Chapter 434

One-Way Repeated Measures Contrasts

Introduction

This module calculates the power of a test of a contrast among the means in a one-way repeated measures design using either the multivariate test or the univariate, repeated-measures F-test as described by Maxwell and Delaney (2003) and Davis (2002).

A repeated measures design is one in which subjects are observed at a fixed set of time points. These time points do not have to be equally spaced, but they must be identical for all subjects. A hypothesis about a contrast of the means can be tested using Hotelling's T^2 test or the univariate F test.

Why Do We Promote the Multivariate Test Instead of the Univariate Test?

As we researched this procedure, we found that there were two possible methods to test the contrast: multivariate and univariate. The univariate test was at first attractive because it gives more denominator degrees of freedom. Unfortunately, the univariate approach can only by used when all correlations among the measurements at the different time points are equal (called sphericity or compound symmetry). This assumption is seldom met in practice.

Therefore, Maxwell and Delaney (2003) recommend the multivariate approach because it makes no special assumption about the correlation pattern, it forms a specific error term that depends only on the time points used in the contrast, and its type-1 error rate is exact.

For a very comprehensive discussion of when to use the each test, we refer you to Maxwell and Delaney (2003).

Technical Details

The formulas used to perform a Hotelling's T^2 power analysis provide exact answers if the above assumptions are met. These formulas can be found in many places. We use the results in Davis (2002). We refer you to that reference for more details.

Technical Details for the Multivariate Test

In one-factor case, a sample of N subjects are measured at M time points. We assume that all N subjects have the same multivariate normal distribution with mean vector μ and variance covariance matrix Σ and that Hotelling's T^2 is used for testing the null hypothesis that $\sum c_i\mu_i = 0$ versus the alternative that $\sum c_i\mu_i = A \neq 0$. The value of T^2 is computed using the formula

$$F_{1,N-1} = T^2_{1,N-1} = N\frac{\left(C'\bar{y}\right)^2}{C'SC}$$

where $\bar{y}$ is the *M*-dimensional vector of sample means, *C* is *M*-dimensional vector of contrast coefficients, and *S* is the sample variance-covariance matrix. Rather than use the T^2 distribution, we use the well-known relationship between it and the F-distribution with $df_1 = 1$ and $df_2 = (N - 1)$.

To calculate power we need the non-centrality parameter for this distribution. This non-centrality parameter is defined as follows

$$\lambda = N\frac{\left(C'\mu\right)^2}{C'\Sigma C} = N\Delta^2$$

We define Δ as the *effect size* because it provides an expression for the magnitude of the contrast of the means under the alternative hypothesis.

Using this non-centrality parameter, the power of the Hotelling's T^2 may be calculated for any value of the means and standard deviations. Since there is a simple relationship between the non-central T^2 and the non-central F, calculations are based on the non-central F using the formula

$$Power = 1 - P\left(F' > F_{1-\alpha,1,(N-1),0} | F' \sim F_{1,(N-1),\lambda}\right)$$

Technical Details for the Univariate Test

The univariate approach calculates an F-test for testing the hypothesis using

$$F_{1,(M-1)(N-1)} = \frac{MS_C}{MS_{TxS}}$$

where MS_C is the mean square for the contrast (including an adjustment for the sum of squared contrast coefficients) and MS_{TxS} is time period-by-subject interaction mean square.

To calculate power we need the non-centrality parameter. This is defined as

$$\lambda = N\frac{\left(C'\mu\right)^2}{C'\Sigma C} = N\Delta^2$$

where Σ meets the compound symmetry requirement defined below. We define Δ as the *effect size* because it provides an expression for the magnitude of the contrast of the means under the alternative hypothesis.

PASS requires the input of σ_Y and ρ. These can be estimated from a repeated measures ANOVA table which provides values for MS_S (mean square of subjects) and MS_{ST} (mean square of subject-time interaction). The parameters can then be calculated as follows

$$\hat{\rho} = \frac{F-1}{F-1+M}, \qquad F = \frac{MS_{ST}}{MS_S}, \qquad \hat{\sigma}_Y^2 = \frac{MS_{ST}}{1-\hat{\rho}}$$

It is useful to note the following expectations

$$E(MS_{ST}) = \sigma_Y^2(1-\rho)$$

$$E(MS_S) = \sigma_Y^2(1+(M-1)\rho)$$

The power is calculated using a formula similar to the one above based on the non-central F distribution

$$Power = 1 - P\left(F' > F_{1-\alpha,1,(M-1)(N-1),0} | F' \sim F_{1,(M-1)(N-1),\lambda}\right)$$

Covariance Patterns

In a repeated measures design with N subjects, each measured M times, observations within a single subject may be correlated, and a pattern for their covariance must be specified. In this case, the overall covariance matrix will have the block-diagonal form:

$$\mathbf{V} = \begin{pmatrix} \mathbf{V}_1 & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_2 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{V}_3 & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{V}_N \end{pmatrix},$$

where $\mathbf{V}_i$ is the M x M covariance submatrices corresponding to the i^{th} subject. The **0**'s represent M x M matrices of zeros giving zero covariances for observations on different subjects. This routine allows the specification of four different covariance matrix types: All ρ's Equal, AR(1), Banded(1), and Banded(2).

All ρ's Equal (Compound Symmetry)

A compound symmetry covariance model assumes that all covariances are equal, and all variances on the diagonal are equal. That is

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & \rho & \rho & \cdots & \rho \\ \rho & 1 & \rho & \rho & \cdots & \rho \\ \rho & \rho & 1 & \rho & \cdots & \rho \\ \rho & \rho & \rho & 1 & \cdots & \rho \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho & \rho & \rho & \rho & \cdots & 1 \end{pmatrix}_{M \times M}$$

where σ^2 is the subject variance and ρ is the correlation between observations on the same subject.

AR(1)

An AR(1) (autoregressive order 1) covariance model assumes that all variances on the diagonal are equal and that covariances t time periods apart are equal to $\rho^t\sigma^2$. **This choice is recommended**. That is

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & \rho^2 & \rho^3 & \ldots & \rho^{M-1} \\ \rho & 1 & \rho & \rho^2 & \ldots & \rho^{M-2} \\ \rho^2 & \rho & 1 & \rho & \ldots & \rho^{M-3} \\ \rho^3 & \rho^2 & \rho & 1 & \ldots & \rho^{M-4} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho^{M-1} & \rho^{M-2} & \rho^{M-3} & \rho^{M-4} & \ldots & 1 \end{pmatrix}_{M \times M}$$

where σ^2 is the residual variance and ρ is the correlation between observations on the same subject.

Banded(1)

A Banded(1) (banded order 1) covariance model assumes that all variances on the diagonal are equal, covariances for observations one time period apart are equal to $\rho\sigma^2$, and covariances for measurements greater than one time period apart are equal to zero. That is

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & 0 & 0 & \cdots & 0 \\ \rho & 1 & \rho & 0 & \cdots & 0 \\ 0 & \rho & 1 & \rho & \cdots & 0 \\ 0 & 0 & \rho & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}_{M \times M}$$

where σ^2 is the residual variance and ρ is the correlation between observations on the same subject.

Banded(2)

A Banded(2) (banded order 2) covariance model assumes that all variances on the diagonal are equal, covariances for observations one or two time periods apart are equal to $\rho\sigma^2$, and covariances for measurements greater than two time period apart are equal to zero. That is

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & \rho & 0 & \cdots & 0 \\ \rho & 1 & \rho & \rho & \cdots & 0 \\ \rho & \rho & 1 & \rho & \cdots & 0 \\ 0 & \rho & \rho & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}_{M \times M}$$

where σ^2 is the residual variance and ρ is the correlation between observations on the same subject.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Design tab. For more information about the options of other tabs, go to the Procedure Window chapter.

Design Tab

The Design tab contains most of the parameters and options that you will be concerned with.

Solve For

Solve For

This option specifies the parameter to be solved for. When you choose to solve for *Sample Size*, the program searches for the lowest sample size that meets the alpha and beta criterion you have specified for each of the terms. The "solve for" parameter is displayed on the vertical axis of the plot.

Test Statistic

Test Statistic Type

Specify the type of test you will use to perform the analysis.

Multivariate Test (Recommended)

The results assume that a multivariate Hotelling's T^2 will be used to analyze the data. This test requires no special assumptions about the structure of variance-covariance matrix. Often, an AR(1) correlation model is used. The denominator of the test is specific to the contrast. The error degrees of freedom is N-1. Maxwell and Delaney (2003) list this as their test of choice.

Univariate Test

Use a univariate F-Test to analyze the data. *This test requires the restrictive assumption that all off-diagonal elements of the variance-covariance matrix are equal.* This is difficult to justify in repeated measures designs where the correlations among observations on the same subject tend to diminish as time points are further apart.

The denominator of the test is an average of the variances. Sometimes it is too large. Sometimes it is too small.

The error degrees of freedom is (N-1)(M-1).

Power and Alpha

Power

This option specifies one or more values for power. Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal means when in fact the means are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when a true null hypothesis is rejected. In this procedure, a type-I error occurs when you reject the null hypothesis of equal means when in fact the means are equal.

Values must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

Note that this is a two-sided test.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Subjects)

Enter a value for the sample size (N), the number of individuals in the study. Each subject is measured 2 or more times. You may enter a single value or a range of values. A separate power calculation will be made for each value of N you enter.

Examples

10 to 100 by 10

10 30 80 90

10, 30, 80, 90

Effect Size – Number of Repeated Measurements

M (Measurements)

Enter a single value for the number of time points (repeated measurements) at which each subject is observed. It is assumed that the measurements are made at the same time points or under the same general conditions for each subject.

The value must be an integer greater than or equal to 2. The usual range is between 2 and 20.

Effect Size – Means

These options specify the means under the alternative hypothesis, one for each time point. The test statistic tests whether a contrast of these means is equal to zero. The contrast value $\Sigma Ci\mu i$ is tested against zero by the hypothesis test. The statistical hypotheses are

H0: Contrast Value = $\Sigma Ci\mu i = 0$

versus

Ha: Contrast Value = $\Sigma Ci\mu i \neq 0$.

The means represent the alternative hypothesis. They are not necessarily the means you expect. Rather, they represent the size of the contrast that you want to detect.

Means

This option specifies how you want to enter the M means for the alternative hypothesis. Your choices are List, Range, or Sequence.

Means = List

Enter a set of M hypothesized means, μi, one mean for each time point. Enter numbers separated by blanks or commas. If the number of means entered is less than M, the last mean entered will be copied forward. If too many means are entered, the extra will be ignored.

Means = Range

Enter a range of values by specifying the first and the last. The other means will be generated equi-spaced between these values.

Means = Sequence

Enter a set of means by specifying the first mean and a step-size (increment) to be added for each succeeding mean.

K (Mean Multiplier)

Enter one or more values for K, the mean multiplier. A separate power calculation is conducted for each value of K. In each analysis, all means are multiplied by K. In this way, you can determine how sensitive the power values are to the magnitude of the means without the need to change them all individually.

Note that multiplying each mean by K results in the contrast value being multiplied by K. For example, if the original means are 0, 1, and 2, setting this option to ‘1 2’ results in two sets of means used in separate analyses: 0, 1, and 2 in the first analysis and 0, 2, and 4 in the second analysis.

If you want to ignore this option, enter a one.

Effect Size – Contrast

These options specify the coefficients (Ci's) of the contrast that you are using. Each coefficient is multiplied by the corresponding mean in the mean-list above and then summed. The contrast value is $\Sigma C_i \mu_i$.

You can either specify a type of contrast and PASS will generate the coefficients for you, or you can select *Custom List* and enter the coefficient values directly.

Contrast Coefficients = Custom List

Enter a set of contrast coefficients (Ci's) separated by commas or blanks. Each coefficient is multiplied by the corresponding mean in the mean-list above and then summed. By definition, these coefficients must sum to zero. That is, $\Sigma C_i = 0$. It is recommended that $\Sigma|C_i| = 2$.

If the number of items in the list is less than M, 0's are added. If the number of items in the list is greater than M, extra items are ignored.

Examples

-1 0 0 1

-3 1 1 1

-1 -1 1 1

-1 4 -2 -3 2

Contrast Coefficients = Linear Trend

A set of coefficients is generated appropriate for testing the alternative hypothesis that there is a linear (straight-line) trend across the means. These coefficients assume that the means are equally spaced across time.

Contrast Coefficients = Quadratic

A set of coefficients is generated appropriate for testing the alternative hypothesis that the means follow a quadratic model. These coefficients assume that the means are equally spaced across time.

Contrast Coefficients = Cubic

A set of coefficients is generated appropriate for testing the alternative hypothesis that the means follow a cubic model. These coefficients assume that the means are equally spaced across time.

Contrast Coefficients = First vs Rest

A set of coefficients is generated appropriate for testing the alternative hypothesis that the first mean is different from the average of the remaining means. For example, if there were four groups, the generated coefficients would be *-3*, *1*, *1*, *1*.

Effect Size – σ (Standard Deviation)

Pattern of σ's Across Time

Specify whether the σi's vary across the measurement points or are the same.

Equal: $\sigma = \sigma_1 = \sigma_2 = ... = \sigma_M$

The σi's are constant across time. This assumption is required by the univariate F-test.

Unequal: $\sigma_1 \neq \sigma_2 \neq ... \neq \sigma_M$

The σi's can vary across time. Most researchers would agree that this is a more reasonable assumption in most cases. Because of this, the multivariate tests are easier to justify.

σ (Standard Deviation) – Equal Pattern

This is the between subject standard deviation of the response variable (Y) at a particular time point. It is assumed to be the same for all time points. As a standard deviation, the number(s) must be greater than zero.

This represents the variability from subject to subject that occurs when the subjects are treated identically.

You can enter a list of values separated by blanks or commas, in which case, a separate analysis will be calculated for each value.

σ Button

You can press the σ button and select 'Covariance Matrix' to obtain help on estimating the standard deviation from an ANOVA table.

Examples

1,4,7,10

1 4 7 10

1 to 10 by 3

σi's (σ1, σ2, …, σM) – Unequal Pattern

Specify how you want to enter the M σi's. Your choices are List, Range, or Sequence.

σi's (σ1, σ2, …, σM) = List (Unequal Pattern)

Enter a list of σi's in the box that appears to the right.

σi's (σ1, σ2, …, σM) = Range (Unequal Pattern)

Enter a range of σi's by specifying the first and the last. The other σi's will be generated between these values using a straight-line trend.

σi's (σ1, σ2, …, σM) = Sequence (Unequal Pattern)

Enter a range of σi's by specifying the first and the step-size (increment) to be added for each succeeding σi.

h (σi Multiplier)

Enter a list of h values. A separate analysis is made for each value of h. For each analysis, the σi's entered above are all multiplied by h. Thus σ1, σ2, ..., σM become hσ1, hσ2, ..., hσM. Hence, using this parameter, you can perform a sensitivity analysis about the value(s) of the standard deviation.

Note that the resulting values must all be positive, so all h's must be greater than 0.

If you want to ignore this option, enter a 1.

Effect Size – ρ (Correlation Between Measurements)

Pattern of ρ's Across Time

Specify the correlation structure of the covariance matrix. The number of diagonal elements in the matrix is equal to M.

Possible options are

- **All ρ's Equal**
  All variances on the diagonal of the within-subject variance-covariance matrix are equal to σ^2, and all covariances are equal to $\rho\sigma^2$.

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & \rho & \rho & \cdots & \rho \\ \rho & 1 & \rho & \rho & \cdots & \rho \\ \rho & \rho & 1 & \rho & \cdots & \rho \\ \rho & \rho & \rho & 1 & \cdots & \rho \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho & \rho & \rho & \rho & \cdots & 1 \end{pmatrix}_{M \times M}$$

- **AR(1)**
  All variances on the diagonal of the within-subject variance-covariance matrix are equal to σ^2, and the covariance between observations t time periods apart is $\rho^t\sigma^2$.

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & \rho^2 & \rho^3 & \ldots & \rho^{M-1} \\ \rho & 1 & \rho & \rho^2 & \ldots & \rho^{M-2} \\ \rho^2 & \rho & 1 & \rho & \ldots & \rho^{M-3} \\ \rho^3 & \rho^2 & \rho & 1 & \ldots & \rho^{M-4} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho^{M-1} & \rho^{M-2} & \rho^{M-3} & \rho^{M-4} & \ldots & 1 \end{pmatrix}_{M \times M}$$

- **Banded(1)**
  All variances on the diagonal of the within-subject variance-covariance matrix are equal to σ^2, and the covariance between observations one time period apart is $\rho\sigma^2$. Covariances between observations more than one time period apart are equal to zero.

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & 0 & 0 & \cdots & 0 \\ \rho & 1 & \rho & 0 & \cdots & 0 \\ 0 & \rho & 1 & \rho & \cdots & 0 \\ 0 & 0 & \rho & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}_{M \times M}$$

- **Banded(2)**
  All variances on the diagonal of the within-subject variance-covariance matrix are equal to σ^2, and the covariance between observations one or two time periods apart are $\rho\sigma^2$. Covariances between observations more than two time periods apart are equal to zero.

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & \rho & \rho & 0 & \cdots & 0 \\ \rho & 1 & \rho & \rho & \cdots & 0 \\ \rho & \rho & 1 & \rho & \cdots & 0 \\ 0 & \rho & \rho & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}_{M \times M}$$

ρ (Correlation)

This is the correlation, ρ, between measurements on the same subject taken at the first and second time points.

At least one value must be entered. If multiple values are entered, a separate analysis is performed for each value.

Range

$0 < \rho < 1$ (negative values are not used). A value near 0 indicates low correlation. A value near 1 indicates high correlation.

Recommended

The value of this parameter depends on time or location pattern at which measurements are taken. In their book on sample size, Machin and Campbell comment that values between 0.60 and 0.75 are typical. Some authors recommend using 0.2 when nothing is known about the actual value.

Examples

0.5

0.5 0.6 0.7

0 to 0.9 by 0.1

Example 1 – Determining Sample Size

Researchers are planning a study of the impact of a new drug on heart rate. They want to evaluate the difference in heart rate between subjects that have taken the specific drug 30 minutes before exercise and the same subjects, two days later, who exercise without the drug.

Their experimental protocol calls for a baseline heart rate measurement, followed by exercise, followed by three additional measurements 5 minutes apart. They expect a quadratic pattern in the means and want to be able to detect a 10% difference in heart rate between the two treatments.

Similar studies have found a standard deviation of the difference between scores at each time point to be between 7 and 9, and a correlation between adjacent differences on the same individual to be 0.6. The researchers assume that a first-order autocorrelation pattern adequately models the data. Since the covariances will not be equal, they decide to use the multivariate test statistic.

They decided to use a mean pattern of 0, -4, -3, and 0 to represent the differences at the four time points. They decide to look at three values of K: 1, 2, and 3.

The test will be conducted at the 0.05 significance level. What sample size is necessary to achieve 90% power over a range of possible means and standard deviations?

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **One-Way Repeated Measures Contrasts** procedure window by expanding **Means**, then **Repeated Measures**, and then clicking on **One-Way Repeated Measures Contrasts**. You may then make the appropriate entries as listed below, or open **Example 1** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Sample Size**
Test Statistic Type	**Multivariate**
Power	**0.90**
Alpha	**0.05**
M (Measurements)	**4**
Means	**List**
Means – List Values	**0 -4 -3 0**
K (Mean Multipliers)	**1 2 3**
Contrast Coefficients	**Quadratic**
Pattern of σ's Across Time	**Equal**
σ (Standard Deviation)	**7 9**
Pattern of ρ's Across Time	**AR(1)**
ρ (Correlation)	**0.6**

Annotated Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for a One-Way Repeated Measures Contrast Test
Test Type: Multivariate T² Test
Means: 0 -4 -3 0
Contrast: Quadratic {1 -1 -1 1}
ρ's: AR(1)
σi's: All Equal

Power	Subjects N	Time Points M	Effect Size Multiplier K	Mean Contrast Σcμ	Std Dev σ	Auto Corr ρ	Alpha
0.9023	21	4	1.00	7.000	7.000	0.600	0.050
0.9079	34	4	1.00	7.000	9.000	0.600	0.050
0.9055	7	4	2.00	14.000	7.000	0.600	0.050
0.9036	10	4	2.00	14.000	9.000	0.600	0.050
0.9556	5	4	3.00	21.000	7.000	0.600	0.050
0.9216	6	4	3.00	21.000	9.000	0.600	0.050

References
Maxwell, S.E. and Delaney, H.D. 2003. Designing Experiments and Analyzing Data, 2nd Edition. Psychology Press. New York.
Davis, Charles S. 2002. Statistical Methods for the Analysis of Repeated Measurements. Springer. New York, New York.
Rencher, Alvin C. 1998. Multivariate Statistical Inference and Applications. John Wiley. New York, New York.
Vonesh, E.F. and Schork, M.A. 1986. 'Sample Sizes in the Multivariate Analysis of Repeated Measurements.' Biometrics, Volume 42, pages 601-610.
Overall, J.E. and Doyle, S.R. 1994. 'Estimating Sample Sizes for Repeated Measurement Designs.' Controlled Clinical Trials, Volume 15, pages 100-123.

Report Definitions
Power is the probability of rejecting a false null hypothesis.
N is the number of subjects. Each subject is measured at two or more time points.
M is the number of time points at which each subject is measured.
K is the effect size multiplier. The original means are all multiplied by this value, resulting in a corresponding change in the effect size.
Σcμ is the linear combination of the means and the contrast coefficients. This represents the difference among the means that is being tested.
σ is the standard deviation across subjects at a given time point. It is assumed to be identical for all time points.
ρ is the (auto)correlation between observations made on a subject at the first and second time points.
Alpha is the probability of rejecting a true null hypothesis.

Summary Statements
A single-factor, repeated measures design with a sample of 21 subjects, measured at 4 time points, achieves 90% power to detect a contrast using a multivariate T² test at a 0.050 significance level. The standard deviation across subjects at the same time point is assumed to be 7.000. The correlation matrix from which the covariance matrix is generated follows an AR(1) pattern across time with a correlation of 0.600 between the first and second time point measurements. The value of the contrast applied to the hypothesized means is 7.000.

This report gives the power for each value of the other parameters. The definitions are shown in the report.

Plots Section

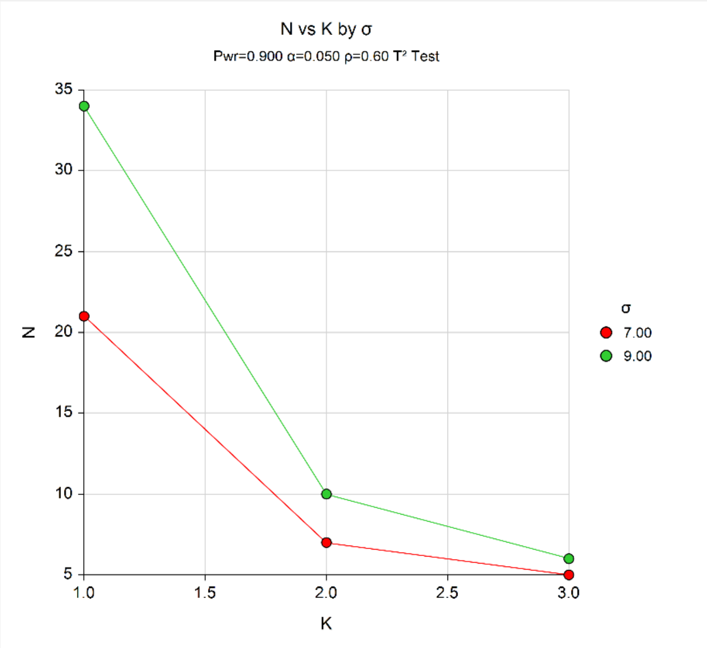

The chart shows the relationship between N, K, and σ when the other parameters in the design are held constant.

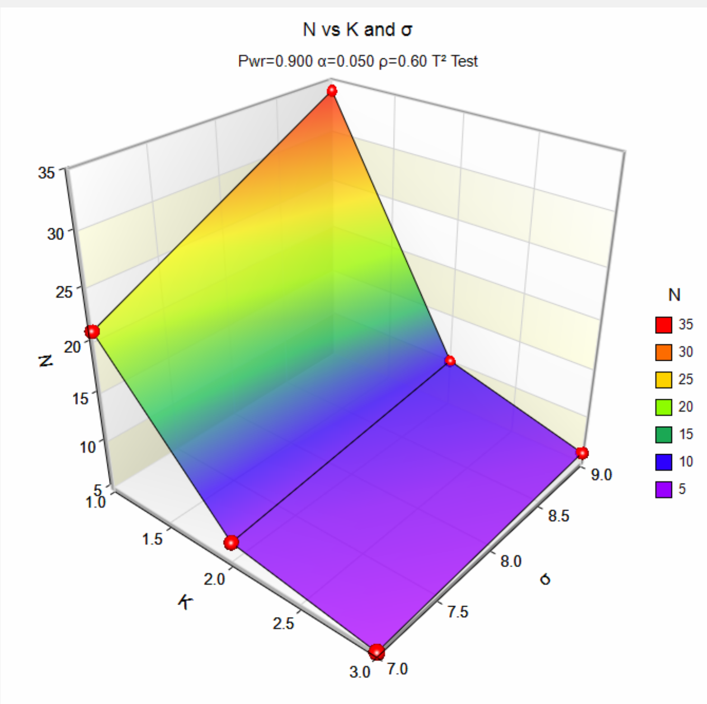

The chart shows a 3D view of the relationship between N, K, and σ when the other parameters in the design are held constant.

Example 2 – Validation using Hand Calculations

We will compute the following example by hand and then compare that with the results that **PASS** obtains. Here are the settings:

Test Statistic	Multivariable
M	3
Means	1, 2, 3
K	1
Coefficients	-2, 1, 1
σ	5
ρ	0.5
Covariance Type	AR(1)
Alpha	0.05
N	100

Using these values, we find the following

C'μ	3
C'ΣC	100
Δ	0.3
$F_{0.95,1,99}$	3.9371169
Non-Centrality	9

Power = 0.8439

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **One-Way Repeated Measures Contrasts** procedure window by expanding **Means**, then **Repeated Measures**, and then clicking on **One-Way Repeated Measures Contrasts**. You may then make the appropriate entries as listed below, or open **Example 2** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Power**
Test Statistic Type	**Multivariate**
Alpha	**0.05**
N (Subjects)	**100**
M (Measurements)	**3**
Means	**List**
Means – List Values	**1 2 3**
K (Mean Multipliers)	**1**
Contrast Coefficients	**List**
Contrast – Custom List	**-2 1 1**
Pattern of σ's Across Time	**Equal**
σ (Standard Deviation)	**5**
Pattern of ρ's Across Time	**AR(1)**
ρ (Correlation)	**0.5**

Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for the a One-Way Repeated Measures Contrast Test
Test Type: Multivariate T^2 Test
Means: 1 2 3
Contrast: -2 1 1
ρ's: AR(1)
σi's: All Equal

Power	Subjects N	Time Points M	Effect Size Multiplier K	Mean Contrast Σcμ	Std Dev σ	Auto Corr ρ	Effect Size Δ	Alpha
0.8439	100	3	1.00	3.000	5.000	0.500	0.300	0.050

Note that **PASS's** power of 0.8439 matches what we obtained above by hand.