Combined Split-Plot Design

Introduction

This tutorial illustrates Design-Expert® software tools for applying split-plot design to experiments that combine both mixture and process factors.



Baking cake is the perfect example to test out the tools in Design-Expert (DX) for the combined split-plot design. It involves all of the elements of a good experiment, from mixing together various portions of ingredients, to treating the mixture at temperature, and even adding another mixture (frosting) to the mix (pun intended). One particularly popular dessert, especially in the southern United States, is the Lady Baltimore Cake—rich and delicious with a fluffy frosting full of nuts and raisins. Supposedly the cake was first baked by Alicia Rhett Mayberry of Charleston, South Carolina for novelist Owen Wister, best known for *The Virginian*, published in 1902. Wister described the cake in his next book—*Lady Baltimore* (1906).

Many recipes for the Lady Baltimore Cake can be found via internet, e.g., the one pictured here detailed by wikiHow at www.wikihow.com/Bake-a-Lady-Baltimore-Cake. These recipes vary, of course. For example, the wikiHow instructions specify the use of all purpose (plain) flour whereas the recipe shown below demands specialized cake flour.

Cake	Frosting and Filling
3 cups sifted cake flour	2 egg whites, unbeaten
3 teaspoons baking powder	1 ½ cups sugar
½ teaspoon salt	5 tablespoons water
½ cup butter	1 ½ teaspoons light corn syrup
1 ½ cups sugar	½ teaspoon vanilla
1 ¼ cups milk	6 dried figs; Candied Cherries
1 teaspoon vanilla	½ cup raisings
4 egg whites	½ cup nuts, chopped

Immediately an experimentally inquisitive mind must question whether it really is necessary to use the fancier flour. Perhaps a combination of flours might do. Also, it may help to adjust the ratio of flour to sugar. Let's investigate a blend of these three ingredients—all-purpose flour, cake flour and sugar--to develop a better-tasting cake. The amounts will be expressed in ounces, holding all other ingredients constant.

Cake is great but it really comes up short without a good frosting. To tailor this second mixture formulation, the key ingredients of water, corn syrup, and vanilla will be varied, again holding the other ingredients constant.

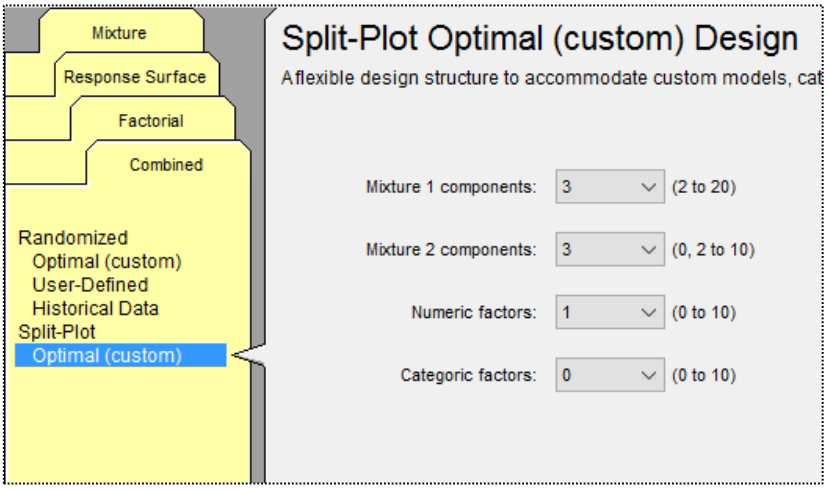
Last but not least, the amount of frosting and filling (F&F) will be varied (the Lady Baltimore Cake features the same mixture in and out). This will be measured in cups—a numerical factor.

Split-Plot and Experimental Setup

As we just explained in the Introduction, the Lady Baltimore Cake experiment encompasses two mixture formulations (cake and frosting) and an amount (F&F). That makes this a "combined" design, which, as you can imagine, results in many runs to provide enough data for a response surface method (RSM) optimization. Luckily, our bakery has an oven that can bake 12 cakes, i.e., a dozen, at a time, in conjunction with a large mixer that can mix that much batter. Obviously, therefore, it will be most convenient to bake the cakes in batches of a dozen each. That's where the split-plot design steps in. Normally, an experiment would be completely randomized and thus require a new cake recipe for every run, but, as you will soon see, the split plot sorts cake batters into convenient groups for hard-to-change factors (HTC) within which will be randomized the easy-to-change (ETC) factors. Let's get going on this combined design so you can see how to go about making the best Lady Baltimore Cake in the land.

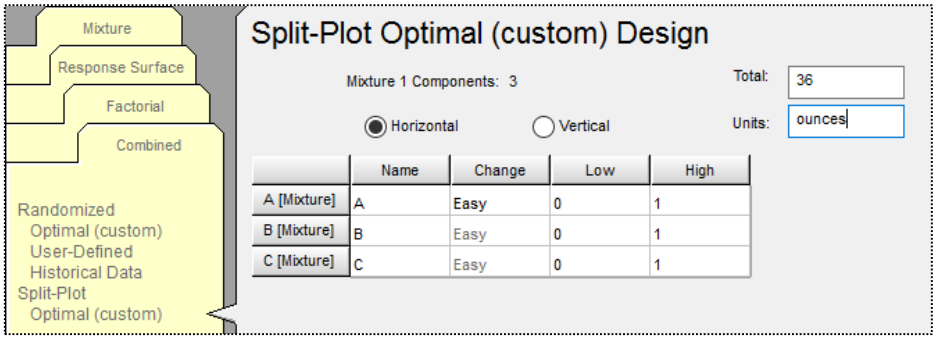
Design the Experiment

To set up the experiment, open the DX software. Then click the blank-sheet icon 🗋 on your toolbar (or choose **File, New Design**). Next click on the **Combined** tab and in the Split-Plot section near the bottom choose the **Optimal (custom)** design. Select **3** from the Mixture 1 components droplist. Then, choose **3** for Mixture 2 components, and **1** for the Numeric factors. Click **Continue** in the bottom right corner to advance to the next page.



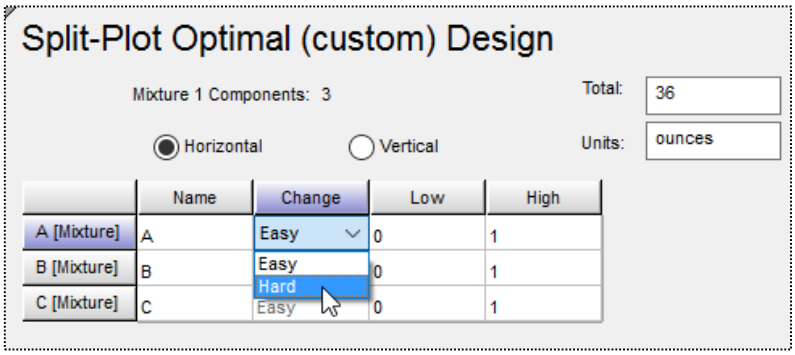
Setting up the combined split-plot design with two mixtures and 1 numeric factor.

Here, you will enter the components for Mixture 1, the cake recipe. Only all-purpose flour, cake flour, and sugar are being investigated. That portion of the recipe comes to 36 total ounces. Enter **36** in the Total box and **ounces** for Units.



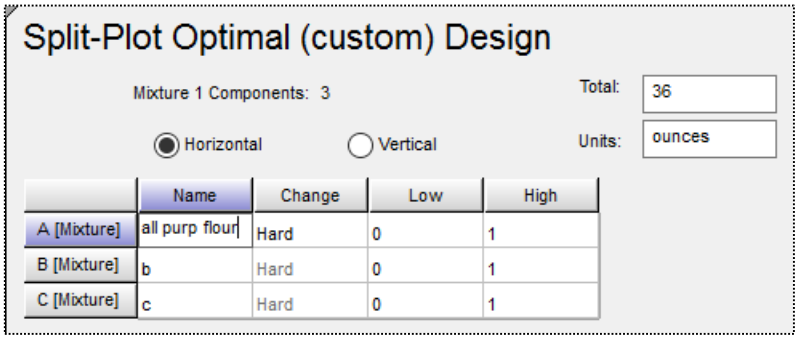
Entering the total and units for Mixture 1, the cake recipe.

That constant total of 36 ounces will be added to the rest of the unchanged recipe, so the proportions are the same every run. The cake mixture is the HTC portion of the design, so switch the change column in A from Easy to **Hard** using the droplist. In a mixture, all of the components need to be either Easy or Hard, so changing the first one changes them all to Hard.



Change the cake mixture to HTC

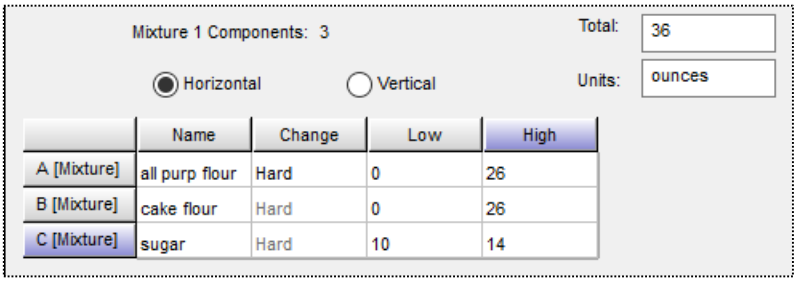
Click back on the component A name to change it to **all purp flour** (all-purpose flour). Note that all the components went to lower-case. HTC factors are lowercase in DX, differentiating them from the usual uppercase labels for ETC factors. This will really come in handy when manipulating the graphs and labels.



Change component a to all purpose ("purp") flour, note the lowercase for the HTC factors

Tab over and enter **0** for the Low and **26** for the high. Then enter **cake flour** as component b with a low of **0** and high of **26,** and **sugar** for component c with a low of **10** and high of **14**. Varying the levels of these components in relation to the rest

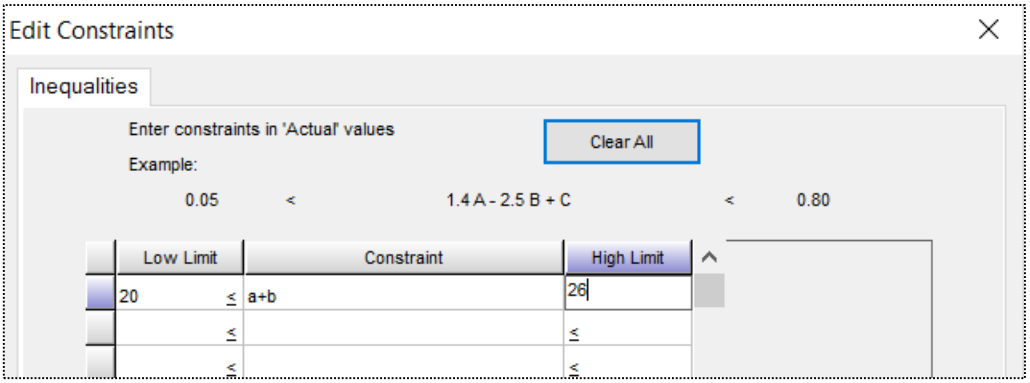
will determine if the cake flour is necessary and what level of sugar is best. The design should now look like this.



	Name	Change	Low	High
A [Mixture]	all purp flour	Hard	0	26
B [Mixture]	cake flour	Hard	0	26
C [Mixture]	sugar	Hard	10	14

Cake mixture components and levels.

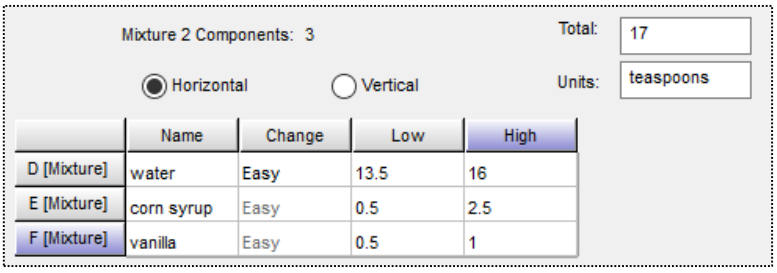
There is one part of the recipe that shouldn't be strayed from too much: the total amount of flour. To ensure there's enough flour, add a constraint by clicking on the **Edit Constraint** button near the bottom of the page. The total flour consists of **a+b**, so enter that into the middle constraints column. To maintain the right amount of flour, enter a Low Limit of **20** and a High Limit of **26**.



Enter the constraint to keep total flour between 20 and 26 ounces

Press **OK** and then **Continue** to move on to Mixture 2.

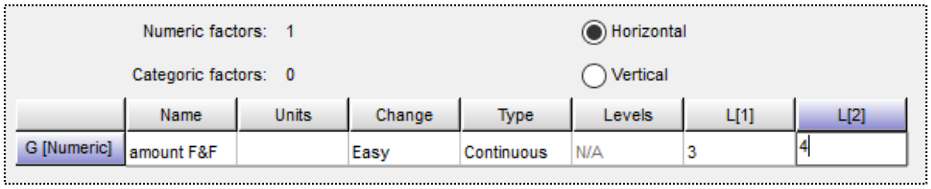
Enter the total for the filling and frosting mixture of **17** and units of **teaspoons**. We're only experimenting on the small (potent) ingredients of the recipe. This mixture is ETC, so no need to adjust the change column. Change the component D name to **water**, with a low of **13.5** and high of **16**. Component B should be **corn syrup** with a low of **0.5** and high of **2.5**, and component C is **vanilla** with a low of **0.5** and high of **1**.



	Name	Change	Low	High
D [Mixture]	water	Easy	13.5	16
E [Mixture]	corn syrup	Easy	0.5	2.5
F [Mixture]	vanilla	Easy	0.5	1

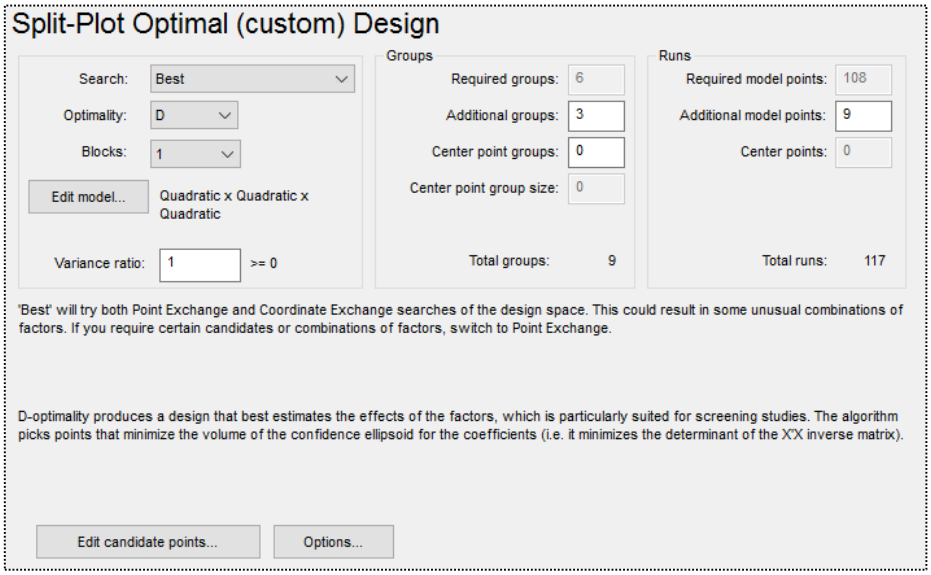
The component names and levels for frosting and filling

Click **Continue** to move on to the numeric factor. This is the numeric factor. Change the name to **amount F&F** (filling and frosting), with a low level, L(1), of **3** and a high level, L(2), of **4**.



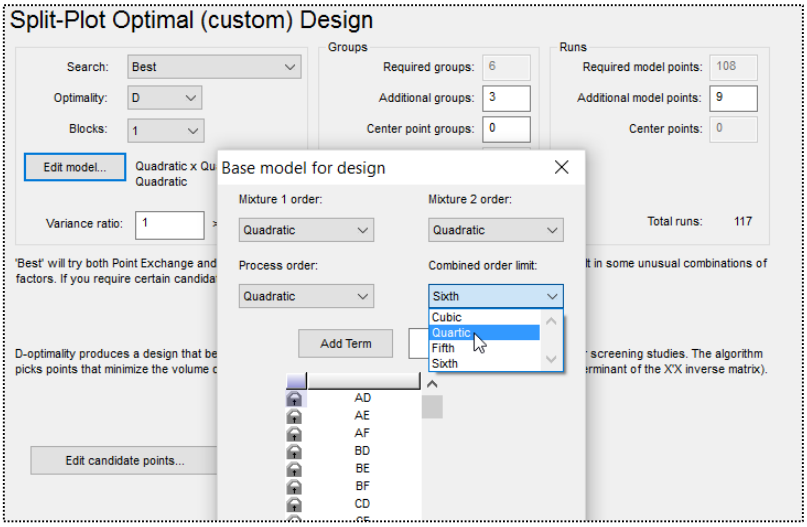
Entering the name and levels for the numeric factor

Click **Continue** and the design options appear. The optimal design has many options to choose from, and with the combined design there are even more. Look near the "Edit Model…" button and you see "Quadratic x Quadratic x Quadratic", so Mix 1 model is Quadratic, as is the Mix 2 model and the numerical model.



Entering the name and levels for the numeric factor

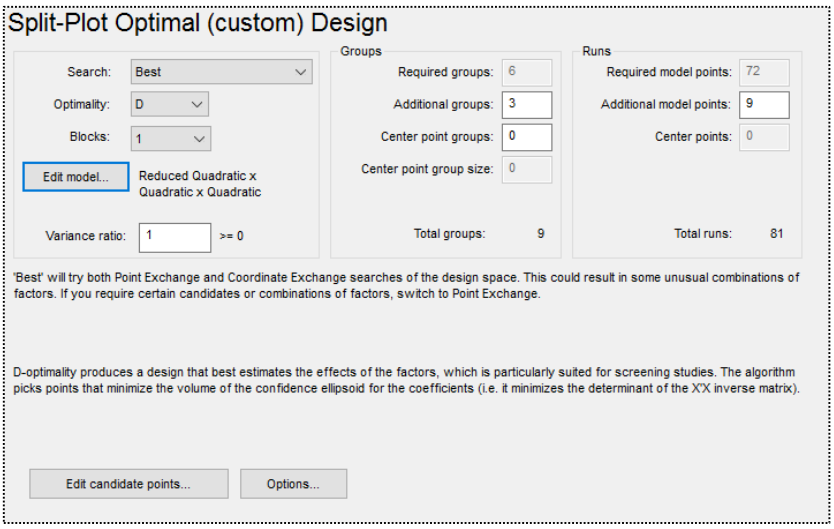
The combined model will multiply the terms from each of these models together, resulting in 108 coefficients, or 108 required model points (see the upper right). That combined model is very complex and allows a very intricate model that may be overkill. To save runs, click the **Edit model…** button. One way to save runs would be to change the model order of one of the individual models, Mix 1, Mix 2, or the process model. Choosing a simpler model, say linear would result in fewer terms, and runs. However, instead of simplifying the individual models, getting rid of some of the very high order interactions caused by the multiplication of these models is a better bet. To do that, change the Combined order limit: to **quartic**.



Change the combined model order to quartic to eliminate higher order interactions and reduce runs

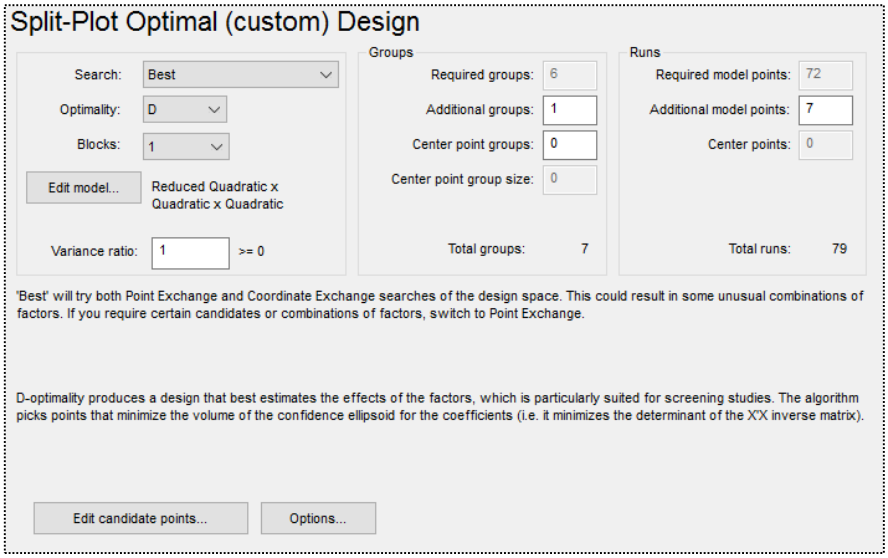
This will take out all of the terms in the model that involve fifth or sixth order terms, things like $ABEFG^2$, or $ADFG^2$, with are sixth and fifth order, respectively. These aren't absolutely necessary to get a good picture of the system, and eliminating them will save some runs. Only 4th order (quartic) terms will remain.

Press **OK** and the required model points goes down from 108 to 72, a nice savings.



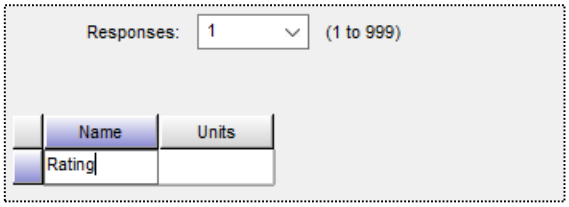
With the simpler model, there are only 72 required model points

Now, look at the Groups column in the middle. These are the "whole plot groups", where the HTC factor levels are held constant. In other words, these groups contain runs that will be mixed together and baked in one big batch, many cakes at a time. By default, there are 9 groups, so with 81 total runs, that's 9 cakes per group. Remember, 12 cakes can be baked at once, so reduce the Additional groups to **1** and press the **tab** key to update. Say **no** to the warning to keep only seven groups. That leaves a total of 79 runs in 7 groups, or a little more than 11 cakes per group, making much better use of the large oven and mixer.



Changing the additional groups to 1 results in 79 total runs and 7 groups

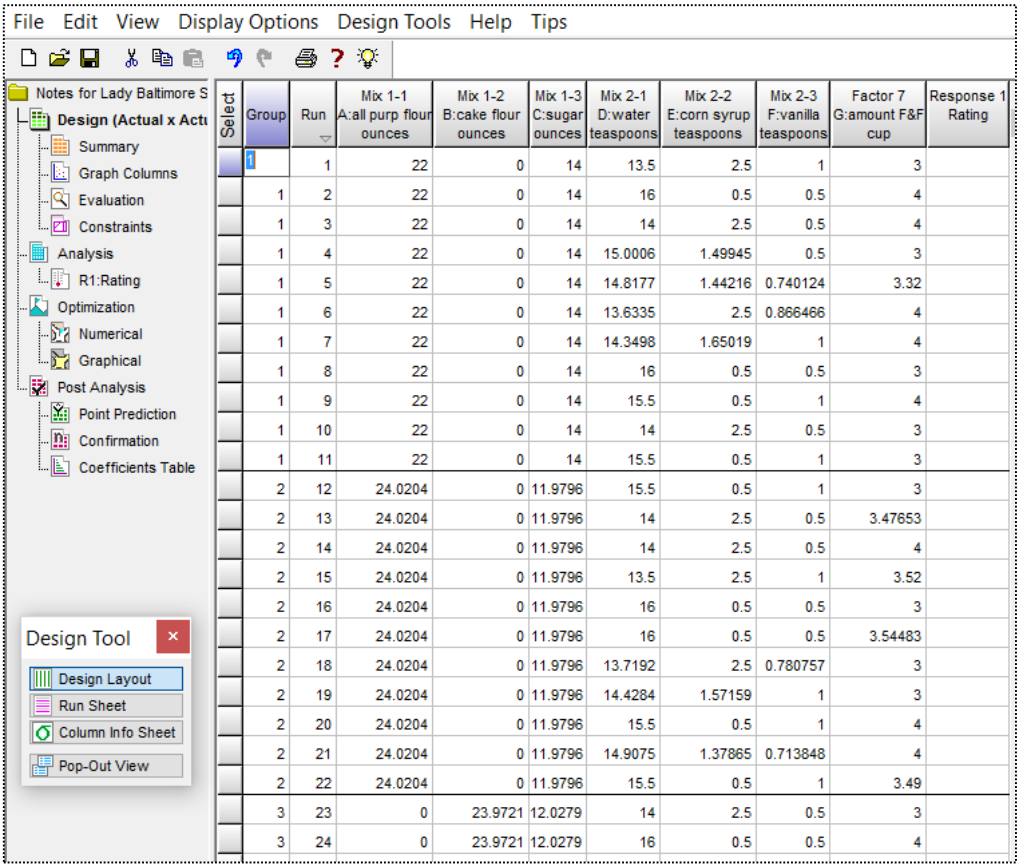
Click **Continue**. For the R1 Name, enter **Rating**, which in this case will be measured on a 100 point scale—the higher the better the taste.



Change the response name to Rating

Click **Finish** to build the design. It may take a few minutes. The program is going through many trials (20 by default) to pick the right set of runs to fit the model terms as precisely and efficiently as possible. After the iterations, the best design (judged by the statistical criteria chosen in the options) is presented.

You will get a warning to reset the factor levels between groups. Just click **OK** to bypass this warning, for now. The runs on your screen will most likely be different due to the randomizing (where not restricted) of the design. The first two groups are shown below. For group 1, cakes will be baked with 22 ounces of all-purpose flour, 0 ounces of cake flour, and 14 ounces of sugar. Coincidentally, group 2 contains another set of runs with no cake flour, but a little less sugar. The great thing about the split-plot design is that these groups contain 11 straight runs with the same cake recipe. That allows 11 cakes to be mixed and baked all in one big batch, an enormous time savings.



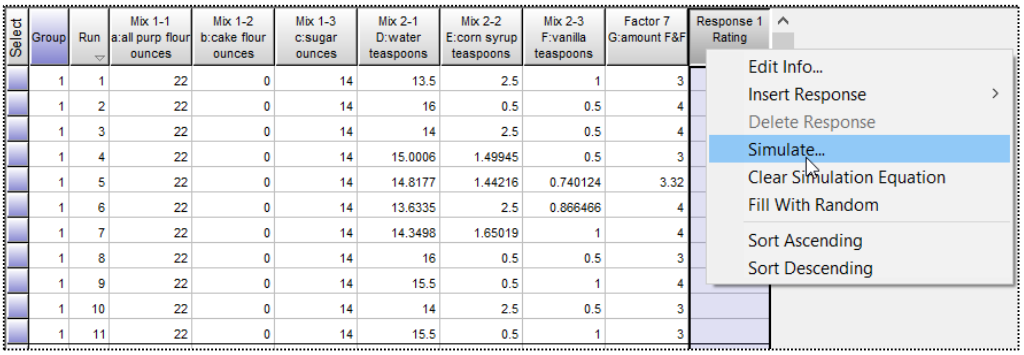
Select	Group	Run	Mix 1-1 A:all purp flour ounces	Mix 1-2 B:cake flour ounces	Mix 1-3 C:sugar ounces	Mix 2-1 D:water teaspoons	Mix 2-2 E:corn syrup teaspoons	Mix 2-3 F:vanilla teaspoons	Factor 7 G:amount F&F cup	Response 1 Rating
	1	1	22	0	14	13.5	2.5	1	3	
	1	2	22	0	14	16	0.5	0.5	4	
	1	3	22	0	14	14	2.5	0.5	4	
	1	4	22	0	14	15.0006	1.49945	0.5	3	
	1	5	22	0	14	14.8177	1.44216	0.740124	3.32	
	1	6	22	0	14	13.6335	2.5	0.866466	4	
	1	7	22	0	14	14.3498	1.65019	1	4	
	1	8	22	0	14	16	0.5	0.5	3	
	1	9	22	0	14	15.5	0.5	1	4	
	1	10	22	0	14	14	2.5	0.5	3	
	1	11	22	0	14	15.5	0.5	1	3	
	2	12	24.0204	0	11.9796	15.5	0.5	1	3	
	2	13	24.0204	0	11.9796	14	2.5	0.5	3.47653	
	2	14	24.0204	0	11.9796	14	2.5	0.5	4	
	2	15	24.0204	0	11.9796	13.5	2.5	1	3.52	
	2	16	24.0204	0	11.9796	16	0.5	0.5	3	
	2	17	24.0204	0	11.9796	16	0.5	0.5	3.54483	
	2	18	24.0204	0	11.9796	13.7192	2.5	0.780757	3	
	2	19	24.0204	0	11.9796	14.4284	1.57159	1	3	
	2	20	24.0204	0	11.9796	15.5	0.5	1	4	
	2	21	24.0204	0	11.9796	14.9075	1.37865	0.713848	4	
	2	22	24.0204	0	11.9796	15.5	0.5	1	3.49	
	3	23	0	23.9721	12.0279	14	2.5	0.5	3	
	3	24	0	23.9721	12.0279	16	0.5	0.5	4	

First 2 groups of the split-plot design (your runs will differ)

After the big batch of cakes are baked, the frosting and filling recipe specified by the Mix 2 components (D, E, and F) can be whipped up and applied in the proper amount, specified by Factor G.

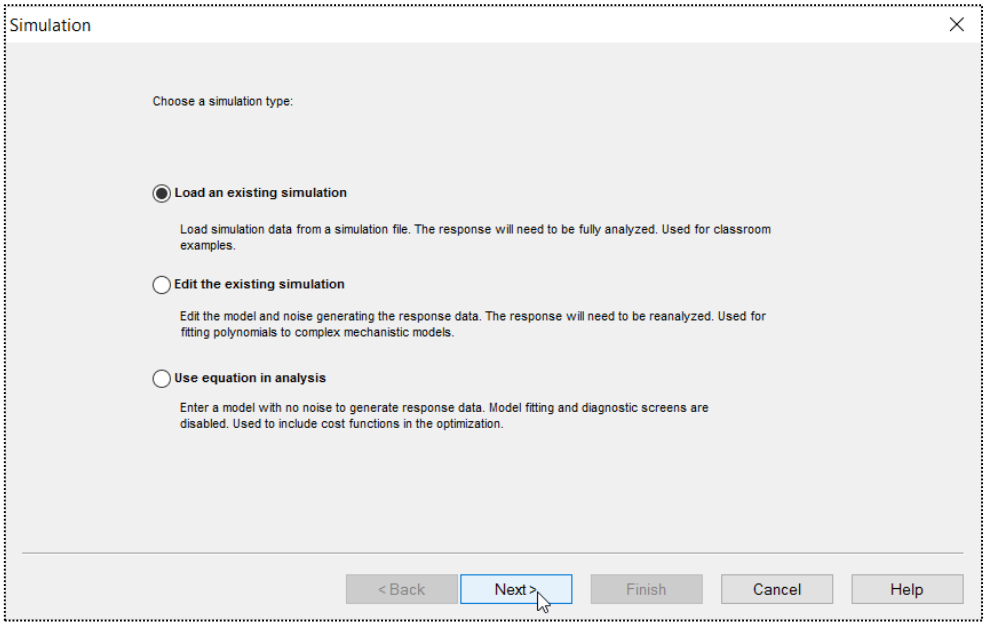
Analyze the Results

A data file with the simulated results has been created. To get the results, **right-click** on the Response 1 column header and select **Simulate…**



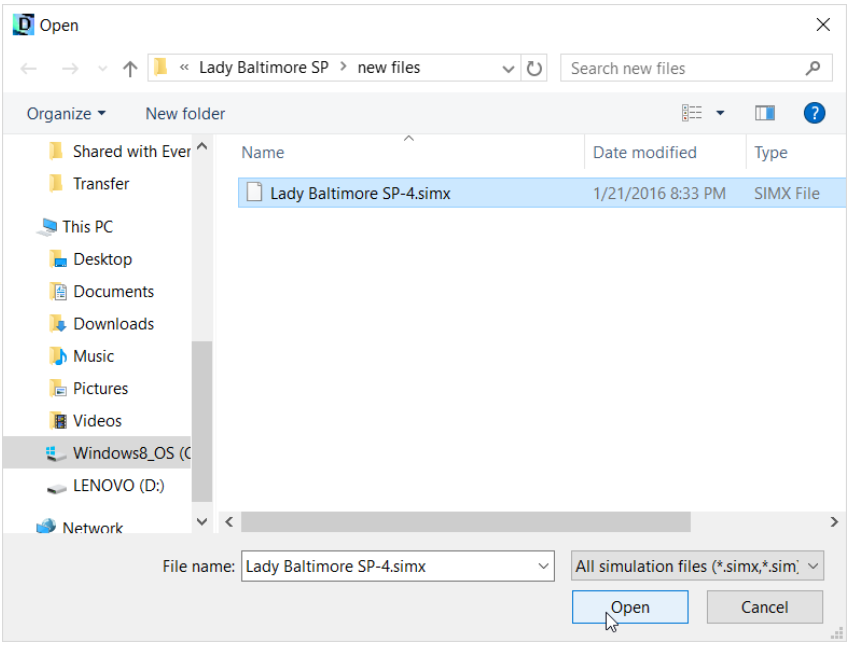
To get the data, right-click on the R1 column header and select Simulate…

Choose **Load an existing simulation** and click **Next**.



Choose Load an existing simulation

Find the **Lady Baltimore SP-4.simx** file and choose **Open** to load the data.



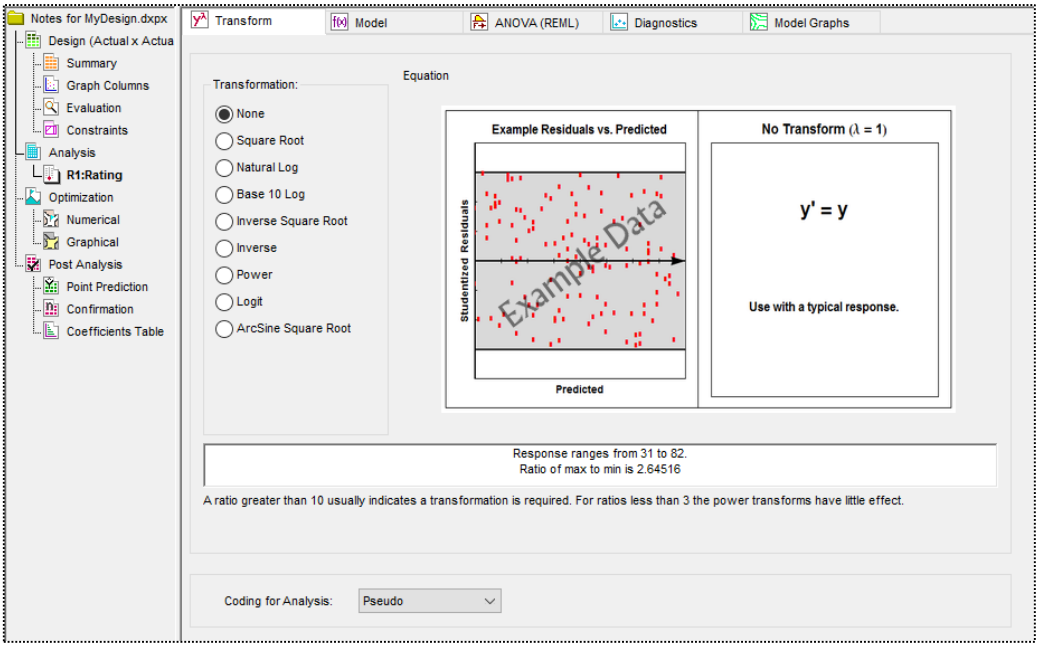
Select Lady Baltimore SP-4.simx and choose Open to get the data.

The design with data should look like the screenshot below. Note that the custom design you built is replaced with the tutorial run order for consistency.

Select	Group	Run	Mix 1-1 a:all purp flour ounces	Mix 1-2 b:cake flour ounces	Mix 1-3 c:sugar ounces	Mix 2-1 D:water teaspoons	Mix 2-2 E:corn syrup teaspoons	Mix 2-3 F:vanilla teaspoons	Factor 7 G:amount F&F cup	Response 1 Rating
	1	1	22	0	14	13.5	2.5	1	3	41
	1	2	22	0	14	16	0.5	0.5	4	37
	1	3	22	0	14	14	2.5	0.5	4	49
	1	4	22	0	14	15.0006	1.49945	0.5	3	43
	1	5	22	0	14	14.8177	1.44216	0.740124	3.32	61
	1	6	22	0	14	13.6335	2.5	0.866466	4	40
	1	7	22	0	14	14.3498	1.65019	1	4	51
	1	8	22	0	14	16	0.5	0.5	3	35
	1	9	22	0	14	15.5	0.5	1	4	46
	1	10	22	0	14	14	2.5	0.5	3	45
	1	11	22	0	14	15.5	0.5	1	3	44
	2	12	24.0204	0	11.9796	15.5	0.5	1	3	36
	2	13	24.0204	0	11.9796	14	2.5	0.5	3.47653	72
	2	14	24.0204	0	11.9796	14	2.5	0.5	4	64

Design with the Rating data (only first 14 runs shown to save space)

To get started with the analysis, click the node labeled **R1: Rating** under the Analysis branch. As with a normal RSM analysis, a new set of tabs appears at the top of your screen and they are arranged from left to right in the order needed to complete the analysis.

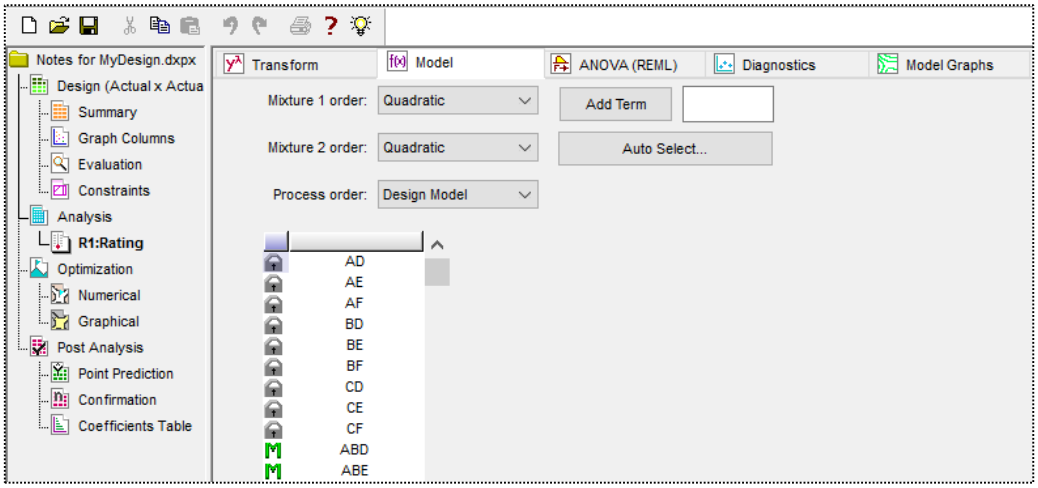


Begin analysis of taste rating

There are a variety of Transforms that can be applied on this page. Not knowing if they will help at this point, click ahead to the ***Model*** tab. There are diagnostics that are checked later that can determine if a transform will help.

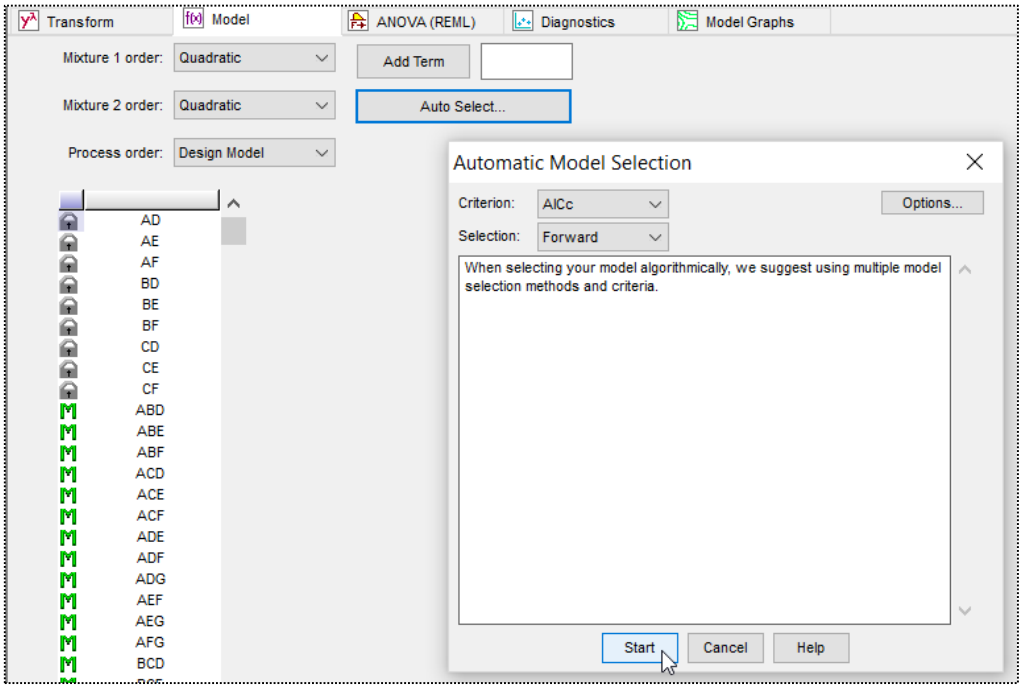
On the model tab, the combined model (reduced to only quartic terms) is presented for consideration (denoted by the green "**M**" next to the terms and the "Design model" in the process order. Clicking ahead to the ANOVA (REML) screen at this

point will evaluate that full designed for model. However, it's best to do some analysis to select the best model from among the possible terms, eliminating insignificant ones.



The designed for model is presented as a starting point

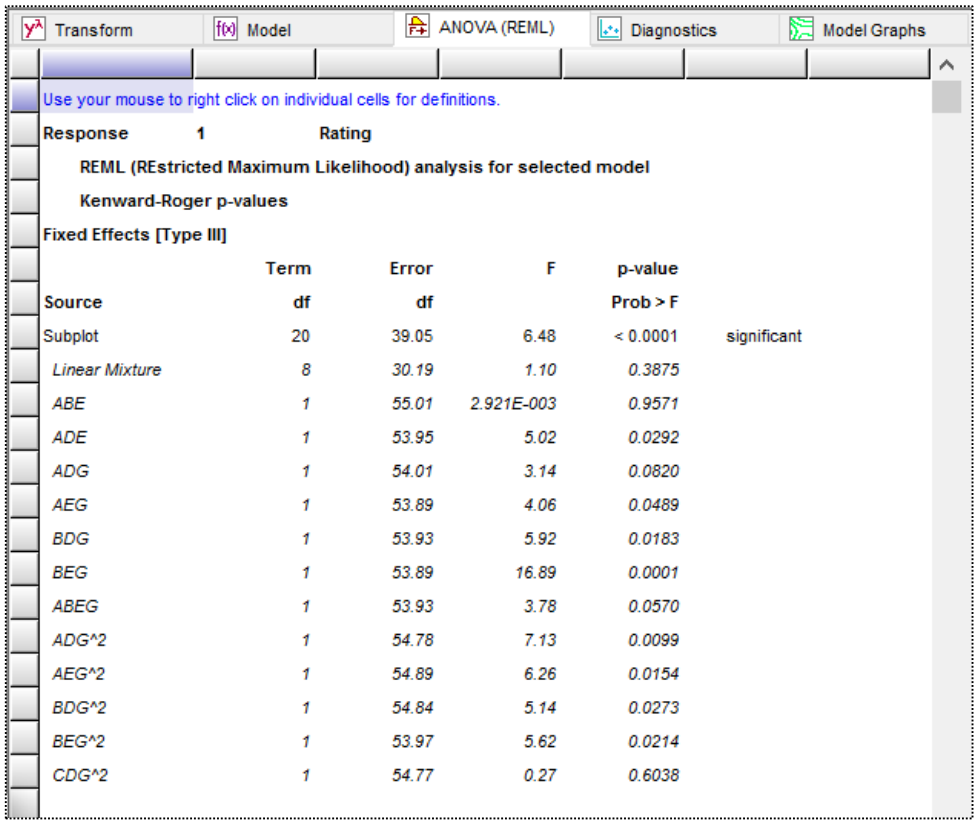
To allow the computer to do this automatically, click on the **Auto Select…** button. Accept the defaults for AICc criterion and forward selection and click the **Start** button to run the analysis. DX will go through the terms in the design model and select which ones improve the AICc criterion the most and add them to the model one at a time until adding terms will no longer improve the criterion.



Using Auto Select… to determine the best model

DX shows you the terms added in selecting the model, showing AICc criterion for each step. Click the Help button for more details on algorithmic model selection and the criterion used. Otherwise, click **Accept** to continue and evaluate the

resulting model. To get the results, click on **ANOVA (REML)** tab. You will get a warning that the model you have selected is not hierarchical. Be sure to click **Yes** to correct for hierarchy. This will give you a more statistically sound model, ensuring lower order terms are present to support higher order terms, even if they are insignificant. This is good statistical practice. Click on the help button in the warning box for more information. The model statistics are then shown.



Transform | Model | ANOVA (REML) | Diagnostics | Model Graphs

Use your mouse to right click on individual cells for definitions.

Response 1 Rating

REML (REstricted Maximum Likelihood) analysis for selected model

Kenward-Roger p-values

Fixed Effects [Type III]

Source	Term df	Error df	F	p-value Prob > F	
Subplot	20	39.05	6.48	< 0.0001	significant
Linear Mixture	*8*	*30.19*	*1.10*	*0.3875*	
ABE	*1*	*55.01*	*2.921E-003*	*0.9571*	
ADE	*1*	*53.95*	*5.02*	*0.0292*	
ADG	*1*	*54.01*	*3.14*	*0.0820*	
AEG	*1*	*53.89*	*4.06*	*0.0489*	
BDG	*1*	*53.93*	*5.92*	*0.0183*	
BEG	*1*	*53.89*	*16.89*	*0.0001*	
ABEG	*1*	*53.93*	*3.78*	*0.0570*	
ADG^2	*1*	*54.78*	*7.13*	*0.0099*	
AEG^2	*1*	*54.89*	*6.26*	*0.0154*	
BDG^2	*1*	*54.84*	*5.14*	*0.0273*	
BEG^2	*1*	*53.97*	*5.62*	*0.0214*	
CDG^2	*1*	*54.77*	*0.27*	*0.6038*	

Model statistics from the REML analysis

This is not your standard ANOVA analysis, which relies on randomization for validity. The analysis done for split-plot designs in DX is a form of maximum likelihood estimation, more specifically, restricted maximum likelihood (REML), as noted at the top of the results table.

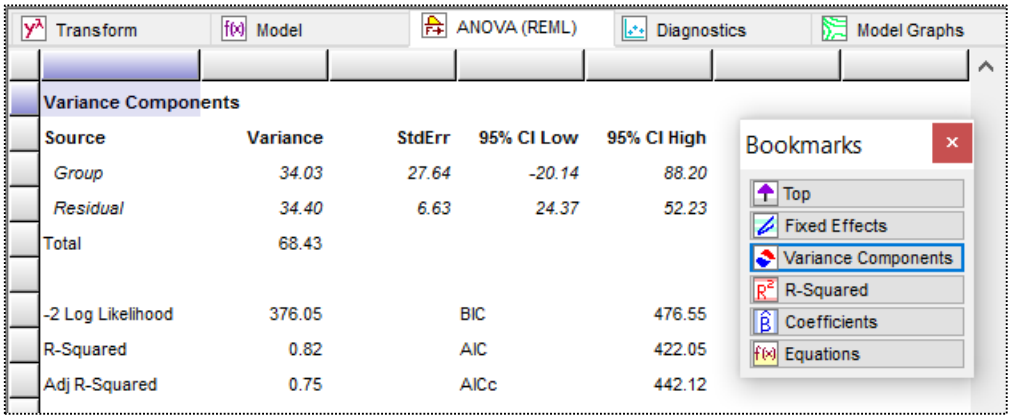
☒ ***Details on split-plot analysis***: The aim of maximum likelihood estimation is to find the parameter value(s) that makes the observed data most likely. Restricted maximum likelihood estimation, which is generally used unless you click on the Analysis menu available on the Model screen to change the method, is another way to estimate variances. In the split plot case, REML estimates the Group variance for the whole plot factors and the residual variance for the subplot factors. Once the variances are estimated, Generalized Least Squares (GLS) is used to estimate the factor effects. The Kenward-Roger's method is then used to produce F-tests and the corresponding p-values. You can learn even more by clicking on the lightbulb icon for screen tips and following the links.

The big difference between the statistics on this table and a normal ANOVA is the grouping of variance terms into a Whole-plot section for the HTC factors and a subplot section for ETC factors. However, for this design, there are no Whole-Plot terms selected. In other words, there are no terms consisting of just A, B, and C that are significant. There are terms involving A, B, and C, but they are always

interacting with the frosting mixture and frosting amount (G) terms and are thus part of the subplot. That's somewhat expected, because in a split-plot design, the subplot terms (and their interactions) have more power and can be detected more easily. In fact, these subplot interactions can often make up for a lack of power for the HTC factors.

Looking at the subplot as a whole, the model has a quite significant F value (p-value of <0.0001). Most of the terms are also significant (at the 0.05 alpha level) or needed for hierarchy. For example, the insignificant term ABE is needed for the significant ABEG term.

Next, press the **Variance Components** on the floating Bookmarks tool and DX presents various statistics to augment the REML analysis.

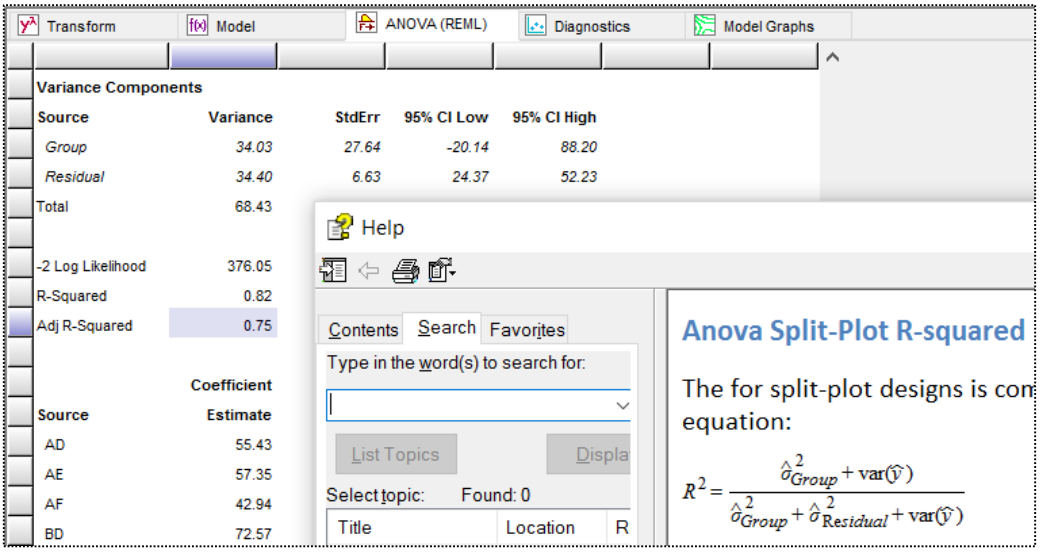


Variance Components and model criterion

Here, you will see more details on the variance components and likelihood ratios for the selected model, including the information criterion (AIC, BIC, and AICc). More can be learned on those in the help menus.

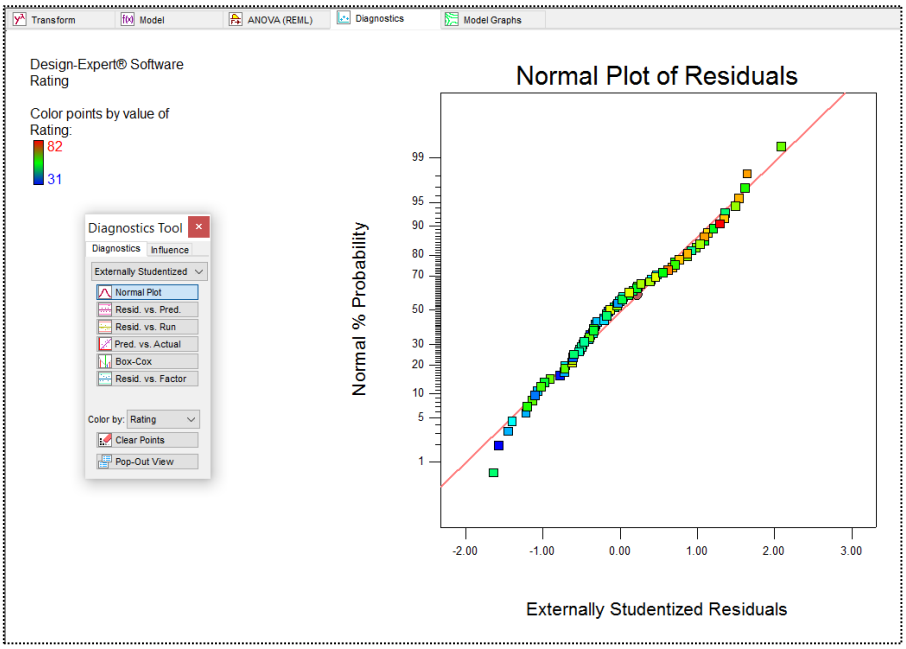
One important number to look at is the Adj. R-squared (adjusted R-squared). This number goes from 0 to 1, with 1 being the best. In this case, the 0.75 Adj. R-Squared shows that the selected model captures most of the variation in the data (~75%).

To learn more about the various model criterion, this would be a great place to exercise DX's context sensitive help. Just click on a number you are interested in to highlight it and then press the **F1** key (or right-click and select Help). For example, look at the information obtained about the Adj R-Squared criterion.



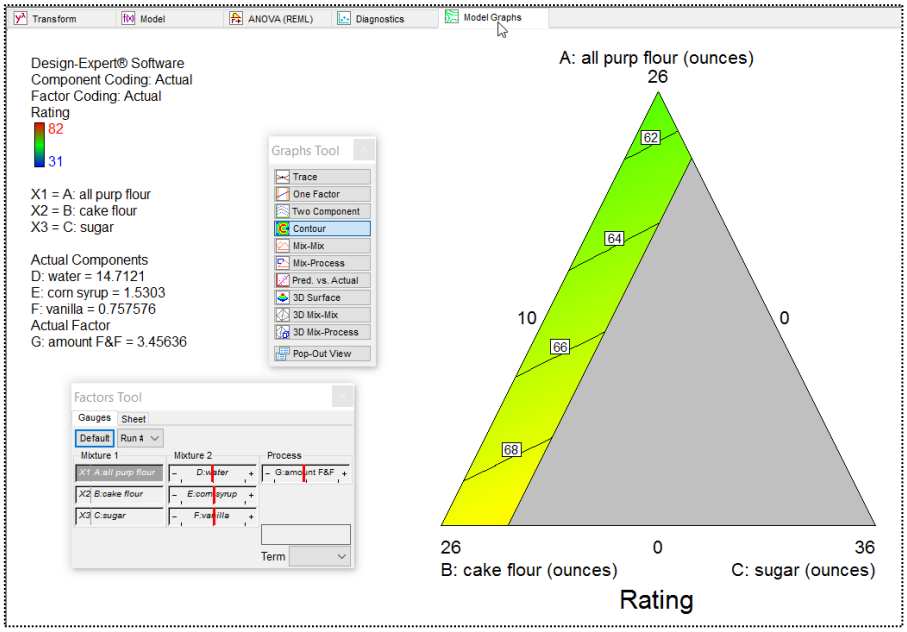
Using the context sensitive Help (F1) to see more details

That's enough on the model statistics. It seems we have quite a strong model. Click the **Diagnostics** tab and examine the graphs of residuals.



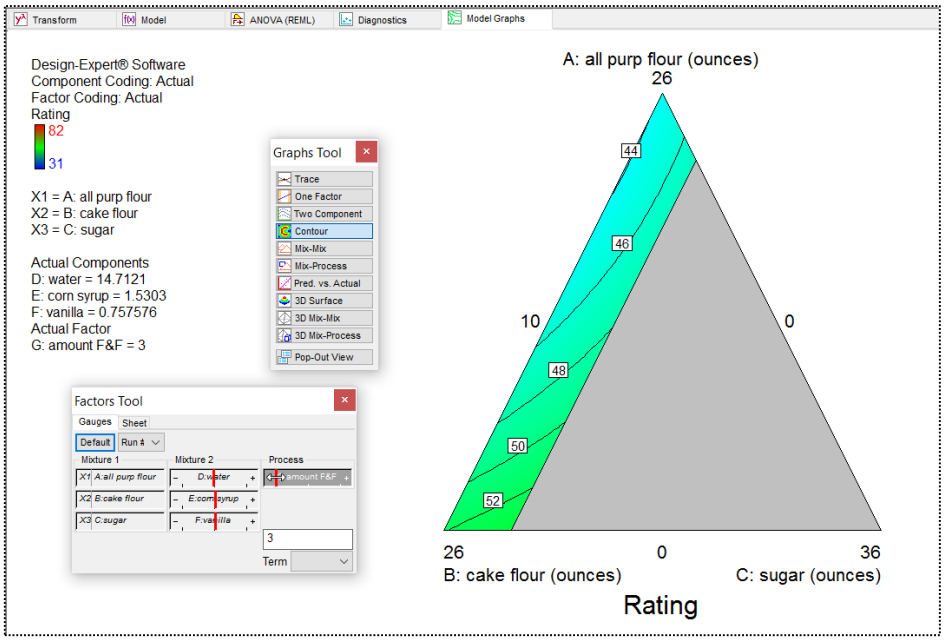
Normal plot of residuals – looks good

The residuals graphs found via the floating diagnostics tool are important to check, but these have been covered extensively in other tutorials. For instance, see the *Multi-Factor RSM* tutorial available on our tutorials page: www.statease.com/software/dx10-tut.html. In this case, the diagnostics all look good, so press on to the **Model Graphs** to view the response in mixture (triangular) space.



Model Graph with view of contour plot for mixture portion of combined design

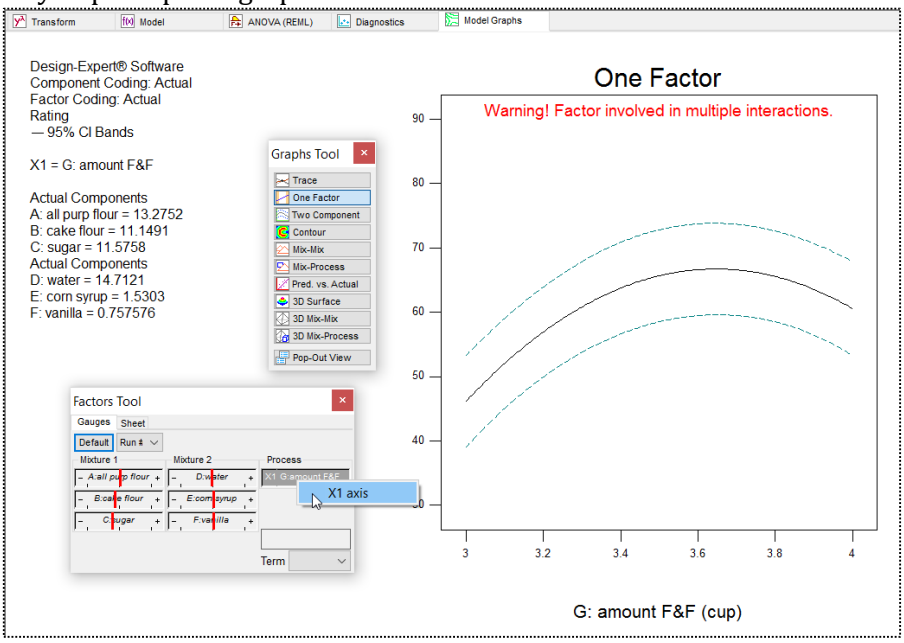
Remember that this mixture interacts with both Mixture 2 (Frosting) and the Process factor (F&F amount), represented on the floating Factors Tool. Click and drag the red bars for those factors and the response graph will change. For example, **drag the amount of F&F** from the Process portion of the factors tool to the **left, low level**. Reducing the amount of frosting seems to reduce the taste rating across the board, which makes sense. Who wants cake without much frosting?



Dragging the amount of F&F (frosting) to the left (low) produces lower taste ratings.

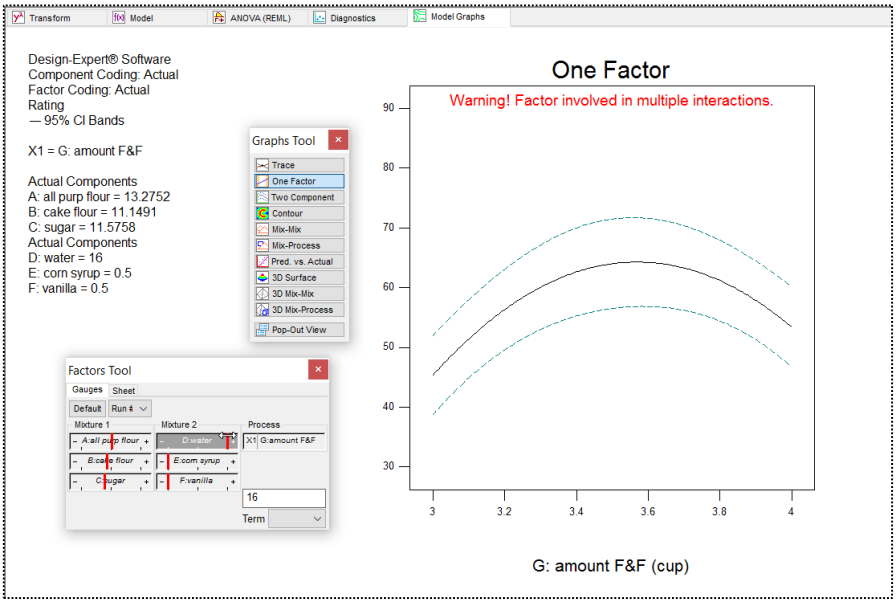
Let's see what this process factor looks like on its own. **Right-click** on the **amount F&F** process factor in the floating Factors Tool and select **X1 axis**. Note that "One

Factor" is now highlighted in the graphs tool. Clicking on that button is the other way to pull up this graph.



Raising the amount of F&F (frosting) improves the taste to a point, but too much is undesirable.

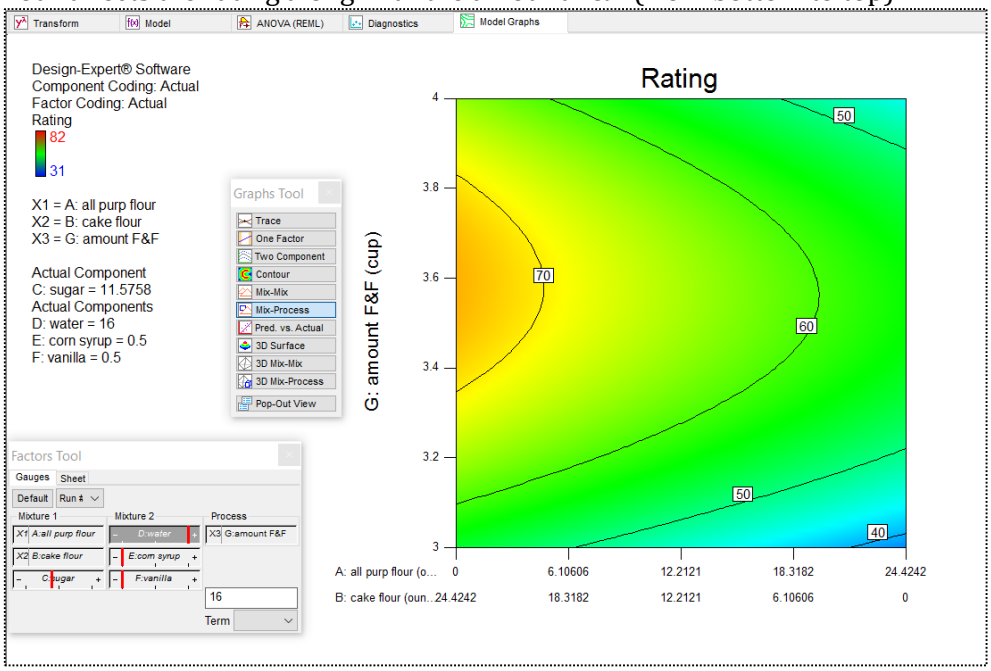
This will put the amount F&F on the x-axis. As seen in the prior graph, raising the amount of F&F improves the taste, but only to a point. If the amount is raised to the highest level, the taste rating goes down. This is only part of the story, as you may note from the red Warning atop the graph! The amount of F&F interacts with the two mixtures. To see how changing the frosting mixture affects the graph, drag the red bars for Mixture 2 from the floating Factors tool. For example, **drag water to its high level** and see how the amount F&F graph responds.



Raising the proportion of water in Mix 2 affects the process plot

Note that as more water is added, the corn syrup and vanilla must be removed to maintain the constant total for the frosting mixture. Adding all that water doesn't change the overall rating much and there is still a peak (slightly sharper) for the optimum amount of F&F.

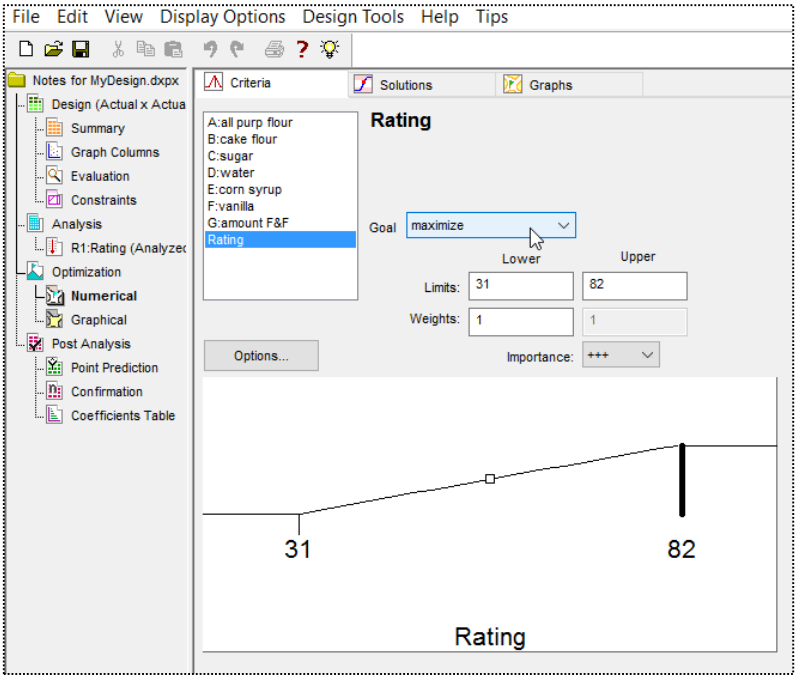
In a combined design like this, there are many other interesting graphs to investigate. For instance, click on the **Mix-Process** button on the floating graphs tool. This reveals see how substituting all-purpose flour (from left to right) for cake flour affects the rating along with the amount F&F (from bottom to top).



Mix-process plot showing flour on the x-axis and amount F&F on the y-axis

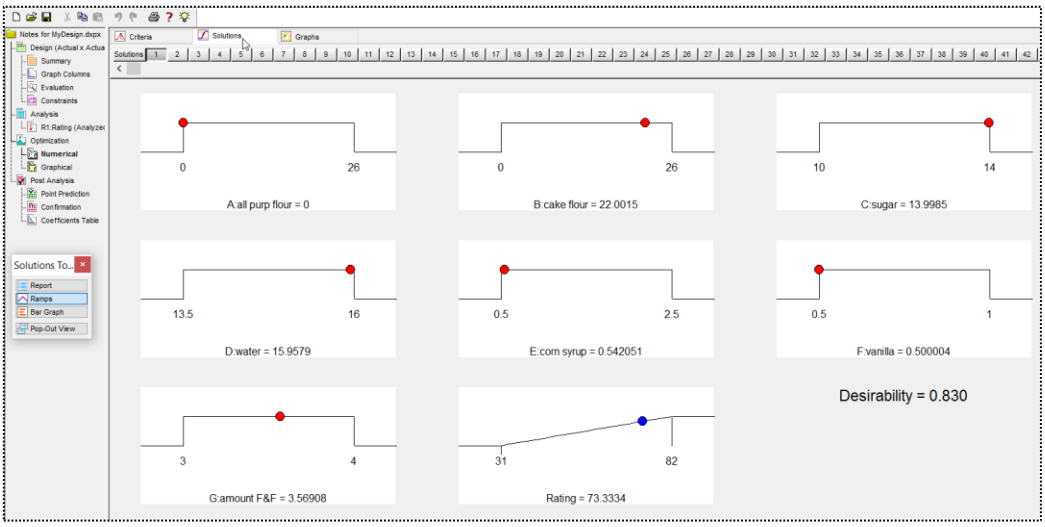
Find the Optimal Solution

The goal of the experimental program is to learn how to customize the Lady Baltimore cake recipe to get the highest overall rating. To find optimal combinations of formulas and processing, click the optimization node labeled **Numerical**. Then select **Rating**. Choose **maximize** from the Goal droplist and leave everything else at the default settings. Your screen should now look like that below.



Numerical Optimization Settings

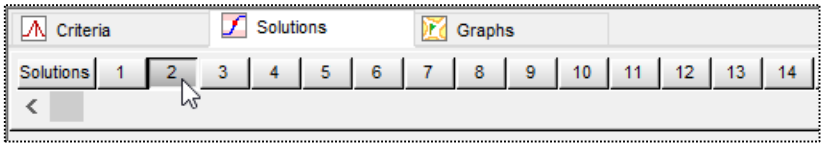
Click the **Solutions** tab. The solutions are presented in the ramps view, by default.



Numerical Optimization Settings (yours may vary slightly)

The ramps view makes it easy to see the levels of each component/factor and the resulting rating (73). Unfortunately, all-purpose flour is set to the low level. I guess the recipe makers knew what they were doing when they called for cake flour. The amount of F&F (G), is set to the upper middle level seen previously when investigating that factor on the one factor plot.

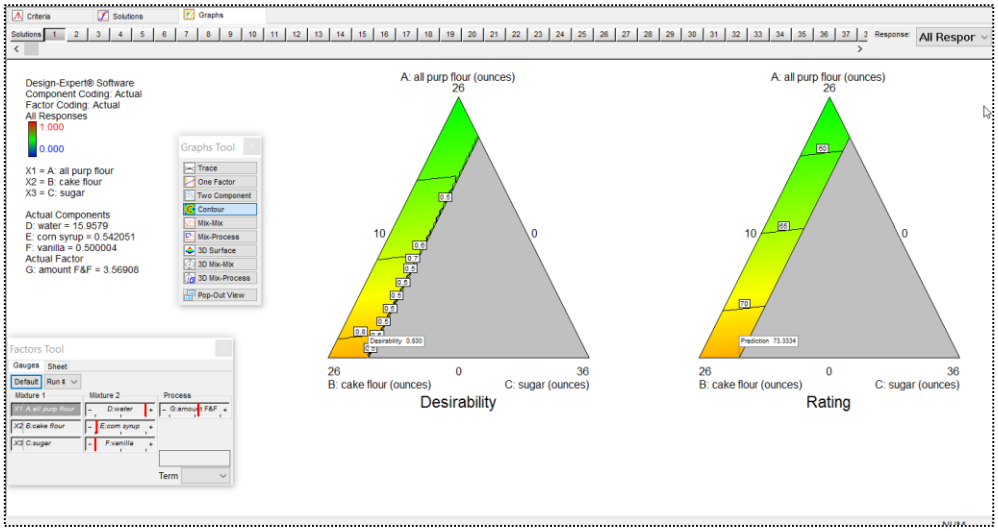
That's solution number one, with the highest rating. To investigate a few other options, **click** through the other solutions in the top **solutions bar**.



Click other solutions to see how the settings change here or on the Graphs tab

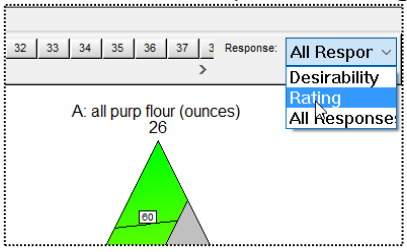
Once you click in the solutions bar, you can easily toggle through the other solutions using the left and right arrow keys. Even going to solutions with lower ratings, it seems the all-purpose flour must be set to a low level to get good ratings.

Click on the **Graphs** tab to investigate the graphs at the optimal solution (be sure to select solution number 1 on the solutions bar). There will be a flag planted at the optimum (at low levels of all-purpose flour). By default, all responses are shown side by side, including the desirability plot used to search for the optimum.



Click the graphs tab to investigate the optimal solutions (flags) on various graphs.

To concentrate on just the rating, select **Rating** from the Response droplist.



To see just the Rating graph, use the Response droplist.

Having explored many graphs already, the recipes have been pretty well optimized. We'll leave you on your own if you'd like to investigate the graphs some more.

Final Comments

A split-plot design can be applied to save experimental effort. This can even be applied in the case of a complex combined design like this, involving two mixtures and a process factor. Here, it allowed the bakers to make 11 cakes at a time in big batches, instead of changing the batter every run and baking one cake at a time. That's a big savings of time and budget, allowing the cake recipe to be fully optimized. With split-plot designs, you can have your cake and eat it, too!

