Chapter 569

One-Way Repeated Measures

Introduction

This module calculates the power for a one-way repeated measures design. It computes power for both the univariate (F-test and F-test with Geisser-Greenhouse correction) and multivariate (Wilks' lambda, Pillai-Bartlett trace, and Hotelling-Lawley trace) approaches. It can also be used to calculate the power of *crossover* designs.

Repeated measures designs are popular because they allow a subject to serve as their own control. This usually improves the precision of the experiment. However, when the analysis of the data uses the traditional *F*-tests, additional assumptions concerning the structure of the error variance must be made. When these assumptions do not hold, the Geisser-Greenhouse correction provides reasonable adjustments so that significance levels are accurate.

An alternative to using the *F*-test with repeated measures designs is to use one of the multivariate tests: Wilks' lambda, Pillai-Bartlett trace, or Hotelling-Lawley trace. These alternatives are appealing because they do not make the strict, often unrealistic, assumptions about the structure of the variance-covariance matrix. Unfortunately, they may have less power than the *F*-test and they cannot be used in all situations.

Impact of Treatment Order, Unequal Variances, and Unequal Correlations

Treatment Order

It is important to understand under what conditions the order of treatment application matters since the design will may include subjects with different treatment arrangements, although each subject will eventually receive all treatments. Usually, the design requires a *washout* period between each two treatment applications so that the effect of one treatment does not *carryover* to the next treatment. It is always a good research practice to try several different treatment orders. **PASS** lets you study this by allowing you to set the means and standard deviations in any order you like.

A brief investigation showed that the order of treatment application only matters when the variances at each treatment are different. The pattern of the correlation matrix does have an impact, but only when the variances are different.

Unequal Variances

PASS lets you investigate the impact of unequal variances on power. Of course, different variance patterns lead to different powers. However, the special point to understand is that when the variances are different, different treatment orders result in different powers as well.

Unequal Correlations

PASS lets you investigate the impact of unequal correlations on power. Of course, different correlation patterns lead to different powers. However, the special point to understand is that when the variances are equal, different treatment orders do not result in different powers. Thus, when the variances are the same, the power values for different treatment orders will remain constant across different treatment orderings.

Conclusions

The above considerations result in the following strategy when using this procedure to power analyze (find the sample size) for a particular design.

1. If the variance is constant across treatments, you can analyze any order and the results will stand for all orderings. For example, for the case M = 3, there are six possible orders: ABC, ACB, BAC, BCA, CAB, and CBA. In this case, the resulting sample size will be identical for any order.
2. If the variances are different across treatments (for example, if the variances are treatments with larger means are assumed to be proportionally larger), you should analyze each order separately and then use the average sample size or the maximum sample size. The point is, you cannot simply analyze one order.

Assumptions

The following assumptions are made when using the *F*-test to analyze a factorial experimental design.

1. The response variable is continuous.
2. The residuals follow the normal probability distribution with mean equal to zero and constant variance.
3. The subjects are independent.

Since in a within-subject design responses coming from the same subject are not independent, assumption 3 must be modified for responses within a subject. Independence between subjects is still assumed.

4. The within-subject covariance matrix is equal for all subjects. In this type of experiment, the repeated measurements on a subject may be thought of as a multivariate response vector having a certain covariance structure.
5. When using an *F*-test, the within-subject covariance matrix is assumed to be *circular*. One way of defining circularity is that the variances of differences between any two measurements within a subject are constant for all measurements. Since responses that are close together in time (or space) often have a higher correlation than those that are far apart, it is common for this assumption to be violated. This assumption is not necessary for the validity of the three multivariate tests: Wilks' lambda, Pillai-Bartlett trace, or Hotelling-Lawley trace.

Advantages of Within-Subjects Designs

Because the response to stimuli usually varies less within an individual than between individuals, the within-subject variability is usually less than (or at most equal to) the between-subject variability. By reducing the underlying variability, the same power can be achieved with a smaller number of subjects.

Disadvantages of Within-Subjects Designs

1. *Practice effect*. In some experiments, subjects systematically improve as they practice the task being studies. In other cases, subjects may systematically get worse as the get fatigued or bored with the experimental task. Note that only the treatment administered first is immune to practice effects. Hence, experimenters should try to balance the number of subjects receiving each treatment first.
2. *Carryover effect*. In many drug studies, it is important to wash out the influence of one drug completely before the next drug is administered. Otherwise, the influence of the first drug carries over into the response to the second drug.
3. *Statistical analysis.* The statistical model is more restrictive than in a one-way factorial design since the individual responses must have certain mathematical properties.

Even in the face of all these disadvantages, repeated measures (within-subject) designs are popular in many areas of research. It is important that you recognize these problems going in so you can make sure that the design is appropriate, rather than learning of them later after the research has been conducted.

Technical Details

General Linear Multivariate Model

This section provides the technical details of the repeated measures designs that can be analyzed by **PASS**. The approximate power calculations outlined in Muller, LaVange, Ramey, and Ramey (1992) are used. Using their notation, for *N* subjects, the usual general linear multivariate model is

$$Y = XM + R$$

where each row of the residual matrix *R* is distributed as a multivariate normal

$$row_k(R) \sim N_p(0, \Sigma)$$

Note that *p* is the number of levels of the within-subject factor, *Y* is the matrix of responses, *X* is the design matrix, *M* is the matrix of regression parameters (means), and *R* is the matrix of residuals.

Hypotheses about various sets of regression parameters are tested using

$$H_0: \Theta = \Theta_0$$

$$MD = \Theta$$

where *D* is an orthonormal contrast matrix and Θ_0 is a matrix of hypothesized values, usually zeros.

Tests of the main effect may be constructed with suitable choices for *D*. These tests are based on

$$\widehat{M} = (X`X)^{-}X`Y$$

$$\widehat{\Theta} = \widehat{M}D$$

$$H_{b\times b} = \left(\widehat{\Theta} - \Theta_0\right)`[C(X`X)^{-}C`]^{-1}\left(\widehat{\Theta} - \Theta_0\right)$$

$$E_{b\times b} = D`\widehat{\Sigma}D \cdot (N - r)$$

$$T_{b\times b} = H + E$$

where *r* is the rank of *X* which is, in this case, *p* - 1. Also, C is the scalar 1 and can be ignored in this case.

Geisser-Greenhouse F-Test

Upon the assumption that Σ has compound symmetry, a size α test of $H_0: \Theta = \Theta_0$ is given by the *F* ratio

$$F = \frac{\text{tr}(H)/(p-1)}{\text{tr}(E)/[(p-1)(N-p+1)]}$$

with degrees of freedom given by

$$df_1 = p - 1$$

$$df_2 = (p-1)(N-p+1)$$

and noncentrality parameter

$$\lambda = df_1(F)$$

The assumption that Σ has compound symmetry is usually not viable. Box (1954a,b) suggested that adjusting the degrees of freedom of the above *F*-ratio could compensate for the lack of compound symmetry in Σ. His adjustment has become known as the Geisser-Greenhouse adjustment. Under this adjustment, the modified degrees of freedom and noncentrality parameter are given by

$$df_1 = (p-1)\varepsilon$$

$$df_2 = (p-1)(N-p+1)\varepsilon$$

$$\lambda = (df_1)F\varepsilon$$

where

$$\varepsilon = \frac{\text{tr}\left(D`\widehat{\Sigma}D\right)^2}{b\ \text{tr}\left(D`\widehat{\Sigma}DD`\widehat{\Sigma}D\right)}$$

The range of ε is $\frac{1}{b-1}$ to 1. When $\varepsilon = 1$, the matrix is *spherical*. When $\varepsilon = \frac{1}{b-1}$, the matrix differs maximally from sphericity.

The critical value F_{Crit} is computed using the expected value of ε to adjust the degrees of freedom. That is, the degrees of freedom of F_{Crit} are given by

$$df_1 = (p-1)E(\varepsilon)$$

$$df_2 = (p-1)(N-p+1)E(\varepsilon)$$

where

$$E(\hat{\varepsilon}) = \begin{cases} \varepsilon + \dfrac{g_1}{N-r} & \text{if } \varepsilon > \dfrac{g_1}{N-r} \\ \varepsilon/2 & \text{otherwise} \end{cases}$$

$$g_1 = \sum_{i=1}^{T} f_{ii}\,\xi_i^2 + \sum\sum_{i \neq j} \frac{f_i \xi_i \xi_j}{(\xi_i - \xi_j)}$$

$$f_i = \frac{\partial \varepsilon}{\partial \xi_i} = \frac{2\sum \xi_j}{df_1 \sum \xi_j^2} - \frac{2\lambda_i \left(\sum \xi_j\right)^2}{df_1 \left(\sum \xi_j^2\right)^2}$$

$$f_{ii} = \frac{\partial^{(2)} \varepsilon}{\partial \xi_i^{(2)}} = 2h_1 - 8h_2 + 8h_3 - 2h_4$$

$$h_1 = \frac{2}{df_1 \sum \xi_j^2}$$

$$h_2 = \frac{\xi_i \left(\sum \xi_j\right)}{df_1 \left(\sum \xi_j^2\right)^2}$$

$$h_3 = \frac{\xi_i^2 \left(\sum \xi_j\right)^2}{df_1 \left(\sum \xi_j^2\right)^3}$$

$$h_4 = \frac{\left(\sum \xi_j\right)^2}{df_1 \left(\sum \xi_j^2\right)^2}$$

and the $\xi_j' s$ are the ordered eigenvalues of $D'\Sigma D$.

Wilks' Lambda Approximate F-Test

The hypothesis $H_0: \Theta = \Theta_0$ may be tested using Wilks' likelihood ratio statistic *W*. This statistic is computed using

$$W = |ET^{-1}|$$

An *F* approximation to the distribution of *W* is given by

$$F_{df_1,df_2} = \frac{\eta/df_1}{(1-\eta)/df_2}$$

where

$$\lambda = df_1 F_{df_1,df_2}$$

$$\eta = 1 - W$$

$$df_1 = p - 1$$

$$df_2 = N - p + 1$$

Pillai-Bartlett Trace Approximate F-Test

The hypothesis $H_0: \Theta = \Theta_0$ may be tested using the Pillai-Bartlett Trace. This statistic is computed using

$$T_{PB} = \text{tr}(HT^{-1})$$

A non-central *F* approximation to the distribution of T_{PB} is given by

$$F_{df_1,df_2} = \frac{\eta/df_1}{(1-\eta)/df_2}$$

where

$$\lambda = df_1 F_{df_1,df_2}$$

$$\eta = \frac{T_{PB}}{s}$$

$$s = 1$$

$$df_1 = p - 1$$

$$df_2 = N - p + 1$$

Hotelling-Lawley Trace Approximate F-Test

The hypothesis $H_0: \Theta = \Theta_0$ may be tested using the Hotelling-Lawley Trace. This statistic is computed using

$$T_{HL} = \text{tr}(HE^{-1})$$

An *F* approximation to the distribution of T_{HL} is given by

$$F_{df_1,df_2} = \frac{\eta/df_1}{(1-\eta)/df_2}$$

where

$$\lambda = df_1 F_{df_1,df_2}$$

$$\eta = \frac{\frac{T_{HL}}{s}}{1+\frac{T_{HL}}{s}}$$

$$s = 1$$

$$df_1 = p - 1$$

$$df_2 = N - p + 1$$

The M (Mean) Matrix

In the general linear multivariate model presented above, *M* represents a matrix of regression coefficients. Since you must provide the elements of *M*, we will discuss its meaning in more detail. Although other structures and interpretations of *M* are possible, in this module we assume that the elements of *M* are the cell means. The row of *M* represents the single group and the columns of *M* represent the within-group categories.

Consider now an example in which *p* = 4. That is, there is one group of subjects that are each measured four times. The matrix *M* would appear as follows.

$$M = [\mu_1 \quad \mu_2 \quad \mu_3 \quad \mu_4]$$

To calculate the power of this design, you would need to specify appropriate values of all four means under the alternative hypothesis.

Specifying the M Matrix

When computing the power in a repeated measures analysis of variance, the specification of the *M* matrix is one of your main tasks. The program cannot do this for you. The calculated power is directly related to your choice. So, your choice for the elements of *M* must be selected carefully and thoughtfully. When authorization and approval from a government organization is sought, you should be prepared to defend your choice of *M*. In this section, we will explain how you can specify *M*.

Before we begin, it is important that you have in mind exactly what *M* is. *M* is a row of means that represent the size of the differences among the means that you want the study or experiment to detect. *M* gives the means under the alternative hypothesis. Under the null hypothesis, these means are assumed to be equal.

The D Matrix for Within-Subject Contrasts

The *D* matrix is comprised of contrasts that are applied to the columns of *M*. The choice of D does not matter as long as it is orthogonal, so an appropriate matrix will be generated for you.

Power Calculations

To calculate statistical power, we must determine distribution of the test statistic under the alternative hypothesis which specifies a different value for the regression parameter matrix *B*. The distribution theory in this case has not been worked out, so approximations must be used. We use the approximations given by Mueller and Barton (1989) and Muller, LaVange, Ramey, and Ramey (1992). These approximations state that under the alternative hypothesis, F_U is distributed as a noncentral *F* random variable with degrees of freedom and noncentrality shown above. The calculation of the power of a particular test may be summarized as follows

1. Specify values of M, Σ.
2. Determine the critical value using F_{crit} = FINV (1 - α, df_1 ,df_2), where FINV ( ) is the inverse of the central *F* distribution and α is the significance level.
3. Compute the noncentrality parameter λ.
4. Compute the power as Power = 1 - NCFPROB (F_{crit}, df_1, df_2, λ), where NCFPROB ( ) is the noncentral *F* distribution.

Covariance Matrix Assumptions

The following assumptions are made when using the *F*-test. These assumptions are not needed when using one of the three multivariate tests.

In order to use the *F* ratio to test hypotheses, certain assumptions are made about the distribution of the residuals e_{ijk}. Specifically, it is assumed that the residuals for each subject, $e_{ij1}, e_{ij2}, \cdots, e_{ijT}$, are distributed as a multivariate normal with means equal to zero and covariance matrix Σ_{ij}. Two additional assumptions are made about these covariance matrices. First, they are assumed to be equal for all subjects. That is, it is assumed that $\Sigma_{11} = \Sigma_{12} = \cdots = \Sigma_{Gn} = \Sigma$. Second, the covariance matrix is assumed to have a particular form called *circularity*. A covariance matrix is *circular* if there exists a matrix *A* such that

$$\Sigma = A + A' + \lambda I_T$$

where I_T is the identity matrix of order *T*, and λ is a constant.

This property may also be defined as

$$\sigma_{ii} + \sigma_{jj} - 2\sigma_{ij} = 2\lambda$$

One type of matrix that is circular is one that has *compound symmetry*. A matrix with this property has all elements on the main diagonal equal and all elements off the main diagonal equal. An example of a covariance matrix with compound symmetry is

$$\Sigma = \begin{bmatrix} \sigma^2 & \rho\sigma^2 & \rho\sigma^2 & \cdots & \rho\sigma^2 \\ \rho\sigma^2 & \sigma^2 & \rho\sigma^2 & \cdots & \rho\sigma^2 \\ \rho\sigma^2 & \rho\sigma^2 & \sigma^2 & \cdots & \rho\sigma^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho\sigma^2 & \rho\sigma^2 & \rho\sigma^2 & \cdots & \sigma^2 \end{bmatrix}$$

or, with actual numbers,

$$\begin{bmatrix} 9 & 2 & 2 & 2 \\ 2 & 9 & 2 & 2 \\ 2 & 2 & 9 & 2 \\ 2 & 2 & 2 & 9 \end{bmatrix}$$

An example of a matrix which does not have compound symmetry but is still circular is

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 2 & 2 & 2 & 2 \\ 3 & 3 & 3 & 3 \\ 4 & 4 & 4 & 4 \end{bmatrix} + \begin{bmatrix} 1 & 2 & 3 & 4 \\ 1 & 2 & 3 & 4 \\ 1 & 2 & 3 & 4 \\ 1 & 2 & 3 & 4 \end{bmatrix} + \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{bmatrix} = \begin{bmatrix} 4 & 3 & 4 & 5 \\ 3 & 6 & 5 & 6 \\ 4 & 5 & 8 & 7 \\ 5 & 6 & 7 & 10 \end{bmatrix}$$

Needless to say, the need to have the covariance matrix circular is a very restrictive assumption.

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:

$$V = \begin{pmatrix} V_1 & 0 & 0 & \cdots & 0 \\ 0 & V_2 & 0 & \cdots & 0 \\ 0 & 0 & V_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & V_N \end{pmatrix},$$

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 two different covariance matrix types: All *ρ*'s Equal and 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 $\rho^t\sigma^2$. That is

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

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

Banded(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.

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 differences in heart rate within subjects after this drug is administered. Each subject will be measured 5 minutes before exercise, 5 minutes after exercise, 10 minutes after exercise, and finally 15 minutes after exercise. Two days later, each subject will be measured in the same way again, except that, this time, the drug will not be administered. The differences between these two sets of measurements will be used as the basic data for the analysis.

They expect a quadratic pattern in the mean differences and want to be able to detect a 10% shift in heart rate between the two treatments. The data analyzed will be the difference between the with-drug and without-drug measurements at each time point.

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 Wilks' Lambda 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

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 1** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab	
Solve For	**Sample Size**
Test Statistic	**Multivariate: Wilks Lambda Test**
Power	**0.90**
Alpha	**0.05**
M (Measurements)	**4**
μi's (μ1, μ2, ..., μM)	**List**
μi's List	**0 -4 -3 0**
K (Means Multiplier)	**1 2 3**
Pattern of σ's Across Time	**Equal: σ = σ1 = σ2 = ⋯ = σM**
σ (Standard Deviation)	**7 9**
Pattern of ρ's Across Time	**AR(1)**
ρ (Correlation)	**0.6**

Output

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

Numeric Results

Numeric Results

Solve For: Sample Size
Test: Wilks Lambda Test
μi's: 0 -4 -3 0
σi's: All Equal
ρ's: AR(1)

Power	Sample Size N	Number of Measurements M	Means Multiplier K	Standard Deviation of Means σm	Standard Deviation σ	Within-Subject Correlation ρ	Alpha
0.9040	26	4	1	1.79	7	0.6	0.05
0.9025	39	4	1	1.79	9	0.6	0.05
0.9176	11	4	2	3.57	7	0.6	0.05
0.9314	15	4	2	3.57	9	0.6	0.05
0.9298	8	4	3	5.36	7	0.6	0.05
0.9433	10	4	3	5.36	9	0.6	0.05

Power The probability of rejecting a false null hypothesis when the alternative hypothesis is true.
N The total number of subjects in the study.
M The number of time points at which each subject is measured.
K The effect size multiplier. The original means are all multiplied by this value, resulting in a corresponding change in the effect size.
σm Standard Deviation of Means. The magnitude of differences among the hypothesized means.
σ The standard deviation across subjects at a given time point.
ρ The (auto)correlation between observations made on a subject at the first and second time points.
Alpha The probability of rejecting a true null hypothesis.

Summary Statements

A single-factor, repeated measures design with a sample of 26 subjects, measured at 4 time points, achieves 90% power to detect differences among the means using a Wilks Lambda Test at a 0.05 significance level. The standard deviation across subjects at the same time point is assumed to be 7. The correlation matrix from which the covariance matrix is generated follows an AR(1) pattern across time with a correlation of 0.6 between the first and second time point measurements. The standard deviation of the hypothesized means is 1.79.

Dropout-Inflated Sample Size

Dropout Rate	Sample Size N	Dropout-Inflated Enrollment Sample Size N'	Expected Number of Dropouts D
20%	26	33	7
20%	39	49	10
20%	11	14	3
20%	15	19	4
20%	8	10	2
20%	10	13	3

Dropout Rate	The percentage of subjects (or items) that are expected to be lost at random during the course of the study and for whom no response data will be collected (i.e., will be treated as "missing"). Abbreviated as DR.
N	The evaluable sample size at which power is computed. If N subjects are evaluated out of the N' subjects that are enrolled in the study, the design will achieve the stated power.
N'	The total number of subjects that should be enrolled in the study in order to obtain N evaluable subjects, based on the assumed dropout rate. After solving for N, N' is calculated by inflating N using the formula N' = N / (1 - DR), with N' always rounded up. (See Julious, S.A. (2010) pages 52-53, or Chow, S.C., Shao, J., Wang, H., and Lokhnygina, Y. (2018) pages 32-33.)
D	The expected number of dropouts. D = N' - N.

Dropout Summary Statements

Anticipating a 20% dropout rate, 33 subjects should be enrolled to obtain a final sample size of 26 subjects.

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.

Edwards, L.K. 1993. Applied Analysis of Variance in the Behavior Sciences. Marcel Dekker. New York.

Muller, K. E., and Barton, C. N. 1989. 'Approximate Power for Repeated-Measures ANOVA Lacking Sphericity.' Journal of the American Statistical Association, Volume 84, No. 406, pages 549-555.

Muller, K. E., LaVange, L.E., Ramey, S.L., and Ramey, C.T. 1992. 'Power Calculations for General Linear Multivariate Models Including Repeated Measures Applications.' Journal of the American Statistical Association, Volume 87, No. 420, pages 1209-1226.

Simpson, S.L., Edwards, L.J., Muller, K.E., Sen, P.K., and Styner, M.A. 2010. 'A linear exponent AR(1) family of correlation structures.' Statistics in Medicine, Volume 29(17), pages 1825-1838.

Naik, D.N. and Rao, S.S. 2001. 'Analysis of multivariate repeated measures data with a Kronecker product structured covariance matrix.' Journal of Applied Statistics, Volume 28 No. 1, pages 91-105.

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

Plots Section

Plots

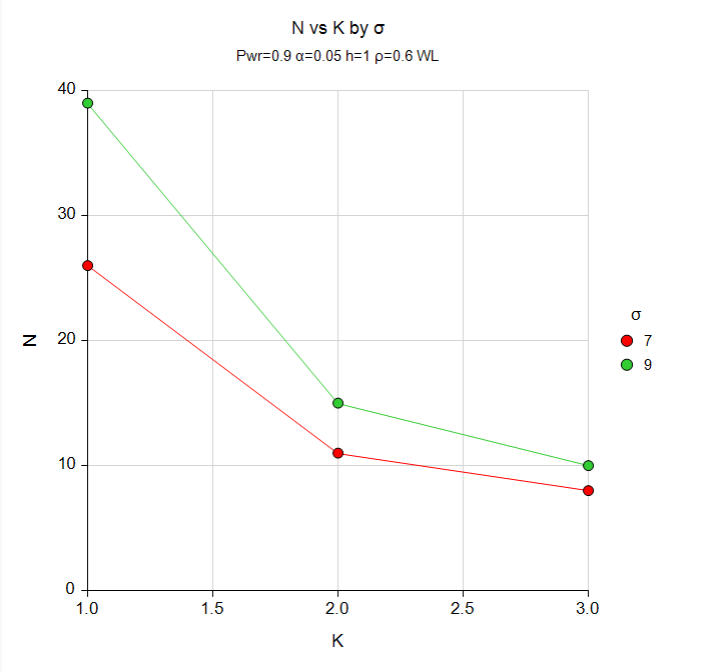

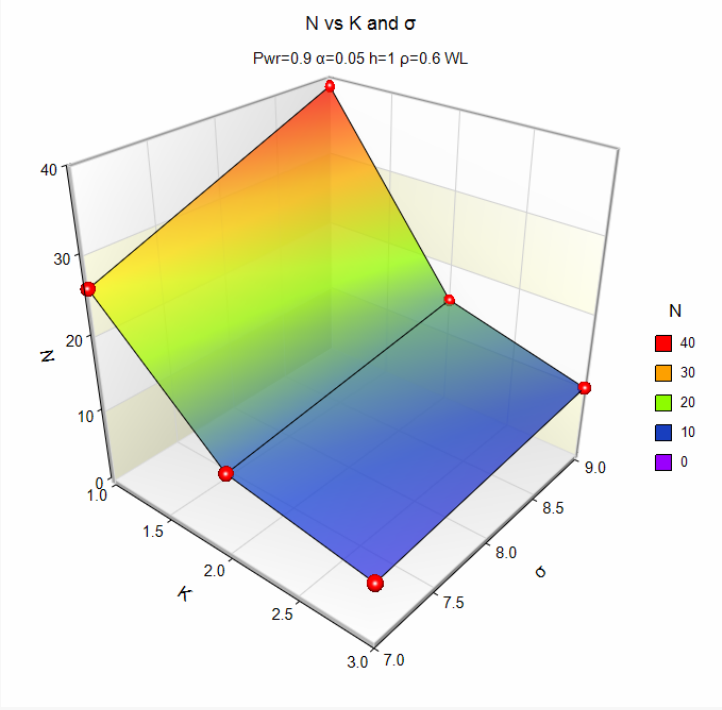

The plots show the relationship between N, K, and σ when the other parameters in the design are held constant.

Example 2 – Validation using PASS's Repeated Measures Procedure

To validate this procedure, we will run the first scenario through **PASS**'s Repeated Measures Analysis procedure. We did this and obtained the following report.

Design Report

Solve For: Sample Size Based on the Wilks' Lambda Approximate F-Test (All Terms)

			Sample Size						
Term*	Test	Power	Average Group n	Total N	Effect Multiplier K	Standard Deviation of Effects σm	Standard Deviation σ	Effect Size σm / σ	Alpha
W1 (4)	Wilks	0.90402	26	26	1	1.78536	2.58147	0.6916	0.05

* The numbers in parentheses represent the number of levels associated with the factor.

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 2** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For	**Sample Size**
Test Statistic	**Multivariate: Wilks Lambda Test**
Power	**0.90**
Alpha	**0.05**
M (Measurements)	**4**
μi's (μ1, μ2, ..., μM)	**List**
μi's List	**0 -4 -3 0**
K (Means Multiplier)	**1**
Pattern of σ's Across Time	**Equal: σ = σ1 = σ2 = ⋯ = σM**
σ (Standard Deviation)	**7**
Pattern of ρ's Across Time	**AR(1)**
ρ (Correlation)	**0.6**

Output

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

Numeric Results

Solve For: Sample Size
Test: Wilks Lambda Test
μi's: 0 -4 -3 0
σi's: All Equal
ρ's: AR(1)

Power	Sample Size N	Number of Measurements M	Means Multiplier K	Standard Deviation of Means σm	Standard Deviation σ	Within-Subject Correlation ρ	Alpha
0.904	26	4	1	1.79	7	0.6	0.05

This new, simplified procedure has obtained identical results with the previously validated procedure.