


# Design and Analysis of Experiments by Douglas Montgomery: *A Supplement for Using JMP®*



**Heath Rushing, Andrew Karl, and James Wisnowski**

# Contents

<b>About This Book .....</b>	<b>ix</b>
<b>About These Authors .....</b>	<b>xiii</b>
<b>Acknowledgments .....</b>	<b>xv</b>
<b>Chapter 1 Introduction .....</b>	<b>1</b>
<b>Chapter 2 Simple Comparative Experiments .....</b>	<b>5</b>
Section 2.2 Basic Statistical Concepts.....	6
Section 2.4.1 Hypothesis Testing.....	10
Section 2.4.3 Choice of Sample Size .....	12
Section 2.5.1 The Paired Comparison Problem.....	17
Section 2.5.2 Advantages of the Paired Comparison Design .....	18
<b>Chapter 3 Experiments with a Single Factor: The Analysis of Variance ....</b>	<b>21</b>
Section 3.1 A One-way ANOVA Example .....	22
Section 3.4 Model Adequacy Checking.....	32
Section 3.8.1 Single Factor Experiment .....	46
Section 3.8.2 Application of a Designed Experiment.....	52
Section 3.8.3 Discovering Dispersion Effects.....	54
<b>Chapter 4 Randomized Blocks, Latin Squares, and Related Designs .....</b>	<b>61</b>
Section 4.2 Creating a Latin Square Design in JMP .....	67
<b>Chapter 5 Introduction to Factorial Designs .....</b>	<b>77</b>
Example 5.1 The Battery Design Experiment .....	78
Example 5.2 A Two-Factor Experiment with a Single Replicate.....	82
Example 5.3 The Soft Drink Bottling Problem .....	84
Example 5.4 The Battery Design Experiment with a Covariate .....	86
Example 5.5 A 3 <sup>2</sup> Factorial Experiment with Two Replicates.....	89
Example 5.6 A Factorial Design with Blocking .....	97

<b>Chapter 6 The 2<sup>k</sup> Factorial Design .....</b>	<b>101</b>
Section 6.2 The 2 <sup>2</sup> design .....	102
Example 6.1 A 2 <sup>3</sup> Design .....	107
Example 6.2 A Single Replicate of the 2 <sup>4</sup> Design .....	109
Example 6.3 Data Transformation in a Factorial Design.....	114
Example 6.5 Duplicate Measurements on the Response .....	118
Example 6.6 Credit Card Marketing .....	125
Example 6.7 A 2 <sup>4</sup> Design with Center Points .....	128
<b>Chapter 7 Blocking and Confounding in the 2<sup>k</sup> Factorial Design.....</b>	<b>131</b>
Example 7.1 A 2 <sup>k</sup> Replicated Factorial Design with Blocking.....	132
Example 7.2 Blocking and Confounding in an Unreplicated Design .....	132
Example 7.3 A 2 <sup>3</sup> Design with Partial Confounding .....	134
<b>Chapter 8 Two-Level Fractional Factorial Designs .....</b>	<b>141</b>
Example 8.1 A Half-Fraction of the 2 <sup>4</sup> Design .....	143
Example 8.2 A 2 <sup>5-1</sup> Design Used for Process Improvement .....	147
Example 8.3 A 2 <sup>4-1</sup> Design with the Alternate Fraction .....	152
Example 8.4 A 2 <sup>6-2</sup> Design.....	153
Example 8.5 A 2 <sup>7-3</sup> Design.....	158
Example 8.6 A 2 <sup>8-3</sup> Design in Four Blocks .....	160
Example 8.7 A Fold-Over 2 <sup>7-4</sup> Resolution III Design .....	164
Example 8.8 The Plackett-Burman Design .....	167
Section 8.7.2 Sequential Experimentation with Resolution IV Designs.....	168
<b>Chapter 9 Three-Level and Mixed-Level Factorial and Fractional Factorial     Designs .....</b>	<b>173</b>
Example 9.1 The 3 <sup>3</sup> Design .....	174
Example 9.2 The 3 <sup>2</sup> Design Confounded in 3 Blocks .....	177
Example 9.3 The Spin Coating Experiment .....	178
Example 9.4 An Experiment with Unusual Blocking Requirements.....	181
<b>Chapter 10 Fitting Regression Models.....</b>	<b>189</b>
Example 10.1 Multiple Linear Regression Model.....	190
Example 10.2 Regression Analysis of a 2 <sup>3</sup> Factorial Design.....	195
Example 10.3 A 2 <sup>3</sup> Factorial Design with a Missing Observation .....	197
Example 10.4 Inaccurate Levels in Design Factors .....	198

Example 10.6 Tests on Individual Regression Coefficients .....	198
Example 10.7 Confidence Intervals on Individual Regression Coefficients .....	199
<b>Chapter 11 Response Surface Methods and Designs .....</b>	<b>201</b>
Example 11.1 The Path of Steepest Ascent .....	202
Example 11.2 Central Composite Design.....	204
Section 11.3.4 Multiple Responses .....	209
Example 11.4 Space Filling Design with Gaussian Process Model .....	214
Example 11.5 A Three-Component Mixture .....	218
Example 11.6 Paint Formulation .....	222
<b>Chapter 12 Robust Parameter Design and Process Robustness Studies .....</b>	<b>227</b>
Example 12.1 Two Controllable Variables and One Noise Variable .....	228
Example 12.2 Two Controllable Variables and Three Noise Variables .....	230
<b>Chapter 13 Experiments with Random Factors.....</b>	<b>239</b>
Example 13.1 A Measurement Systems Capability Study .....	240
Example 13.3 The Unrestricted Model .....	242
Example 13.5 A Three-Factor Factorial Experiment with Random Factors .....	244
Example 13.6 Approximate F Tests .....	245
<b>Chapter 14 Nested and Split-Plot Designs.....</b>	<b>251</b>
Example 14.1 The Two-Stage Nested Design.....	252
Example 14.2 A Nested-Factorial Design.....	254
Section 14.4 The Experiment on the Tensile Strength of Paper .....	256
Example 14.3 A $2^{5-1}$ Split-Plot Experiment.....	259
<b>Chapter 15 Other Design and Analysis Topics.....</b>	<b>263</b>
Example 15.1 Box-Cox Transformation.....	264
Example 15.2 The Generalized Linear Model and Logistic Regression.....	265
Example 15.3 Poisson Regression.....	267
Example 15.4 The Worst Yarn Experiment .....	269
Section 15.2 Unbalanced Data in a Factorial Design.....	270
Example 15.5 Analysis of Covariance.....	271
Section 15.3.4 Factorial Experiments with Covariates .....	273
<b>Index .....</b>	<b>277</b>

**viii** Contents

# 1 Introduction

The analysis of a complex process requires the identification of target quality attributes that characterize the output of the process and of factors that may be related to those attributes. Once a list of potential factors is identified from subject-matter expertise, the strengths of the associations between those factors and the target attributes need to be quantified. A naïve, one-factor-at-a-time analysis would require many more trials than necessary. Additionally, it would not yield information about whether the relationship between a factor and the target depends on the values of other factors (commonly referred to as interaction effects between factors). As demonstrated in Douglas Montgomery's *Design and Analysis of Experiments* textbook, principles of statistical theory, linear algebra, and analysis guide the development of efficient experimental designs for factor settings. Once a subset of important factors has been isolated, subsequent experimentation can determine the settings of those factors that will optimize the target quality attributes. Fortunately, modern software has taken advantage of the advanced theory. This software now facilitates the development of good design and makes solid analysis more accessible to those with a minimal statistical background.

Designing experiments with specialized design of experiments (DOE) software is more efficient, complete, insightful, and less error-prone than producing the same design by hand with tables. In addition, it provides the ability to generate algorithmic designs (according to one of several possible optimality criteria) that are frequently required to accommodate constraints commonly encountered in practice. Once an experiment has been designed and executed, the analysis of the results should respect the assumptions made during the design process. For example, split-plot experiments with hard-to-change factors should be analyzed as such; the constraints of a mixture design must be incorporated; non-normal responses should either be transformed or modeled with a generalized linear model; correlation between repeated observations on an experimental unit may be modeled with random effects; non-constant variance in the response variable

## 2 *Design and Analysis of Experiments by Douglas Montgomery: A Supplement for Using JMP*

across the design factors may be modeled, etc. Software for analyzing designed experiments should provide all of these capabilities in an accessible interface.

JMP offers an outstanding software solution for both designing and analyzing experiments. In terms of design, all of the classic designs that are presented in the textbook may be created in JMP. Optimal designs are available from the JMP Custom Design platform. These designs are extremely useful for cases where a constrained design space or a restriction on the number of experimental runs eliminates classical designs from consideration. Multiple designs may be created and compared with methods described in the textbook, including the Fraction of Design Space plot. Once a design is chosen, JMP will randomize the run order and produce a data table, which the researcher may use to store results. Metadata for the experimental factors and response variables is attached to the data table, which simplifies the analysis of these results.

The impressive graphical analysis functionality of JMP accelerates the discovery process particularly well with the dynamic and interactive profilers and plots. If labels for plotted points overlap, can be clicked and dragged. Selecting points in a plot produced from a table selects the appropriate rows in the table and highlights the points corresponding to those rows in all other graphs produced from the table. Plots can be shifted and rescaled by clicking and dragging the axes. In many other software packages, these changes are either unavailable or require regenerating the graphical output.

An additional benefit of JMP is the ease with which it permits users to manipulate data tables. Data table operations such as sub-setting, joining, and concatenating are available via intuitive graphical interfaces. The relatively short learning curve for data table manipulation enables new users to prepare their data without remembering an extensive syntax. Although no command-line knowledge is necessary, the underlying JMP scripting language (JSL) scripts for data manipulation (and any other JMP procedure) may be saved and edited to repeat the analysis in the future or to combine with other scripts to automate a process.

This supplement to *Design and Analysis of Experiments* follows the chapter topics of the textbook and provides complete instructions and useful screenshots to use JMP to solve every example problem. As might be expected, there are often multiple ways to perform the same operation within JMP. In many of these cases, the different possibilities are illustrated across different examples involving the relevant operation. Some theoretical results are discussed in this supplement, but the emphasis is on the practical application of the methods developed in the textbook. The JMP DOE functionality detailed here represents a fraction of the software's features for not only DOE, but also for most other

areas of applied statistics. The platforms for reliability and survival, quality and process control, time series, multivariate methods, and nonlinear analysis procedures are beyond the scope of this supplement.



4 *Design and Analysis of Experiments by Douglas Montgomery: A Supplement for Using JMP*

# Simple Comparative Experiments

# 2

Section 2.2 Basic Statistical Concepts.....	6
Section 2.4.1 Hypothesis Testing.....	10
Section 2.4.3 Choice of Sample Size .....	12
Section 2.5.1 The Paired Comparison Problem .....	17
Section 2.5.2 Advantages of the Paired Comparison Design .....	18

The problem of testing the effect of a single experimental factor with only two levels provides a useful introduction to the statistical techniques that will later be generalized for the analysis of more complex experimental designs. In this chapter, we develop techniques that will allow us to determine the level of statistical significance associated with the difference in the mean responses of two treatment levels. Rather than only considering the difference between the mean responses across the treatments, we also consider the variation in the responses and the number of runs performed in the experiment. Using a t-test, we are able to quantify the likelihood (expressed as a p-value) that the observed treatment effect is merely noise. A “small” p-value (typically taken to be one smaller than  $\alpha = 0.05$ ) suggests that the observed data are not likely to have occurred if the null hypothesis (of no treatment effect) were true.

A related question involves the likelihood that the null hypothesis is rejected given that it is false (the power of the test). Given a fixed significance level,  $\alpha$  (our definition of what constitutes a “small” p-value), theorized values for the pooled standard deviation, and a minimum threshold difference in treatment means, it is possible to solve for the minimum sample size that is necessary to achieve a desired power. This procedure is useful for determining the number of runs that must be included in a designed experiment.

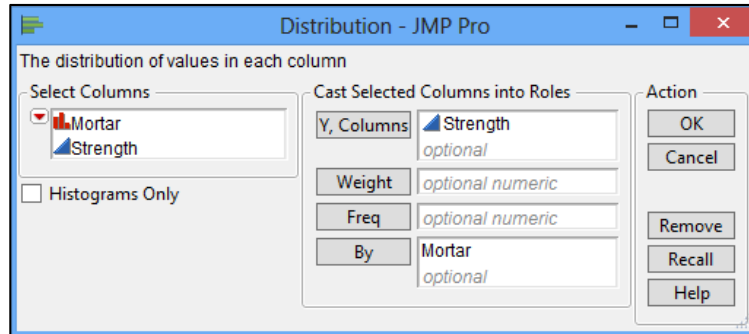
In the first example presented in this chapter, a scientist has developed a modified cement mortar formulation that has a shorter cure time than the unmodified formulation. The scientist would like to test if the modification has affected the bond strength of the mortar. To study whether the two formulations, on average, produce bonds of different strengths, a two-sided t-test is used to analyze the observations from a randomized experiment with 10 measurements from each formulation. The null hypothesis of this test is that the mean bond strengths produced by the two formulations are equal; the alternative hypothesis is that mean bond strengths are not equal.

We also consider the advantages of a paired t-test, which provides an introduction to the notion of blocking. This test is demonstrated using data from an experiment to test for similar performance of two different tips that are placed on a rod in a machine and pressed into metal test coupons. A fixed pressure is applied to the tip, and the depth of the resulting depression is measured. A completely randomized design would apply the tips in a random order to the test coupons (making only one measurement on each coupon). While this design would produce valid results, the power of the test could be increased by removing noise from the coupon-to-coupon variation. This may be achieved by applying both tips to each coupon (in a random order) and measuring the difference in the depth of the depressions. A one-sample t-test is then used for the null hypothesis that the mean difference across the coupons is equal to 0. This procedure reduces experimental error by eliminating a noise factor.

This chapter also includes an example of procedures for testing the equality of treatment variances, and a demonstration of the t-test in the presence of potentially unequal group variances. This final test is still valid when the group variances are equal, but it is not as powerful as the pooled t-test in such situations.

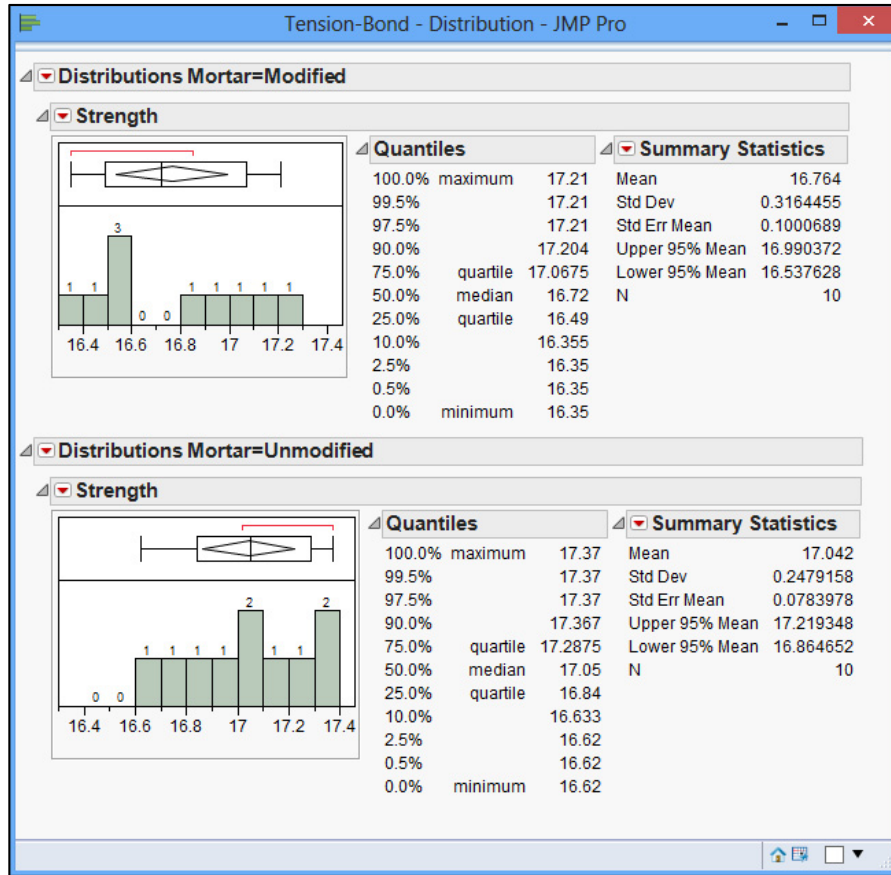
## Section 2.2 Basic Statistical Concepts

1. Open Tension-Bond.jmp.
2. Select **Analyze > Distribution**.
3. Select *Strength* for **Y, Columns**.
4. Select *Mortar* for **By**. As we will see in later chapters, these fields will be automatically populated for data tables that were created in JMP.



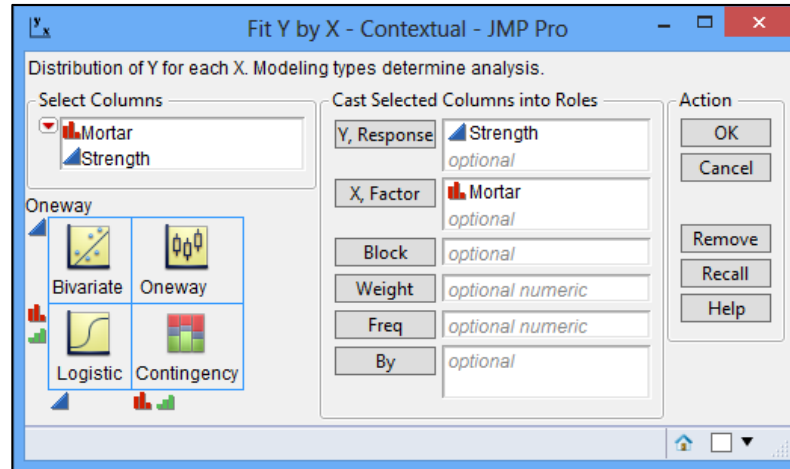
5. Click **OK**.
6. Click the red triangle next to Distributions Mortar=Modified and select **Uniform Scaling**.
7. Repeat step 6 for Distributions Mortar=Unmodified.
8. Click the red triangle next to Distributions Mortar=Modified and select **Stack**.
9. Repeat step 8 for Distributions Mortar=Unmodified.
10. Hold down the *Ctrl* key and click the red triangle next to Strength. Select **Histogram Options > Show Counts**. Holding down *Ctrl* applies the command to all of the histograms created by the Distribution platform; it essentially “broadcasts” the command.




8 Design and Analysis of Experiments by Douglas Montgomery: A Supplement for Using JMP



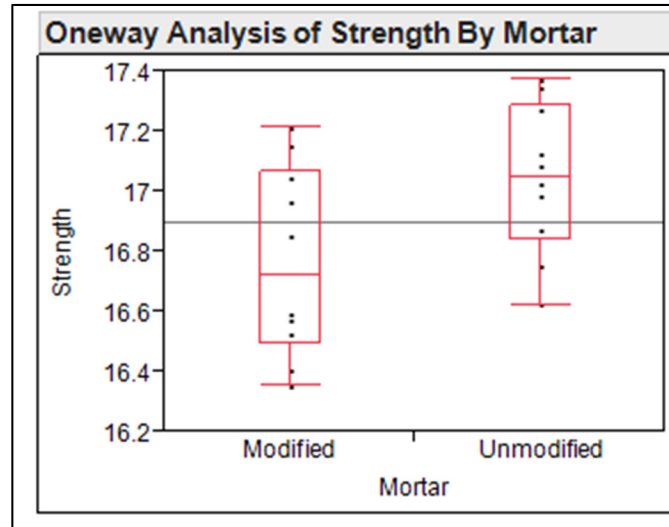
It appears from the overlapped histograms that the unmodified mortar tends to produce stronger bonds than the modified mortar. The unmodified mortar has a mean strength of 17.04 kgf/cm<sup>2</sup> with a standard deviation of 0.25 kgf/cm<sup>2</sup>. The modified mortar has a mean strength of 16.76 kgf/cm<sup>2</sup> with a standard deviation of 0.32 kgf/cm<sup>2</sup>. A naïve comparison of mean strength indicates that the unmodified mortar outperforms the modified mortar. However, the difference in means could simply be a result of sampling fluctuation. Using statistical theory, our goal is to incorporate the sample standard deviations (and sample sizes) to quantify how likely it is that the difference in mean strengths is due only to sampling error. If it turns out to be unlikely, we will conclude that a true difference exists between the mortar strengths.

11. Select **Analyze > Fit Y by X**.
12. Select *Strength* for **Y, Response** and *Mortar* for **X, Grouping**.



The Fit Y by X platform recognizes this as a one-way ANOVA since the response, *Strength*, is a continuous factor, and the factor *Mortar* is a nominal factor. When JMP is used to create experimental designs, it assigns the appropriate variable type to each column. For imported data, JMP assigns a modeling type—continuous , ordinal , or nominal —to each variable based on attributes of that variable. A different modeling type may be specified by right-clicking the modeling type icon next to a column name and selecting the new type.

13. Click **OK**.
14. To create box plots, click the red triangle next to One-way Analysis of Strength by Mortar and select **Quantiles**.




The median modified mortar strength (represented by the line in the middle of the box) is lower than the median unmodified mortar strength. The similar length of the two boxes (representing the interquartile ranges) indicates that the two mortar formulations result in approximately the same variability in strength.

15. Keep the Fit Y by X platform open for the next exercise.

### Section 2.4.1 Hypothesis Testing

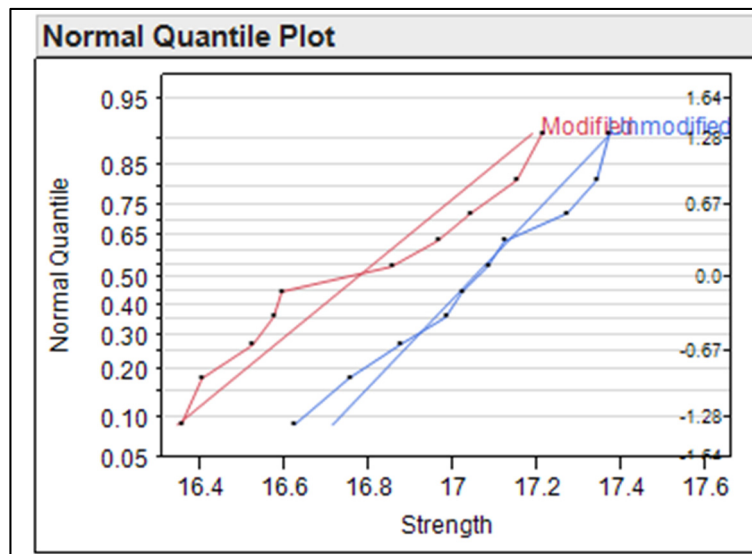
1. Return to the Fit Y by X platform from the previous exercise.
2. Click the red triangle next to One-way Analysis of Strength by Mortar and select **Means/Anova/Pooled t**.



The **t-test** report shows the two-sample t-test assuming equal variances. Since we have a two-sided alternative hypothesis, we are concerned with the p-value labeled  $\text{Prob} > |t| = 0.0422$ . Since we have set  $\alpha = 0.05$ , we reject the null hypothesis that the mean strengths produced by the two formulations of mortar are equal and conclude that the mean strength of the modified mortar and the mean strength of the unmodified mortar are (statistically) significantly different. In practice, our next step would be to decide from a subject-matter perspective if the difference is practically significant.

Before accepting the conclusion of the t test, we should use diagnostics to check the validity of assumptions made by the model. Although this step is not shown for every example in the text, it is an essential part of every analysis. For example, a quantile plot may be used to check the assumptions of normality and identical population variances. Though not shown here, a plot of the residuals against run order could help identify potential violations of the assumed independence across runs (the most important of the three assumptions).

- Click the red triangle next to One-way Analysis of Strength by Mortar and select **Normal Quantile Plot > Plot Quantile by Actual**.



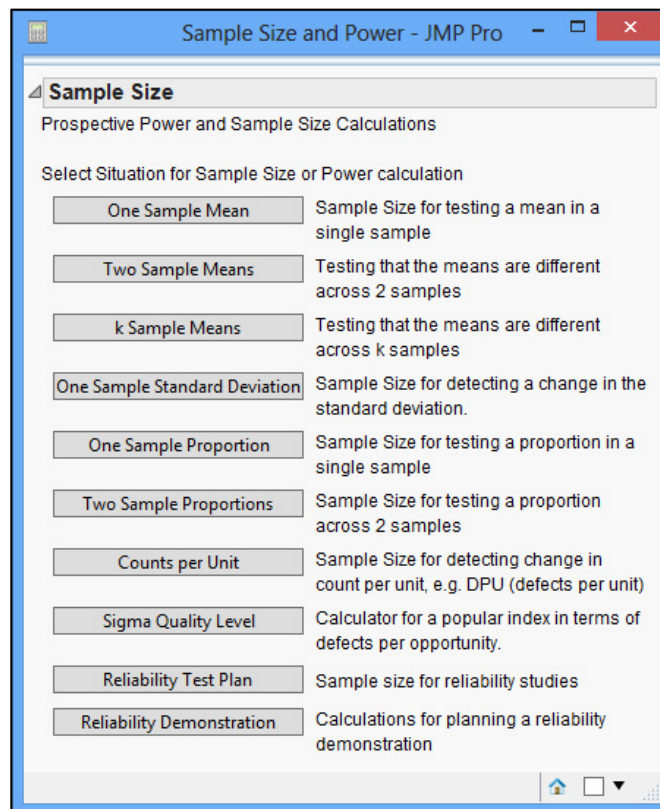
The points fall reasonably close to straight lines in the plot, suggesting that the assumption of normality is reasonable. The slopes of the lines are proportional to the standard deviations in each comparison group. These slopes appear to be similar, supporting the decision to assume equal population variances.



4. Select **Window > Close All**.

### Section 2.4.3 Choice of Sample Size

1. To determine the necessary sample size for a proposed experiment, select **DOE > Sample Size and Power**.



2. Click **Two Sample Means**.
3. Enter 0.25 for **Std Dev**, 0.5 for **Difference to detect**, and 0.95 in **Power**. Notice that the **Difference to detect** requested here is the actual difference between group means, not the scaled difference,  $\delta$ , described in the textbook.

**Sample Size**

Two Means  
Testing if two means are different from each other.

Alpha

Std Dev

Extra Parameters

Supply two values to determine the third.  
Enter one value to see a plot of the other two.

Difference to detect

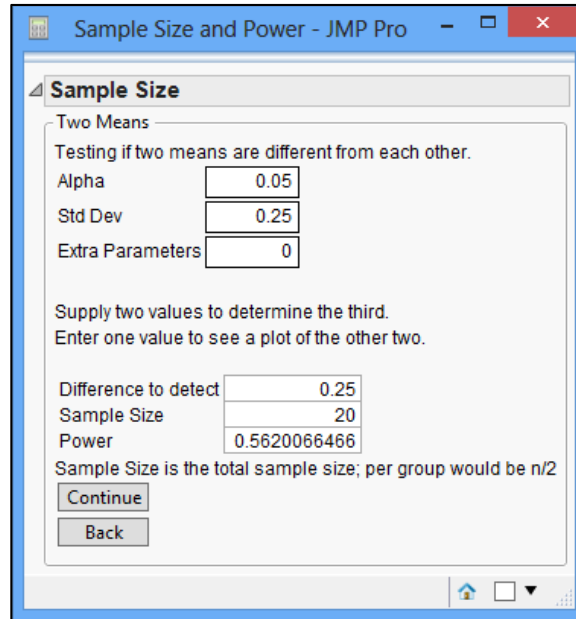
Sample Size

Power

Sample Size is the total sample size; per group would be n/2

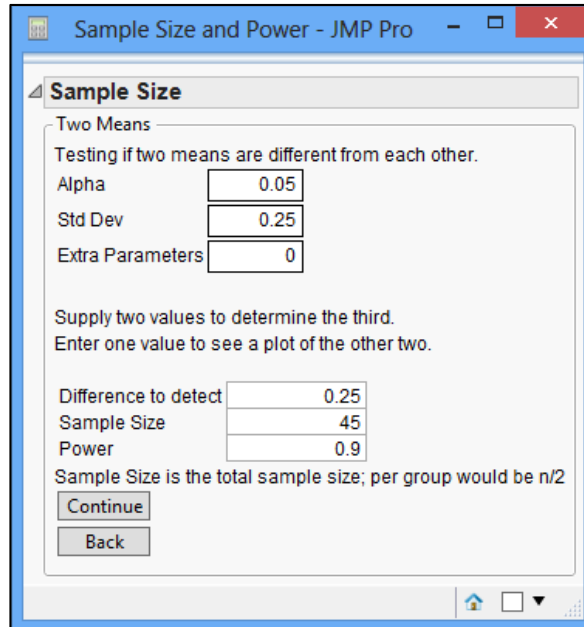
4. Click **Continue**. A value of 16 then appears in **Sample Size**. Thus, we should allocate 8 observations to each treatment ( $n_1 = n_2 = 8$ ).
5. Suppose we use a sample size of  $n_1 = n_2 = 10$ . What is the power for detecting difference of 0.25 kgf/cm<sup>2</sup>? Delete the value 0.95 from the **Power** field, change **Difference to detect** to 0.25, and set **Sample Size** to 20.
6. Click **Continue**.

14 *Design and Analysis of Experiments by Douglas Montgomery: A Supplement for Using JMP*



The power has dropped to 0.56. That is, if the model assumptions hold and the true pooled standard deviation is 0.25, only 56% of the experiments (assuming that we repeat this experiment several times) with 10 measurements from each group would successfully detect the difference of 0.25 kgf/cm<sup>2</sup>. What sample size would be necessary to achieve a power of 0.9 for this specific difference to detect?

7. Clear the **Sample Size** field and enter 0.9 for **Power**.
8. Click **Continue**.

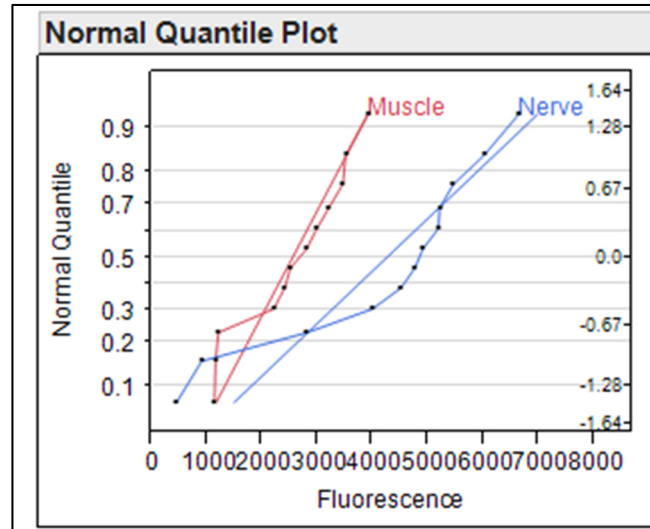


The required total sample size is 45. This means that we need at least 22.5 observations per group. Rounding up, we see that we need at least 23 observations from each group to achieve a power of at least 0.9. We could have left the **Power** field blank, specifying only that the **Difference to detect** is 0.25. The Sample Size and Power platform would then have produced a power curve, displaying **Power** as a function of **Sample Size**.

9. Select **Window > Close All**.

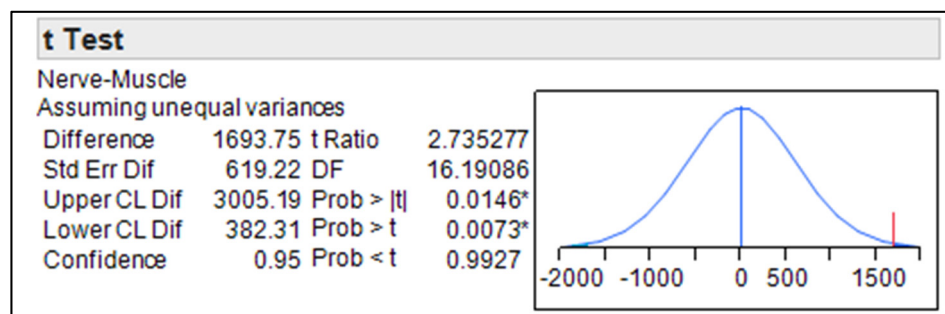
#### Example 2.1 Hypothesis Testing

1. Open Fluorescence.jmp.
2. Click **Analyze > Fit Y by X**.
3. Select *Fluorescence* for **Y, Response** and *Tissue* for **X, Factor**.
4. Click **OK**.
5. Click the red triangle next to One-way Analysis of Fluorescence by Tissue and select **Normal Quantile Plot > Plot Quantile by Actual**.



Since the slopes of the lines in the normal quantile plots are proportional to the sample standard deviations of the treatments, the difference between the slopes of the lines for *Muscle* and *Nerve* indicates that the variances may be different between the groups. As a result, we will use a form of the t-test that does not assume that the population variances are equal. Formal testing for the equality of the treatment variances is illustrated in Example 2.3 at the end of this chapter.

- Click the red triangle next to One-way Analysis of Fluorescence by Tissue and select **t-Test**.



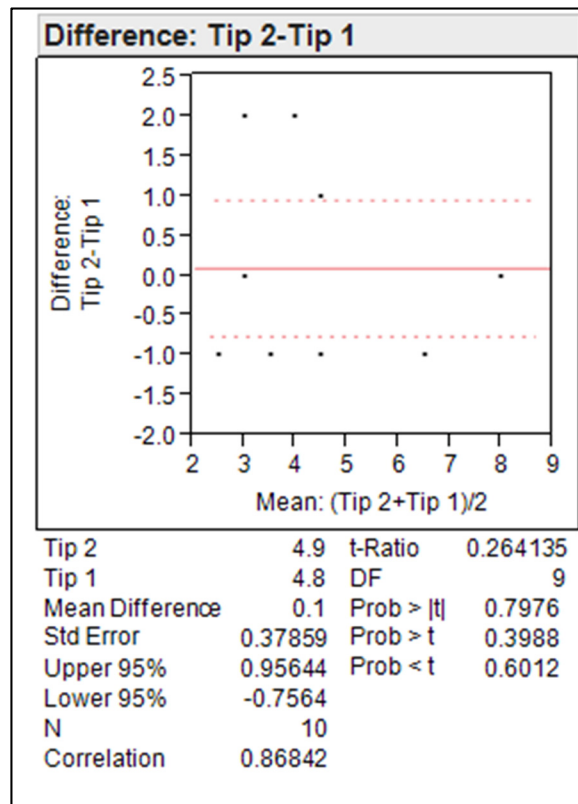
The p-value for the one-sided hypothesis test is 0.0073, which is less than the set  $\alpha$  of 0.05. We therefore reject the null hypothesis and conclude that the mean normalized fluorescence for nerve tissue is greater than the mean normalized fluorescence for

muscle tissue. Subject matter knowledge would need to determine if there is a practical difference; confidence intervals for the differences (reported in JMP) can be beneficial for this assessment.

7. Select **Window > Close All**.

### Section 2.5.1 The Paired Comparison Problem

1. Open Hardness-Testing.jmp
2. Select **Analyze > Matched Pairs**.
3. Select *Tip 1* and *Tip 2* for **Y, Paired Response**.
4. Click **OK**.

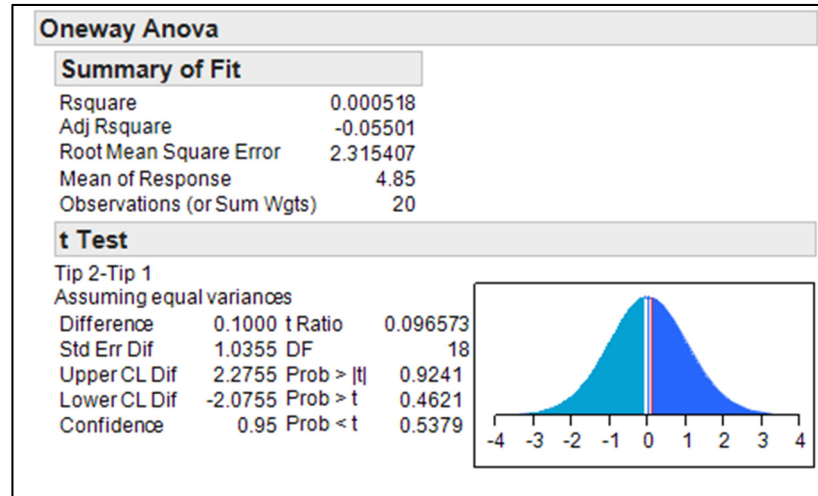


The p-value,  $\text{Prob} > |t| = 0.7976$ , indicates that there is no evidence of a difference in the performance of the two tips. This p-value is larger than the standard significance level of  $\alpha = 0.05$ .

5. Leave *Hardness-Testing.jmp* open for the next exercise.

### Section 2.5.2 Advantages of the Paired Comparison Design

1. Return to the *Hardness-Testing* table opened in the previous example.
2. Select **Tables > Stack**. This will create a file in long format with one observation per row. Most JMP platforms expect data to appear in long format.
3. Select *Tip 1* and *Tip 2* for **Stack Columns**.
4. Type “Depth” in the **Stacked Data Column** field.
5. Type “Tip” in the **Source Label Column** field.
6. Type “Hardness-Stacked” in the **Output table name** field.
7. Click **OK**.
8. *Hardness-Stacked* is now the current data table. Select **Analyze > Fit Y by X**.
9. Select *Depth* for **Y, Response** and *Tip* for **X, Grouping**.
10. Click **OK**.
11. Click the red triangle next to One-way Analysis of Depth by Tip and select **Means/Anova/Pooled t**.



The root mean square error of 2.315407 is the pooled standard deviation estimate from the t-test. Compared to the standard deviation estimate of 1.20 from the paired difference test, we see that blocking has reduced the estimate of variability considerably. Though we do not work through the details here, it would be possible to perform this same comparison for the Fluorescence data from Example 2.1.


12. Leave *Hardness-Stacked.jmp* and the Fit Y by X output window open for the next exercise.

### Example 2.3 Testing for the Equality of Variances

This example demonstrates how to test for the equality of two population variances. Section 2.6 of the textbook also discusses hypothesis testing for whether the variance of a single population is equal to a given constant. Though not shown here, the testing for a single variance may be performed in the Distribution platform.

1. Return to the Fit Y by X platform from the previous example.
2. Click the red triangle next to One-way Analysis of Depth by Tip and select **Unequal Variances**.






3. Save Hardness-Stacked.jmp.

The p-value for the F test (described in the textbook) for the null hypothesis of equal variances (with a two-sided alternative hypothesis) is 0.8393. The data do not indicate a difference with respect to the variances of depth produced from Tip 1 versus Tip 2. Due to the use of a slightly different data set, the F Ratio of 1.1492 reported here is different from the ratio of 1.34 that appears in the book. Furthermore, the textbook uses a one-sided test with an alternative hypothesis. That hypothesis is that the variance of the depth produced by Tip 1 is greater than that produced by Tip 2. Since the sample standard deviation from Tip 1 is greater than that from Tip 2, the F Ratios for the one- and two-sided tests are both equal to 1.1492, but the p-value for the one-sided test would be 0.4197.

It is important to remember that the F test is extremely sensitive to the assumption of normality. If the population has heavier tails than a normal distribution, this test will reject the null hypothesis (that the population variances are equal) more often than it should. By contrast, the Levene test is robust to departures from normality.

4. Select **Window > Close All**.



# Gain Greater Insight into Your JMP<sup>®</sup> Software with SAS Books.

Discover all that you need on your journey to knowledge and empowerment.

 [support.sas.com/bookstore](https://support.sas.com/bookstore)  
for additional books and resources.

