rate* as a function of *oxidizer*, *binder* and *fuel*:



The height of the surface represents the predicted value \hat{Y} , which is plotted over the range of the experimental components.

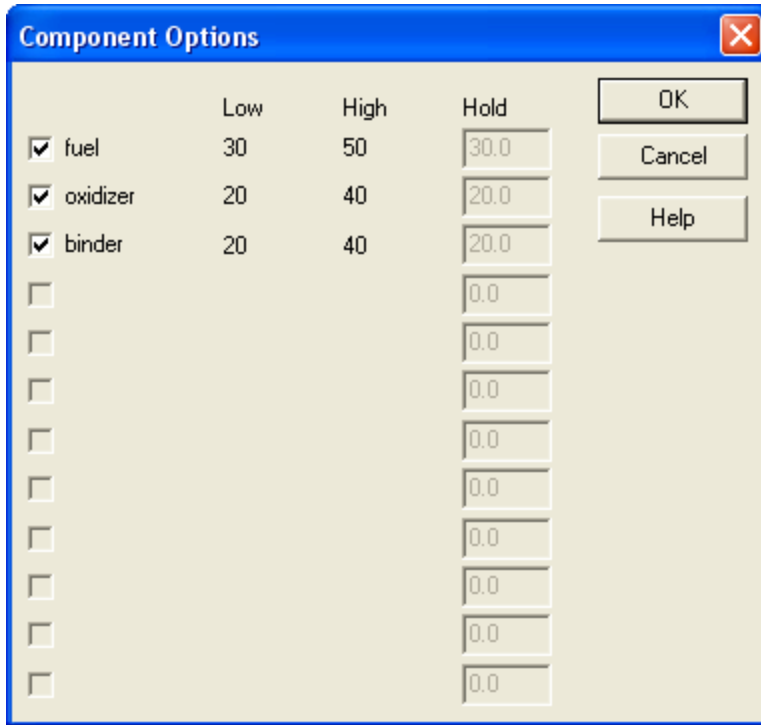
Contour plots draw lines or colors to indicate values of the predicted response. For example, the plot below codes the range of the predicted values for *burn rate* into colors extending from blue at 35 to red at 110:



Pane Options

- **Type:** type of response plot to create.
- **From:** location at which the first contour line is drawn, or the start of the first region.
- **To:** location at which the last contour line is drawn, or the end of the last region.
- **By:** spacing between contour lines or regions.
- **Lines:** contour lines are drawn at levels of the predicted response, as on a topographical map.
- **Painted Regions:** regions are drawn covering various ranges of the predicted response.
- **Continuous:** values are indicated by a continuous range of colors.
- **Resolution:** defines the resolution m of an m -by- m grid of predicted values which is used to draw the surface and contour lines. Increasing the resolution may improve the smoothness and definition of the plots, at the expense of computer time and memory.
- **Horizontal Divisions:** the number of divisions along the first experimental axis. This determines how many vertical lines will be drawn on the surface plot.
- **Vertical Divisions:** the number of divisions along the second experimental axis. This determines how many horizontal lines will be drawn on the surface plot.
- **Contours Below:** requests that a contour plot, of type specified below, be drawn in the bottom face of the 3-D plot.

- **Wire Frame:** requests that the surface be drawn using cross-hatched lines. This is often the most effective choice for black-and-white presentation.
- **Solid:** requests that the surface be drawn using a solid color.
- **Components button:** specifies the components to be plotted on each axis and the levels at which the other components will be held:



Three factors must be checked.

Predictions

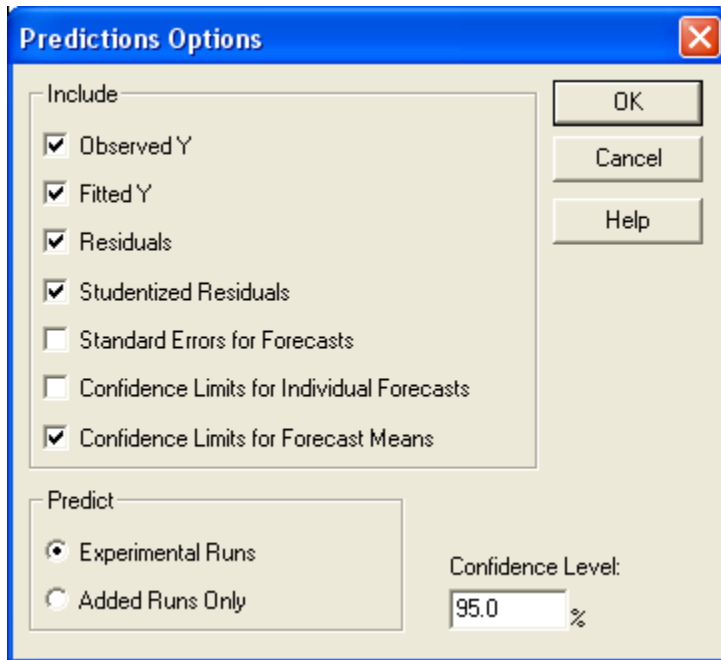
The *Predictions* pane may be used to generate predictions from the fitted model:

Estimation Results for burn rate				
	<i>Observed</i>	<i>Fitted</i>	<i>Lower 95.0% CL</i>	<i>Upper 95.0% CL</i>
<i>Row</i>	<i>Value</i>	<i>Value</i>	<i>for Mean</i>	<i>for Mean</i>
1	32.5	35.4946	18.3296	52.6596
2	54.5	42.7756	25.6106	59.9406
3	64.0	70.3613	53.1963	87.5263
4	44.0	43.1404	19.5715	66.7093
5	63.2	62.0119	38.4429	85.5808
6	94.0	90.7738	67.2049	114.343
7	112.5	102.23	89.3239	115.136
8	67.1	67.9691	57.307	78.6313
9	73.0	79.9833	69.3212	90.6455
10	87.5	95.4691	84.8069	106.131
11	37.9	35.4946	18.3296	52.6596
12	32.5	42.7756	25.6106	59.9406
13	78.5	70.3613	53.1963	87.5263
14	98.5	102.23	89.3239	115.136
15	103.6	102.23	89.3239	115.136
16		95.9457	84.0766	107.815

The table may include all rows in the datasheet, or only rows for which the value of the response variable Y has not been entered. The latter feature allows the analyst to make predictions at combinations of X that were not included in the experiment. For example, the above table shows the result of adding a 16th row with $fuel = 35$, $oxidizer = 30$, and $binder = 25$. The predicted value of $burn\ rate$ is 95.95. The 95% confidence interval for the mean value of $burn\ rate$ at that same combination of the components ranges from 86.26 to 105.63.

One other noticeable entry in the above table is the Studentized residual for row #2. The Studentized residual measures the difference between the observed response and the predicted response, in units of its standard error, when the observation in question is not used to fit the model. The Studentized residual for observation #2 equals 2.3. Values in excess of 3.0 are unusual and would typically require further scrutiny.

Pane Options



- **Include:** items to include in the table:
 1. *Observed Y* - the observed response values Y_i .
 2. *Fitted Y* - the predicted values \hat{Y}_i calculated from the fitted model.
 3. *Residuals* - the residuals e_i .
 4. *Studentized Residuals* - a type of standardized residual, where each residual is divided by an estimate of its standard error. STATGRAPHICS computes *Studentized deleted residuals*, in which each observation is removed one at a time and the model refit without that data value. The deleted residual then equals the observed response minus the value predicted from a model fit *without* that observation, i.e.,

$$d_i = Y_i - \hat{Y}_{(i)} \quad (21)$$

The Studentized residual is calculated from

$$e_i^* = \frac{d_i}{s(d_i)} \quad (22)$$

where

$$s^2(d_i) = MSE_{(i)} \left(1 + X_i' (X_{(i)}' X_{(i)})^{-1} X_i \right) \quad (23)$$

The deleted residuals should follow a t distribution with $n - p - 1$ degrees of freedom, where p is the number of estimated coefficients in the fitted model.

1. *Standard Errors for Forecasts* - the standard error for new observations at a selected combination of the experimental components X_h , given by

$$\sqrt{MSE \left(1 + X_h' (X' X)^{-1} X_h \right)} \quad (24)$$

2. *Confidence Limits for Individual Forecasts* - confidence limits for new observations at a selected combination of the experimental components X_h , given by

$$\hat{Y}_h \pm t_{n-p} \sqrt{MSE \left(1 + X_h' (X' X)^{-1} X_h \right)} \quad (25)$$

3. *Confidence Limits for Forecast Means* - confidence limits for the mean response at a selected combination of the experimental components X_h , given by

$$\hat{Y}_h \pm t_{n-p} \sqrt{MSE \left(X_h' (X' X)^{-1} X_h \right)} \quad (26)$$

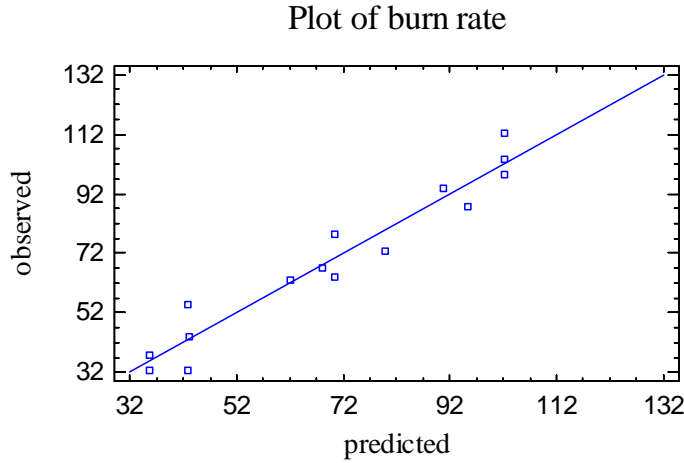
- **Predict** - whether forecasts are displayed for all of the runs in the experiment data file, or only for runs that have a missing value in the response column.
- **Confidence level** - the confidence levels for the intervals.

Diagnostic Plots

Several plots are also provided under *Diagnostic Plots* to examine the residuals from the fitted model. The *Pane Options* dialog box displays the various choices, which include the following:

Observed versus Predicted


This plot displays the observed response Y_i versus the fitted values \hat{Y}_i , together with a diagonal line:



If the model fits well, the values should lie close to the line, as in the example above. Curvature around the line may suggest the need to transform the values of Y_i using a logarithm or similar function.

Residual versus Predicted

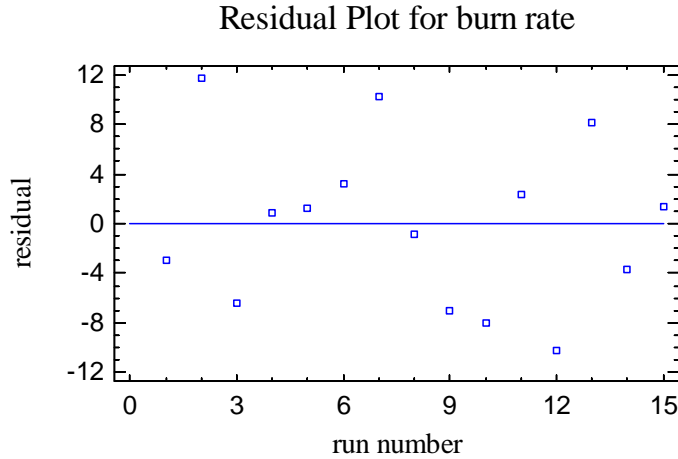
This plot displays the residuals e_i versus the fitted values \hat{Y}_i , with a horizontal line at zero:



The residuals should vary randomly around the line. Changes in the magnitude of the residuals from left to right may signal that the variance of the experimental error varies with the mean level of the response. Such *heteroscedasticity* may frequently be eliminated by a variance-stabilizing transformation such as a logarithm or a square root.

Residuals versus Run Order

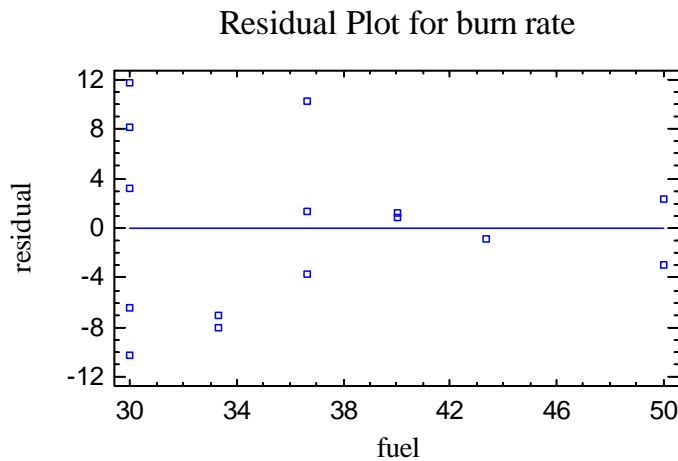
This plot displays the residuals e_i versus run number i , with a horizontal line at zero:



Any non-random pattern may indicate a time trend or other effect. In such cases, addition of a factor to account for the change may improve the fit of the model.

Residuals versus Factor

This plot displays the residuals e_i versus the observed values of a selected experimental component:




Any curvature around the line may suggest the need for a model with quadratic effects. The above plot suggests that the variability may decrease as the percentage of fuel increases.

Normal Probability Plot of Residuals

This plot displays the residuals e_i versus quantiles of a normal distribution, with an optional fitted line as reference:

Normal Probability Plot for Residuals

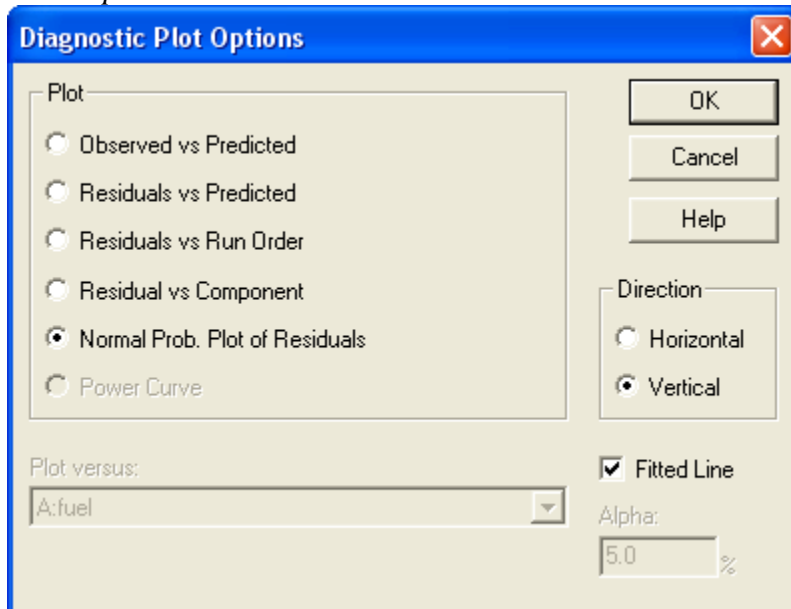


If the experimental error follows a normal distribution, the points should lie along a straight line, as in the plot above.

Power Curve

The power curve is not available for mixture designs.

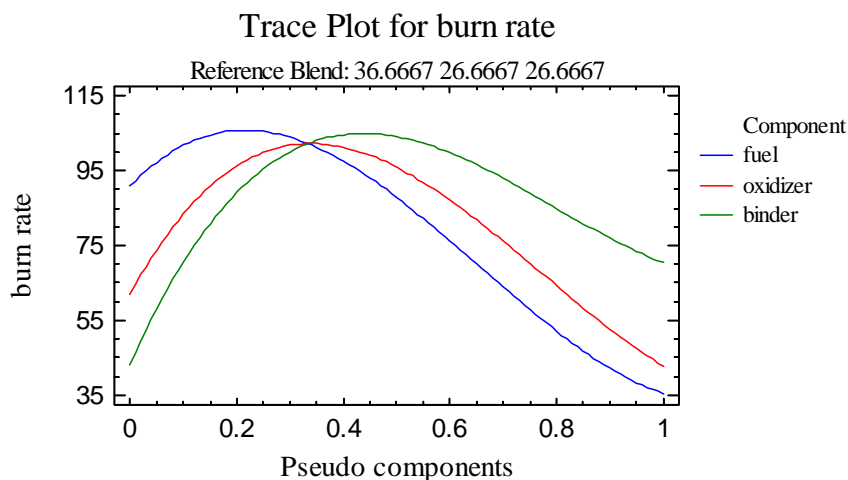
Pane Options



- **Plot:** the type of plot to be created.
- **Plot versus:** selects the experimental factor to be shown in the plot, for those plots where a factor is needed.
- **Direction:** defines the orientation of the normal probability plot.
- **Fitted Line:** specifies whether a line should be fit to the data on the normal probability plot.
- **Alpha:** specifies the α -risk associated with the *Power Curve*. Not used for mixture designs.

Trace Plot

The *Trace Plot* provides a useful diagnostic tool for assessing the importance of the various components:



Beginning at a reference blend (initially the centroid of the design space), it shows how the response changes as each component is increased or decreased. As the selected component is changed, the other components are either decreased or increased proportionately. The shape of the trace shows the relative importance of the components. In the above plot, *binder* and *fuel* have steeper slopes at the reference blend than *oxidizer* and therefore are more influential in affecting the burn rate. A reference blend other than the centroid, such as the optimum combination derived earlier, may be selected using *Pane Options*:

Pane Options


	Low	High	Reference Blend	
fuel	30	50	36.6667	OK
oxidizer	20	40	26.6667	Cancel
binder	20	40	26.6667	Help
			0.0	
			0.0	
			0.0	
			0.0	
			0.0	
			0.0	
			0.0	
			0.0	
			0.0	

Enter a reference blend to be used as the base blend in the trace plot.

Optimization

Step #9: Optimize responses

Once a statistical model has been developed for each response, the analyst may now determine what combination of factors will yield the best results. Pressing the button labeled *Step #9* will search for the best conditions throughout the experimental region. When the optimization is complete, a message similar to that shown below will be displayed:



The dialog box indicates the “Desirability” of the final result, based on a metric designed to balance competing requirements of multiple responses (see the document titled DOE Wizard for full details). The value displayed in this case indicates that the predicted *burn rate* at the optimum factor settings has a desirability of 100%, meaning that it exactly equaled the target value/

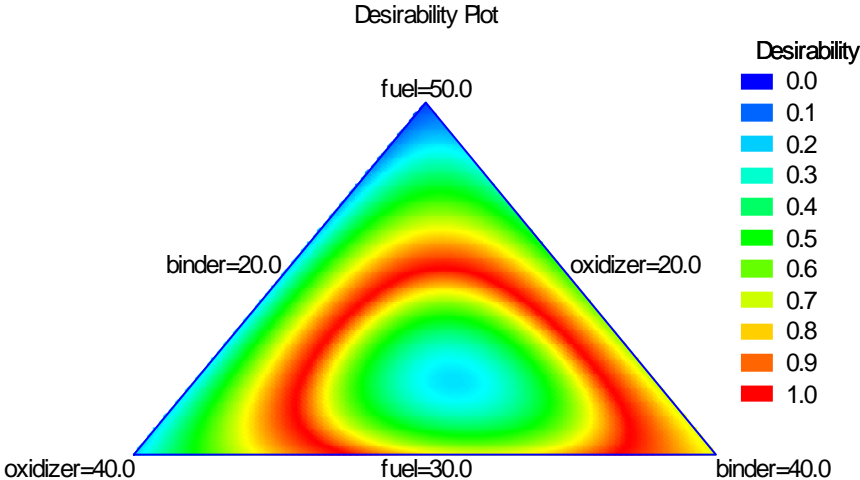
If you press *OK*, additional information will be added to the main *DOE Wizard* window:

<u>Step 9: Optimize the responses</u>				
Response Values at Optimum				
<i>Response</i>	<i>Prediction</i>	<i>Lower 95.0% Limit</i>	<i>Upper 95.0% Limit</i>	<i>Desirability</i>
burn rate	85.0	76.1446	93.8554	1.0

<u>Factor Settings at Optimum</u>	
<i>Factor</i>	<i>Setting</i>
fuel	33.2135
oxidizer	32.786
binder	24.0006

The table shows the estimated response at the optimal settings of the experimental factors. For the rocket propellant data, it is estimated that the mean *burn rate* will equal 85.0 when the factors are set at *fuel* = 33.214%, *oxidizer* = 32.786%, and *binder* = 24.001%. The 95% confidence interval for the mean ranges between 76.14 and 93.86.

If you push the *Tables and Graphs* button on the analysis toolbar, you can display the estimated desirability throughout the experimental region. A useful display for the current design is the contour plot shown below:



The region of highest desirability is shown as dark red. There are lots of different combinations of the 3 components that could achieve the desired burn rate. To select a specific combination, such as that will minimize cost, *cost* could be added to the design as a second response variable.