

Chapter 372

Mixed Models Tests for Two Means at the End of Follow-Up in a 2-Level Hierarchical Design

Introduction

This procedure calculates power and sample size for a two-level longitudinal design in which subjects (level-two unit) are randomly assigned to one of two groups. Each subject is measured at several time points and the goal of the study is to compare the group means of at the final time point. This procedure assumes that the group means are identical at the beginning of the study (which they often are in a randomized trial).

All subjects in a group are assumed to have a common (fixed) slope. Each subject is assigned to receive one of two possible interventions.

Technical Details

Our formulation comes from Ahn, Heo, and Zhang (2015), chapter 5, section 5.4.4, pages 165-167. The longitudinal mixed model that is adopted is a re-parameterized version of fixed-slope model

$$Y_{ij} = \beta_0 + \xi X_{ij} + \tau T_{ij} + \delta_f X_{ij} T_{ij} + u_i + e_{ij}$$

in which the time variable is rescaled using $S_{ij} = T_{ij} - T_{end}$. Therefore, the S_{ij} become $-T_{end}, \dots, 0$.

Let $\beta'_0 = \beta_0 + \tau T_{end}$ and $\delta_e = \xi + \delta_f T_{end}$. We assume that, because of randomization, $\xi = 0$. The mixed model then becomes

$$Y_{ij} = \beta'_0 + \delta_e X_{ij} + \tau S_{ij} + \delta_f X_{ij} S_{ij} + u_i + e_{ij}$$

where

Y_{ij} is the continuous response of the j^{th} measurement in the i^{th} subject.

β_0 is the fixed intercept.

Mixed Models Tests for Two Means at the End of Follow-Up in a 2-Level Hierarchical Design

- X_{ij} is an indicator variable that is 1 if subject i is assigned to group 1 and 0 otherwise.
- T_{ij} is the time value. It is assumed that $T_{ij} = j - 1$ for all i and j .
- S_{ij} is the rescaled time value. It is assumed that $S_{ij} = j - (M - 1)$ for all i and j .
- ξ is the intervention effect at baseline and is assumed to be zero.
- δ_f is the treatment-by-time interaction effect.
- δ_e is the intervention effect at the end of the study. By construction, $\delta_e = \mu_1 - \mu_2$ at the final time point.
- u_i is a random effect (subject-specific intercept) term for the i^{th} cluster that is distributed as $N(0, \sigma_u^2)$.
- e_{ij} is a random effect for the j^{th} subject in the i^{th} cluster that is distributed as $N(0, \sigma_e^2)$.
- σ_u^2 is variance of the subject random effects.
- σ_e^2 is variance of the measurement random effects.
- σ^2 is the variance of Y .
- ρ is the correlation between measurements on the same subject.
- K_1 is the number of subjects in group 1.
- K_2 is the number of subjects in group 2.
- λ is K_1/K_2 .
- M is the number of measurements per subject.
- $V(T)$ is $\sum_{j=1}^M (T_j - \bar{T})^2 / M$. Note that $V(T) = V(S)$ and $\bar{S} = -\bar{T}$.
- $CV(S)$ is $\sqrt{V(S)}/\bar{S}$.
- $C_{(2)}$ is $1 + \frac{(1-\rho)}{CV(S)^2 f}$.

The test of significance of δ_e in the mixed model analysis is the test statistic of interest. It tests the difference of the two means at the final time point. The power calculations assume that the estimated value of δ , called $\hat{\delta}$, has an approximate normal distribution with mean δ and known variance.

The power is calculated using

$$Power = \Phi \left\{ \frac{|\delta_e|}{\sigma} \sqrt{\frac{K_2 M}{(1 + \lambda^{-1}) C_{(2)} f}} - \Phi(1 - \alpha/2) \right\}$$

The values of the other parameters when requested are found using a binary search based on this formula.

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 from the other parameters. The parameters that may be selected are $\mu_1 - \mu_2$, Power, K1, and M.

Under most situations, you will select either *Power* to calculate power or *K1* to calculate the number of clusters. Occasionally, you may want to fix the number of clusters and find the necessary cluster size.

Note that the value selected here always appears as the vertical axis on the charts.

The program is set up to calculate power directly. To find appropriate values of the other parameters, a binary search is made using an iterative procedure until an appropriate value is found. This search considers integer values of M only.

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.

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

If your only interest is in determining the appropriate sample size for a confidence interval, set power to 0.5.

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.

Values must be between zero and one. Usually, the value of 0.05 is used for alpha and this has become a standard. 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.

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 – Subjects and Measurements

K1 (Group 1 Subjects)

This is the number of subjects (level-two units) in group 1 (treatment group).

This value must be a positive number greater than 1.

You can enter a list of values such as "10 20 30". A separate analysis will be run for each element in the list.

Mixed Models Tests for Two Means at the End of Follow-Up in a 2-Level Hierarchical Design**K2 (Group 2 Subjects)**

This is the number of subjects in group 2 (control group).

This value must be a positive number greater than 1.

Using Multiples of K1

If you simply want a multiple of the value for group one, you would enter the multiple followed by "K1", with no blanks. If you want to use K1 directly, you do not have to enter the leading "1". For example, all of the following are valid entries:

10K1 2K1 0.5K1 K1.

You can use a list of values such as "10 20 30" or "K1 2K1 3K1".

M (Measurements Per Subject)

This is the number of measurements (time points) per subject. It is assumed that these measurements are equally spaced through time.

This value must be a positive integer that is at least 2.

List

You can use a list of values such as "4 6 8". A separate analysis will be run for each element in the list.

Effect Size **$\mu_1 - \mu_2$ (Mean Difference)**

Enter one or more values for the mean difference where μ_1 is the mean of group 1 at the final measurement and μ_2 is the mean of group 2 at the final measurement.

This sample size estimate assumes that the mean difference at the beginning is zero. This assumption is reasonable when the subjects are randomly assigned to groups at the beginning of a study. This procedure assumes that the times are 0, 1, 2, ..., (M-1) and the value of T at the last measurement point is (M-1).

The power is the probability of detecting a difference of at least this amount. This value is not necessarily the true difference. Rather, it is the difference you want to detect.

The difference, $\mu_1 - \mu_2$, can be any non-zero value (positive or negative). Since this procedure uses a two-sided test statistic, you will get the same result with either positive or negative values.

Syntax

You can enter a single value such as 1 or a series of values such as 1 2 3 or 1 to 2 by 0.2.

When a series of values is entered, **PASS** will generate a separate calculation result for each value of the series.

 σ (Standard Deviation)

Enter the standard deviation of Y. This standard deviation applies for both groups. A good estimate of this value is the average squared residual from a previous study.

Note that σ must be a positive number. You can enter a single value such as 5 or a series of values such as 1 3 5 7 9 or 1 to 9 by 2.

Press the small 'σ' button to the right to obtain calculation options for estimating the standard deviation.

 ρ (Within Subject Correlation)

Enter one or more values of the within-subject correlation coefficient. It may be interpreted as the correlation between any two observations on the same subject. Possible values are from 0 to just below 1. Typical values are between 0.0001 and 0.5.

Example 1 – Calculating Power

Researchers are planning a study of the impact of a new drug on heart rate. They want to compare the end of follow-up heart rate of subjects who take the new drug and subjects who take a standard drug. Their experimental protocol calls for a baseline heart rate measurement, followed by administration the drug, followed by three additional measurements ten minutes apart. They want to be able to detect a difference of 10 in the final measurement heart-rate means between the two treatments. They want a sensitivity analysis by considering a range of differences from 9 to 11.

Similar studies have found a standard deviation of 9.2. These studies also showed a correlation among measurements on the same individual of 0.5. The two-sided test will be conducted at the 0.05 significance level. The desired power is 90%. They are planning on dividing subjects equally between the treatment and control groups. They want to investigate $K1 = K2 = 10$ to 25 by 5.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Mixed Models Tests for Two Means at the End of Follow-Up in a 2-Level Hierarchical Design** procedure window. 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**.

<u>Option</u>	<u>Value</u>
Design Tab	
Solve For	Power
Alpha.....	0.05
K1 (Group 1 Subjects)	10 15 20 25
K2 (Group 2 Subjects)	K1
M (Measurements Per Subject).....	4
$\mu_1 - \mu_2$ (Mean Difference).....	9 10 11
σ (Standard Deviation)	9.2
ρ (Within Subject Correlation).....	0.5

Annotated Output

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

Numeric Results

Numeric Results for a Two-Level, Longitudinal Mixed Model									
	Total Subjects	Group 1 Subjects	Group 2 Subjects	Measurements /Subject	Mean Diff at End of Follow-Up	Standard Deviation	Correlation Within a Subject		
Power	K	K1	K2	M	$\mu_1 - \mu_2$	σ	ρ	Alpha	
0.6601	20	10	10	4	9.00	9.20	0.500	0.050	
0.8279	30	15	15	4	9.00	9.20	0.500	0.050	
0.9186	40	20	20	4	9.00	9.20	0.500	0.050	
0.9634	50	25	25	4	9.00	9.20	0.500	0.050	
0.7506	20	10	10	4	10.00	9.20	0.500	0.050	
0.8977	30	15	15	4	10.00	9.20	0.500	0.050	
0.9615	40	20	20	4	10.00	9.20	0.500	0.050	
0.9864	50	25	25	4	10.00	9.20	0.500	0.050	
0.8264	20	10	10	4	11.00	9.20	0.500	0.050	
0.9443	30	15	15	4	11.00	9.20	0.500	0.050	
0.9839	40	20	20	4	11.00	9.20	0.500	0.050	
0.9957	50	25	25	4	11.00	9.20	0.500	0.050	

Mixed Models Tests for Two Means at the End of Follow-Up in a 2-Level Hierarchical Design

References

Ahn, C., Heo, M., and Zhang, S. 2015. Sample Size Calculations for Clustered and Longitudinal Outcomes in Clinical Research. CRC Press. New York.

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

K is the total number of subjects.

K1 and K2 are the number of subjects in groups 1 and 2, respectively.

M is the number of equally-spaced measurement times.

$\mu_1 - \mu_2$ is the difference in the two group means at the final measurement time: M - 1.

σ is the standard deviation of the response.

ρ is the correlation of the measurements on a subject.

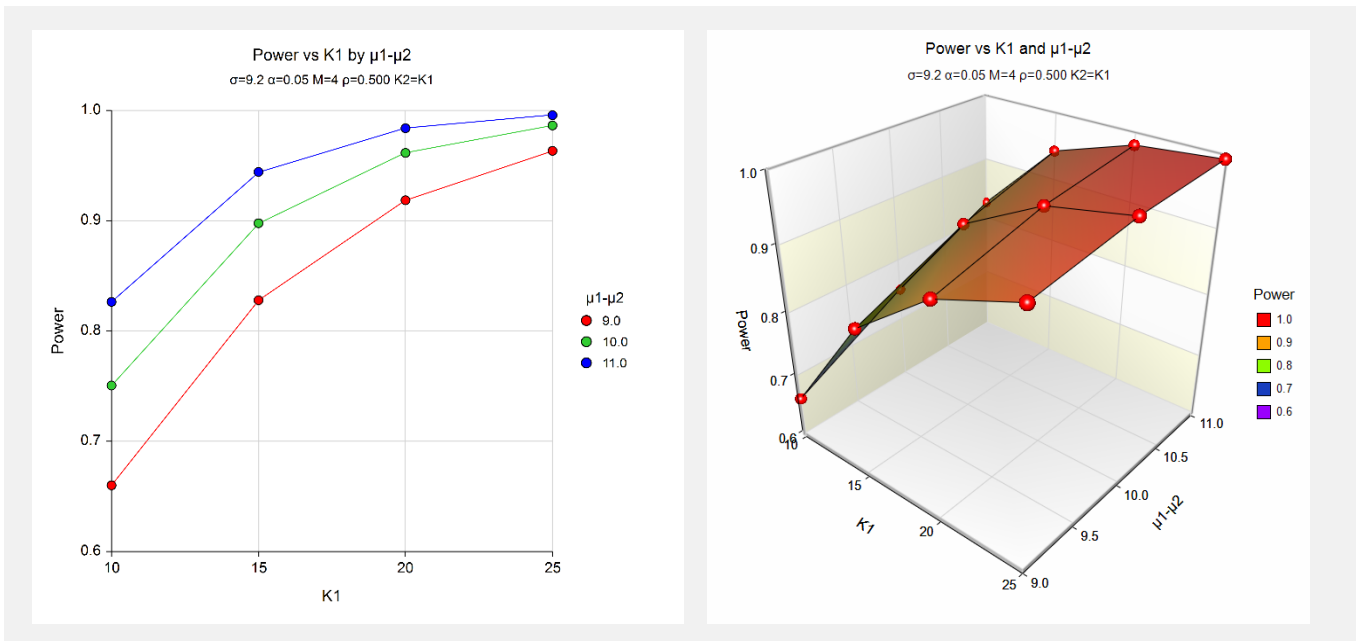
Alpha is the probability of rejecting a true null hypothesis, that is, rejecting when the slopes are actually equal.

Summary Statements

A total of 20 subjects, obtained as 10 in group one and 10 in group two, were each measured at 4 time points. The study achieves 66% power to detect a difference of 9.00 between the (fixed) group means at the last time. The standard deviation is 9.20. The correlation between measurements within a subject is 0.500. A test based on a mixed-model analysis is anticipated at a significance level of 0.050.

This report shows the power for each of the scenarios.

Plots Section



These plots show the power versus K1 for the various differences.

Example 2 – Calculating Sample Size (Number of Subjects)

Continuing with the last example, suppose the researchers want to determine the number of subjects needed to achieve 90% power.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Mixed Models Tests for Two Means at the End of Follow-Up in a 2-Level Hierarchical Design** procedure window. 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**.

<u>Option</u>	<u>Value</u>
Design Tab	
Solve For	K1 (Group 1 Subjects)
Power.....	0.90
Alpha.....	0.05
K2 (Group 2 Subjects)	K1
M (Measurements Per Subject).....	4
$\mu_1 - \mu_2$ (Mean Difference).....	9 10 11
σ (Standard Deviation)	9.2
ρ (Within Subject Correlation).....	0.5

Output

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

Numeric Results

Numeric Results for a Two-Level, Longitudinal Mixed Model								
	Total Subjects	Group 1 Subjects	Group 2 Subjects	Measurements /Subject	Mean Diff at End of Follow-Up	Standard Deviation	Correlation Within a Subject	Alpha
Power	K	K1	K2	M	$\mu_1 - \mu_2$	σ	ρ	
0.9050	38	19	19	4	9.00	9.20	0.500	0.050
0.9154	32	16	16	4	10.00	9.20	0.500	0.050
0.9109	26	13	13	4	11.00	9.20	0.500	0.050

This report shows the number of subjects required for each of the scenarios.

Example 3 – Validation using Ahn, Heo, and Zhang (2015)

Ahn, Heo, and Zhang (2015) page 167 provide a table in which several scenarios are reported. We will validate this procedure by duplicating the first row of their table. The following parameter settings were used: Power = 0.80; $\mu_1 - \mu_2 = 0.4$; $\sigma = 1$; $\rho = 0.1$; $M = 5$; and $\alpha = 0.05$. The reported value of K_1 is 63. The achieved power is 0.801.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Mixed Models Tests for Two Means at the End of Follow-Up in a 2-Level Hierarchical Design** procedure window. You may then make the appropriate entries as listed below, or open **Example 3** by going to the **File** menu and choosing **Open Example Template**.

Design Tab

Solve For **K1 (Group 1 Subjects)**
 Power..... **0.80**
 Alpha..... **0.05**
 K2 (Group 2 Subjects) **K1**
 M (Measurements Per Subject)..... **5**
 $\mu_1 - \mu_2$ (Mean Difference)..... **0.4**
 σ (Standard Deviation) **1**
 ρ (Within Subject Correlation)..... **0.1**

Output

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

Numeric Results

Numeric Results for a Two-Level, Longitudinal Mixed Model									
	Total Subjects	Group 1 Subjects	Group 2 Subjects	Measure- ments /Subject	Mean Diff at End of Follow-Up	Standard Deviation	Correlation Within a Subject		
Power	K	K1	K2	M	$\mu_1 - \mu_2$	σ	ρ	Alpha	
0.8013	126	63	63	5	0.40	1.00	0.100	0.050	

PASS also calculates the value of K_1 to be 63. The calculated power is also 0.801.