Chapter 371

Mixed Models Tests for Two Means at the End of Follow-Up in a 3-Level Hierarchical Design (Level-3 Randomization)

Introduction

This procedure calculates power and sample size for a three-level hierarchical design which is randomized at the third level. Each subject (level-2 unit) is measured at several time points (level-1 units) and the goal of the study is to compare the group means of at the final time point, i.e., at the end of follow-up.

The study is assumed to be longitudinal, so repeated measurements (level-1 units) are nested in patients (level-2 units) which are nested in clinics (level-3 units).

Each clinic (level-3 unit) is randomized into one of two intervention groups; e.g., treatment and control.

Technical Details

Our formulation comes from Ahn, Heo, and Zhang (2015), chapter 6, section 6.4.4, pages 209-211. The longitudinal mixed model that is adopted is a re-parameterized version of fixed-slope model

$$Y_{ij} = \beta_0 + \xi X_{ijk} + \tau T_{ijk} + \delta_f X_{ijk} T_{ijk} + u_i + u_{j(i)} + e_{ijk}$$

in which the time variable is rescaled using $S_k = T_k - T_{end}$. Therefore, the S_k 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

$$Y_{ijk} = \beta_0' + \delta_e X_i + \tau S_k + \delta_f X_i S_k + u_i + u_{j(i)} + e_{ij}$$

where

    Y_{ijk} is the continuous response of the k^{th} level-1 unit of the jth level-2 unit of the ith level-3 unit.

    β_0 is the fixed intercept.

    δ_f is the treatment-by-time interaction effect. It can be ignored since S_k is zero at the final observation because of the re-parameterization.

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. This is the term that will be tested for significance.
X_{ijk}	is an indicator variable that is 1 if i^{th} unit is in group 1 and 0 if it is in group 2.
S_k	is the rescaled time value. It is assumed that $S_k = k - (M - 1)$ for all i and j.
ξ	is the intervention effect at baseline and is assumed to be zero.
u_i	is the level-3 random intercept effect for the i^{th} level-3 unit. It is distributed as $N(0, \sigma_3^2)$.
$u_{j(i)}$	is the level-2 random intercept effect for the i^{th} level-3 unit. It is distributed as $N(0, \sigma_2^2)$.
e_{ijk}	is a random error term which is distributed as $N(0, \sigma_e^2)$.
σ^2	is the variance of Y_{ijk}, where $\sigma^2 = \sigma_e^2 + \sigma_2^2 + \sigma_3^2$.
ρ_1	is the correlation among level-1 units which are in a particular level-2 unit. For fixed models like this, $\rho_1 = Corr(Y_{ijk}, Y_{ijk'}) = (\sigma_2^2 + \sigma_3^2)/\sigma^2$.
ρ_2	is the correlation among level-2 units which are in a particular level-3 unit. For fixed models like this, $\rho_2 = Corr(Y_{ijk}, Y_{ij'k'}) = (\sigma_3^2)/\sigma^2$.
C_1	is the number of level-3 units assigned to group 1.
C_2	is the number of level-3 units assigned to group 2.
K	is the number of level-2 units per level-3 unit.
M	is the number of level-1 units per level-2 unit.
$V(T)$	is $\sum_{j=1}^{M}(T_j - \bar{T})^2 / M$. Note that $V(S) = V(T)$ and $\bar{S} = -\bar{T}$.
$CV(S)$	is $\sqrt{V(S)}/\bar{S}$.
$C_{(3)}$	is $1 + \frac{(1-\rho_1)}{CV(S)^2 f_3}$.

The test of significance of the X_i term in the mixed model analysis is the test statistic of interest. It tests the difference of the two group means. Since these are asymptotic results, the specific type of mixed model is not stated.

Assume that $\delta_e = \mu_1 - \mu_2$ is to be tested using a z-test (large sample). The statistical hypotheses are $H_0: \delta_e = 0$ vs. $H_a: \delta_e \neq 0$. The test statistic is the regression coefficient of the X_i term in a mixed model.

The power can be calculated using

$$Power = \Phi\left\{ \left|\frac{\delta_e}{\sigma}\right| \sqrt{\frac{C_2 KM}{\left[f_3\left(1 + \frac{1}{\lambda}\right)C_{(3)}\right]}} - \Phi(1 - \alpha/2)\right\}$$

where $\lambda = C_1/C_2$ and $f_3 = 1 + M(K - 1)\rho_2 + (M - 1)\rho_1$.

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*, $C1$, K, and M.

Under most situations, you will select either *Power* to calculate power or $C1$ to calculate the number of level-3 unit. Occasionally, you may want to fix $C1$ and $C2$ and calculate K or M.

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 only integer values of $C1$, K, or M.

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*.

PASS Sample Size Software   **NCSS.com**

Mixed Models Tests for Two Means at the End of Follow-Up in a 3-Level Hier. Design (Level-3 Rand.)

Sample Size

C1 (Level 3 Units Assigned to Group 1)

This is the number of level 3 units in group 1 (treatment group). The sample size for this group is equal to Ceiling(C1 * K * M).

Examples of level 3 units are schools, clinics, hospitals, or cities.

A longitudinal study might include clinics as the level 3 units, patients as the level 2 units, and repeated measurements are the level 1 units.

This value must be a positive integer.

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

C2 (Level 3 Units Assigned to Group 2)

This is the number of level 3 units in group 2 (control group). The sample size for this group is equal to Ceiling(C2 * K * M).

Examples of level 3 units are clinics, hospitals, or cities.

A *longitudinal* study might include clinics as the level 3 units, patients as the level 2 units, and repeated measurements are the level 1 units.

This value must be a positive integer.

Using Multiples of C1

If you simply want a multiple of the value for group 1, you would enter the multiple followed by "C1", with no blanks. If you want to use C1 directly, you do not have to enter the leading "1". For example, all of the following are valid entries:

C1 2C1 0.5C1 1.5C1.

You can use a list of values such as "10 20 30" or "C1 2C1 3C1".

K (Level 2 Units Per Level 3 Unit)

This is the average number of level 2 units in each level 3 unit. Examples of level 2 units are classrooms or patients.

A *longitudinal* study might include clinics as the level 3 units, patients as the level 2 units, and repeated measurements are the level 1 units.

This value must be a positive number that is at least 1.

It can be a decimal (fractional) number such as '2.7'. The resulting total number of level 1 units will be rounded up to the next largest integer.

You can use a list of values such as "5 10 15". A separate analysis will be run for each element in the list.

M (Level 1 Units Per Level 2 Unit)

This is the average number of level 1 units in each level 2 unit. Examples of level 1 units are repeated measurements or students.

A *cross-sectional* study might include schools as the level 3 units, classrooms as the level 2 units, and students are the level 1 units. A *longitudinal* study might include clinics as the level 3 units, patients as the level 2 units, and repeated measurements are the level 1 units.

This value must be a positive number that is at least 1.

It can be a decimal (fractional) number such as '2.7'. The resulting total number of level 1 units will be rounded up to the next largest integer. You can use a list of values such as "5 10 15". A separate analysis will be run for each element in the list.

Effect Size

μ1 - μ2 (Mean Difference)

Enter one or more values for the mean difference where μ1 is the mean of group 1 at the final measurement and μ2 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, μ1 - μ2, can be any non-zero value (positive or negative). Since this procedure uses a is 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 an individual observation, Y_{ijk}. This standard deviation applies for both groups.

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.

ρ1 (Correlation Among Level 1 Units)

This is the correlation of each pair of level 1 units in a particular level 2 unit.

Possible values are from 0 to just below 1. For fixed slope models, ρ1 ≥ ρ2. Typical values are between 0.0001 and 0.5.

You may enter a single value or a list of values.

ρ2 (Correlation Among Level 2 Units)

This is the correlation of each pair of level 2 units in a particular level 3 unit.

Possible values are from 0 to just below 1. For fixed slope models, ρ1 ≥ ρ2. Typical values are between 0.0001 and 0.5.

You may enter a single value or a list of values.

PASS Sample Size Software

Mixed Models Tests for Two Means at the End of Follow-Up in a 3-Level Hier. Design (Level-3 Rand.)

NCSS.com

Example 1 – Calculating Power

Suppose that a longitudinal trial is planned in which there will be repeated measurements (level-1) which are nested in patients (level-2) which are nested in clinics (level-3). This analysis will calculate the power for detecting the difference between two interventions in heart rate. The researchers want to compare the end of follow-up heart rate of subjects who take a new drug to that of 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.

The analysis will be a mixed model of continuous data. The following parameter settings are to be used for the power analysis: $\mu_1 - \mu_2 = 9, 10, 11$; $\sigma = 9.2$; $\rho_1 = 0.1$; $\rho_2 = 0.05$; K = 5; M = 5; $\alpha = 0.05$; and C1 = C2 = 2 to 6 by 2.

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 3-Level Hier. Design (Level-3 Rand.)** 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**.

Option	**Value**
Design Tab	
Solve For	**Power**
Alpha...	**0.05**
C1 (Level 3 Units Assigned to Group 1)	**2 4 6**
C2 (Level 3 Units Assigned to Group 2)	**C1**
K (Level 2 Units Per Level 3 Unit)	**5**
M (Level 1 Units Per Level 2 Unit)..................	**5**
μ1 - μ2 (Mean Difference)...............................	**9 10 11**
σ (Standard Deviation)	**9.2**
ρ1 (Correlation Among Level 1 Units)	**0.1**
ρ2 (Correlation Among Level 2 Units)	**0.05**

Annotated Output

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

Numeric Results

Numeric Results for a Three-Level Hierarchical Mixed Model ——————————————————————

Power	Total N	Group 1 Level 3 Units C1	Group 2 Level 3 Units C2	Level 2 Units K	Level 1 Units M	Mean Diff at End of Follow-Up μ1 - μ2	Std Dev σ	Level 1 Corr ρ1	Level 2 Corr ρ2	Alpha
0.6652	100	2	2	5	5	9.00	9.20	0.100	0.050	0.050
0.9215	200	4	4	5	5	9.00	9.20	0.100	0.050	0.050
0.9851	300	6	6	5	5	9.00	9.20	0.100	0.050	0.050
0.7555	100	2	2	5	5	10.00	9.20	0.100	0.050	0.050
0.9633	200	4	4	5	5	10.00	9.20	0.100	0.050	0.050
0.9958	300	6	6	5	5	10.00	9.20	0.100	0.050	0.050

(report continues)

PASS Sample Size Software   **NCSS.com**

Mixed Models Tests for Two Means at the End of Follow-Up in a 3-Level Hier. Design (Level-3 Rand.)

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.
N is the total number of level-1 units.
C1 and C2 are the number of level-3 units assigned to groups 1 and 2, respectively.
K is the average number of level-2 units per level-3 unit.
M is the average number of level-1 units per level-2 unit.
μ1 - μ2 is the difference in the two group means at the final measurement time: M - 1.
σ is the standard deviation of the level-1 responses.
ρ1 is the correlation among level-1 units in a particular level-2 unit.
ρ2 is the correlation among level-2 units in a particular level-3 unit.
Alpha is the probability of rejecting a true null hypothesis, that is, rejecting when the means are actually equal.

Summary Statements ——
The total number of observations, 100, were obtained from 2 level-3 units in group one and 2
level-3 units in group two with an average of 5 level-2 units per level-3 unit and an average
of 5 level-1 units per level-2 unit. The study achieves 67% power to detect a difference of
9.00 between the (fixed) group means at the last time. The standard deviation of level-1 units
is 9.20. The correlation of level-1 units within a level-2 unit is 0.100. The correlation of
level-2 units within a level-3 unit is 0.050. 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
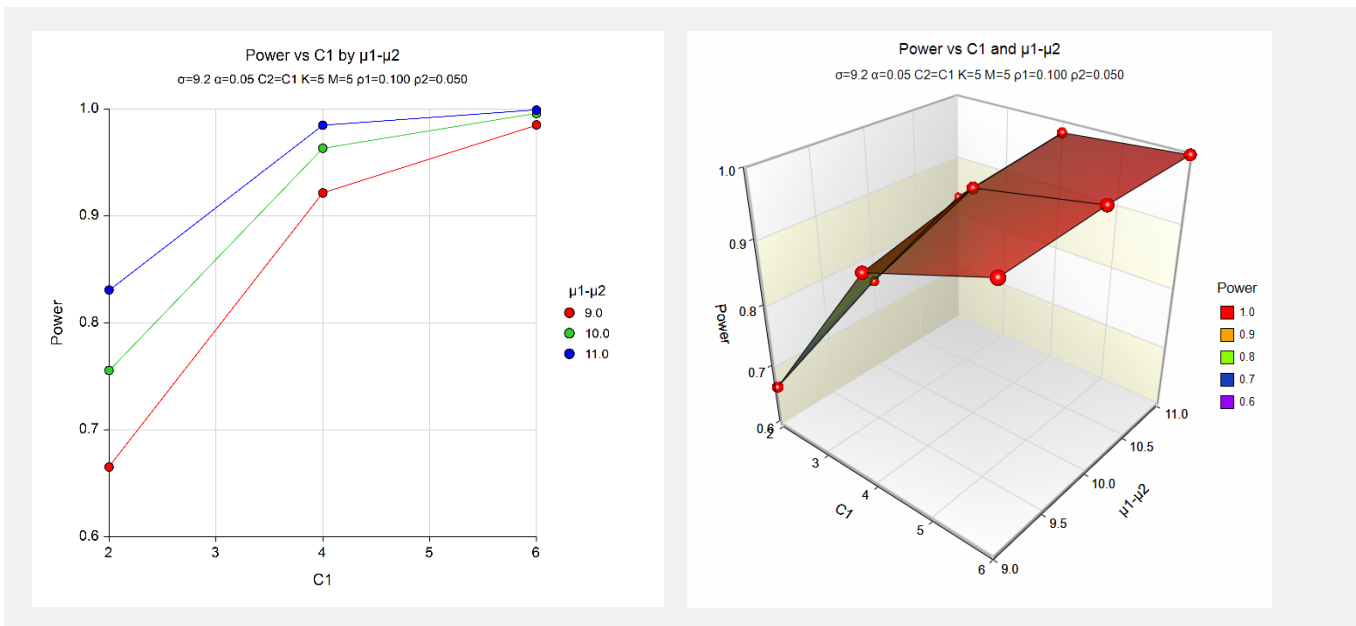
These plots show the power versus the level-3 count for the three values of the difference.

PASS Sample Size Software   **NCSS.com**

Mixed Models Tests for Two Means at the End of Follow-Up in a 3-Level Hier. Design (Level-3 Rand.)

Example 2 – Calculating Sample Size (Number of Clinics)

Continuing with the last example, suppose the researchers want to determine the number of clinics (level-3 units) needed to achieve 90% power for the values of the mean difference.

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 3-Level Hier. Design (Level-3 Rand.)** 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**.

Option	Value
Design Tab	
Solve For	**C1 (Number Level 3 Units)**
Power...	**0.90**
Alpha..	**0.05**
C2 (Level 3 Units Assigned to Group 2)	**C1**
K (Level 2 Units Per Level 3 Unit)	**5**
M (Level 1 Units Per Level 2 Unit)..................	**5**
μ1 - μ2 (Mean Difference)...............................	**9 10 11**
σ (Standard Deviation)	**9.2**
ρ1 (Correlation Among Level 1 Units)	**0.1**
ρ2 (Correlation Among Level 2 Units)	**0.05**

Output

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

Numeric Results

Numeric Results for a Three-Level Hierarchical Mixed Model ──────────────────

Power	Total N	Group 1 Level 3 Units C1	Group 2 Level 3 Units C2	Level 2 Units K	Level 1 Units M	Mean Diff at End of Follow-Up μ1 - μ2	Std Dev σ	Level 1 Corr ρ1	Level 2 Corr ρ2	Alpha
0.9215	200	4	4	5	5	9.00	9.20	0.100	0.050	0.050
0.9011	150	3	3	5	5	10.00	9.20	0.100	0.050	0.050
0.9466	150	3	3	5	5	11.00	9.20	0.100	0.050	0.050

This report shows the power for each of the scenarios.

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

Ahn, Heo, and Zhang (2015) page 211 provide a table in which several scenarios are reported. We will validate this procedure by the first row of the table. The following parameter settings were for the analysis: power = 0.80; $\mu_1 - \mu_2 = 0.3$; $\sigma = 1$; $\rho_1 = 0.1$; $\rho_2 = 0.05$; K = 4; M = 5; and $\alpha = 0.05$. These settings resulted in a value of C1 and C2 (their $N_3^{(0)}$) of 35 and an attained power of 0.806.

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 3-Level Hier. Design (Level-3 Rand.)** 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**.

Option	Value
Design Tab	
Solve For	**C1 (Number Level 3 Units)**
Power...	**0.80**
Alpha..	**0.05**
C2 (Level 3 Units Assigned to Group 2)	**C1**
K (Level 2 Units Per Level 3 Unit)	**4**
M (Level 1 Units Per Level 2 Unit)..................	**5**
µ1 - µ2 (Mean Difference)...............................	**0.3**
σ (Standard Deviation)	**1**
ρ1 (Correlation Among Level 1 Units)	**0.1**
ρ2 (Correlation Among Level 2 Units)	**0.05**

Output

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

Numeric Results

Numeric Results for a Three-Level Hierarchical Mixed Model ─────────────────────

Power	Total N	Group 1 Level 3 Units C1	Group 2 Level 3 Units C2	Level 2 Units K	Level 1 Units M	Mean Diff at End of Follow-Up µ1 - µ2	Std Dev σ	Level 1 Corr ρ1	Level 2 Corr ρ2	Alpha
0.8062	1400	35	35	4	5	0.30	1.00	0.100	0.050	0.050

PASS calculates the same values of C1 and power— 35 and 0.806.