Chapter 565

Randomized Block Analysis of Variance

Introduction

This module analyzes a randomized block analysis of variance with up to two treatment factors and their interaction. It provides tables of power values for various configurations of the randomized block design.

The Randomized Block Design

The randomized block design (*RBD*) may be used when a researcher wants to reduce the experimental error among observations of the same treatment by accounting for the differences among blocks. If three treatments are arranged in two blocks, the *RBD* might appear as follows:

Block A
Treatment 1
Treatment 3
Treatment 2

Block B
Treatment 2
Treatment 1
Treatment 3

This diagram shows the main features of a *RBD*:

1. Each block is divided into *k* sub-blocks, where *k* is the number of treatments.
2. Each block receives all the treatments.
3. The treatments are assigned to the sub-blocks in random order.
4. There is some reason to believe that the blocks are the same internally, but different from each other.

RBD Reduces Random Error

The random error component of a completely randomized design (such as a one-way or a fixed-effects factorial design) represents the influence of all possible variables in the universe on the response except for the controlled (treatment) variables. This random error component is called the standard deviation or σ (sigma).

As we have discussed, the sample size required to meet alpha and beta error requirements depends directly on the standard deviation. As the standard deviation increases, the sample size increases. Hence, researchers are always looking for ways to reduce the standard deviation. Since the random error component contains the variation due to all possible variables other than treatment variables, one of the most obvious ways to reduce the standard deviation is to remove one or more of these *nuisance* variables from the random error component. One of the simplest ways of doing this is by blocking on them.

For example, an agricultural experiment is often blocked on fields so that differences among fields are explicitly accounted for and removed from the error component. Since these field differences are caused by variations in variables such as soil type, sunlight, temperature, and water, blocking on fields removes the influence of several variables.

Blocks are constructed so that the response is as alike (homogeneous) as possible within a block, but as different as possible between blocks. In many situations, there are obvious natural blocking factors such as schools, seasons, individual farms, families, times of day, etc. In other situations, the blocks may be somewhat artificially constructed.

Once the blocks are defined, they are divided into *k* smaller sections called *subblocks*, where *k* is the number of treatment levels. The *k* treatments are randomly assigned to the subblocks, one block at a time. Hence the order of treatment application will be different from block to block.

Measurement of Random Error

The measurement of the random error component (σ) is based on the assumption that there is no fundamental relationship between the treatment variable and the blocking variable. When this is true, the interaction component between blocks and treatment is zero. If the interaction component is zero, then the amount measured by the interaction is actually random error and can be used as an estimate of σ.

Hence, the randomized block design makes the assumption that there is no interaction between treatments and blocks. The block by treatment mean square is still calculated, but it is used as the estimated standard deviation. This means that the degrees of freedom associated with the block-treatment interaction are the degrees of freedom of the error estimate. If the experimental design has *k* treatments and *b* blocks, the interaction degrees of freedom are equal to $(k-1)(b-1)$. Hence the sample size of this type of experiment is measured in terms of the number of blocks.

Treatment Effects

Either one or two treatment variables may be specified. If two are used, their interaction may also be measured. The null hypothesis in the *F* test states that the effects of the treatment variable are zero. The magnitude of the alternative hypothesis is represented as the size of the standard deviation (σ_m) of these effects. The larger the size of the effects, the larger their standard deviation.

When there are two factors, the block-treatment interaction may be partitioned just as the treatment may be partitioned. For example, if we let *C* and *D* represent two treatments, an analysis of variance will include the terms *C*, *D*, and *CD*. If we represent the blocking factor as *B*, there will be three interactions with blocks: *BC*, *BD*, and *BCD*. Since all three of these terms are assumed to measure the random error, the overall estimate of random error is found by averaging (or *pooling*) these three interactions. The pooling of these interactions increases the power of the experiment by effectively increasing the sample size on which the estimate of σ is based. However, it is based on the assumption that $\sigma = \sigma_{BC} = \sigma_{BD} = \sigma_{BCD}$, which may or may not be true.

An Example

Following is an example of data from a randomized block design. The block factor has four blocks (*B1*, *B2*, *B3*, *B4*) while the treatment factor has three levels (low, medium, and high). The response is shown within the table.

Randomized Block Example			
	Treatments		
Blocks	**Low**	**Medium**	**High**
B1	16	19	20
B2	18	20	21
B3	15	17	22
B4	14	17	19

Analysis of Variance Hypotheses

The *F* test for treatments in a randomized block design tests the hypothesis that the treatment effects are zero. (See the beginning of the Fixed-Effects Analysis of Variance chapter for a discussion of the meaning of effects.)

Single-Factor Repeated Measures Designs

The randomized block design is often confused with a single-factor repeated measures design because the analysis of each is similar. However, the randomization pattern is different. In a randomized block design, the treatments are applied in random order within each block. In a repeated measures design, however, the treatments are usually applied in the same order through time. You should not mix the two. If you are analyzing a repeated measures design, we suggest that you use that module of ***PASS*** to do the sample size and power calculations.

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 whether you want to solve for power or number of blocks. This choice controls which options are displayed in the Sample Size section.

If you select number of blocks, another option appears.

Solve For Sample Size 'Based On Term'

Specify which term to use when searching for the number of blocks. The power of this term's test will be evaluated as the search is conducted.

- **All**

  If you want to use all terms, select 'All.' Note that only terms that are active may be selected. If you select a term that is not active, you will be prompted to select another term.

Power and Alpha

Minimum Power

Enter a value for the minimum power to be achieved when solving for the sample size. The resulting sample size is large enough so that the power is greater than this amount. This search is conducted using either all terms or a specific term.

Power is the probability of rejecting the null hypothesis when it is false. Beta, the probability of obtaining a false negative on the statistical test, is equal to 1-power, so specifying power implicitly specifies beta.

Range

Since power is a probability, its valid range is between 0 and 1.

Recommended

Different disciplines have different standards for power. A popular value is 0.9, but 0.8 is also popular.

Alpha for All Terms

Enter a single value of alpha to be used in all statistical tests.

Alpha is the probability of obtaining a false positive on a statistical test. That is, it is the probability of rejecting a true null hypothesis. The null hypothesis is usually that the variance of the means (or effects is) zero.

Range

Since Alpha is a probability, it is bounded by 0 and 1. Commonly, it is between 0.001 and 0.250.

Recommended

Alpha is usually set to 0.05 for two-sided tests such as those considered here.

Design and Effects

Number of Factors

Select the number of factors included in the design. Up to three factors may be selected.

Main Effects

Levels (A and B)

These options specify the number of levels (categories) contained in each active factor. A factor must have at least 2 levels. We have arbitrarily set the maximum at 100, but realistically, the maximum is about 10.

Since the total sample size is equal to the product of the number of levels in each factor and the number of observations per group (n), increasing the number of levels of a factor increases the total sample size of the experiment.

Means

The standard deviation of the means is calculated using the formula

$$\sigma_m = \sqrt{\sum_{i=1}^{k} \frac{\left(\mu_i - \bar{\mu}\right)^2}{k}}$$

where k is the number of means. This option lets you select the method used to enter σ_m: either σ_m directly or as a list of means.

As a Std Dev

Enter value of σ_m in the box to the right. The value entered represents the magnitude of the variation among the means that you want to detect.

You can use the Standard Deviation Estimator window to calculate a value of σ_m for various sets of means. This window is obtained by pressing the 'Sm' icon to the right.

List of Means

Enter a list of means in the box to the right. The standard deviation of these means is calculated using the formula

$$\sigma_m = \sqrt{\sum_{i=1}^{k} \frac{\left(\mu_i - \bar{\mu}\right)^2}{k}}$$

where k is the number of means. Note that the standard deviation will be the same whether you enter means or effects.

Enter a set of means that give the pattern of differences you expect or the pattern that you wish to detect. For example, in a particular study involving a factor with three categories, your research might be meaningful if either of two treatment means is 50% larger than the control mean. If the control mean is 50, then you would enter *50,75,75* as the three means.

Note that it is not the values of the means themselves that is important, but their differences. Thus, the mean values *0,1,2* produce the same results as the values *100,101,102*.

If not enough means are entered to match the number of groups, the last mean is repeated. For example, suppose that four means are needed and you enter *1,2* (only two means). ***PASS*** will treat this as *1,2,2,2*. If too many values are entered, ***PASS*** will truncate the list to the number of means needed.

Interaction

A*B

This check box indicates whether the A*B interaction is included in the analysis of variance model.

Effects

Select the method used to enter the effects.

Std Dev of Effects

Enter the value of σ_m directly. You can use the 'Sm' tool to help you calculate an appropriate value.

Multiple of Another Term

Enter the value of σ_m as a multiple of another term. You will enter a multiplier and select a term.

Multiplier

Enter the value of the multiplier. Usually, this will be a value near one. If you enter '1,' the σ_m value of the other term will be used for this term.

Discussion

The general formula for the calculation of the standard deviation of the interaction effects is

$$\sigma_m = \sqrt{\frac{\sum_{i=1}^{k} (e_i)^2}{k}}$$

where *k* is the number of effects. In the case of a two-way interaction, the standard deviation is calculated using the formula:

$$\sigma_m(AB) = \sqrt{\frac{\sum_{i=1}^{I}\sum_{j=1}^{J}\left(\mu_{ij} - \mu_{i\bullet} - \mu_{\bullet j} + \bar{\mu}\right)^2}{IJ}}$$

where *i* is the factor *A* index (from 1 to *I*), *j* is the factor *B* index (from 1 to *J*), μ_{ij} is the mean in the ij^{th} *cell,* $\mu_{i\bullet}$ is the i^{th} mean of factor *A* across all levels of other factors, $\mu_{\bullet j}$ is the j^{th} mean when factor *B* across all levels of other factors, and $\bar{\mu}$ is the overall mean of the means.

To see how this works, consider the following table of means from an experiment with *I* = 2 and *J* = 3:

		i		
		1	2	
	1	2.0	4.0	| 3.0
j	2	4.0	6.0	| 5.0
	3	6.0	11.0	| 8.5
	---	---	---	| ---
	Total	4.0	7.0	| 5.5

Now, if we subtract the factor *A* means, subtract the factor *B* means, and add the overall mean, we get the interaction effects:

0.5	-0.5
0.5	-0.5
-1.0	1.0

Next, we sum the squares of these six values:

$$(0.5)^2 + (-0.5)^2 + (0.5)^2 + (-0.5)^2 + (-1.0)^2 + (1.0)^2 = 3$$

Next we divide this value by (2)(3) = 6:

$$3/6 = 0.5$$

Finally, we take the square root of this value:

$$\sqrt{0.5} = 0.7071$$

Hence, for this configuration of means,

$$\sigma_m(AB) = 0.7071.$$

Notice that the average of the absolute values of the interaction effects is:

$$[0.5 + 0.5 + 0.5 + 0.5 + 1.0 + 1.0]/6 = 0.6667$$

Note that SD(interaction) is close to the average absolute interaction effect. That is, 0.7071 is close to 0.6667. This will usually be the case. Hence, one way to interpret the interaction-effects standard deviation is as a number a little larger than the average absolute interaction effect.

Standard Deviation

σ (Block-Treament Interaction)

Enter an estimate of the square-root of the pooled block-by-treatment interaction. In a randomized block design in which the block-by-treatment interaction is not included in the model, this is the square root of the mean square error. This value will usually have to be determined from a previous study.

Assuming that each block is divided into several sub-blocks, this is an estimate of the sub-block standard deviation.

Note that ONLY ONE VALUE may be entered.

Sample Size

Number of Blocks

This specifies one or more values for the number of blocks. If a list is entered, a separate calculation is made for each value in the list.

Example 1 – Power after a Study

This example will explain how to calculate the power of *F* tests from data that have already been collected and analyzed. We will analyze the power of the experiment that was given at the beginning of this chapter.These data were analyzed using the analysis of variance procedure in NCSS and the following results were obtained.

Analysis of Variance Table

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
A (Blocks)	3	13.66667	4.555555			
B (Treatment)	2	45.16667	22.58333	19.83	0.002269*	0.991442
AB	6	6.833333	1.138889			
S	0	0				
Total (Adjusted)	11	65.66666				
Total	12					

* Term significant at alpha = 0.05

Means and Effects Section

Term	Count	Mean	Standard Error	Effect
B: Treatment				
High	4	20.5	0.5335937	2.333333
Low	4	15.75	0.5335937	-2.416667
Medium	4	18.25	0.5335937	8.333334E-02

We will now calculate the power of the *F* test. Note that factor *B* in this printout becomes factor *A* on the PASS template.

Setup

To analyze these data, we enter the means for factor *A*. The value of σ is estimated as the square root of the mean square error:

$$\sigma = \sqrt{1.138889} = 1.0672$$

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Randomized Block Analysis of Variance** procedure window by expanding **Means**, then clicking on **Multi-Factor Designs (ANOVA)**, and then clicking on **Randomized Block Analysis of Variance**. 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 for All Terms	**0.05**
Number of Factors	**1**
A - Levels	**3**
A - Means	**List of Means**
A - List of Means	**15.75 18.25 20.50**
σ (Block-Treatment Interaction)	**1.0672**
Number of Blocks	**2 3 4 5**

Annotated Output

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

Numeric Results

Term	Power	Number of Blocks	Units	df1	df2	Standard Deviation of Means (σm)	Block-Treatment Interaction (σ)	Effect Size (σm/σ)	Alpha	Beta
A	0.4213	2	6	2	2	1.940	1.067	1.818	0.05	0.58
A	0.8938	3	9	2	4	1.940	1.067	1.818	0.05	0.11
A	0.9914	4	12	2	6	1.940	1.067	1.818	0.05	0.01
A	0.9996	5	15	2	8	1.940	1.067	1.818	0.05	0.00

References

Odeh, R.E. and Fox, M. 1991. Sample Size Choice. Marcel Dekker, Inc. New York, NY.
Winer, B.J. 1991. Statistical Principles in Experimental Design. Third Edition. McGraw-Hill. New York, NY.

Report Definitions

Power is the probability of rejecting a false null hypothesis.
Blocks are the number of blocks in the design.
Units are the number of experimental units in the design.
df1 is the numerator degrees of freedom.
df2 is the denominator degrees of freedom.
σm is the standard deviation of the group means or effects.
σ is the pooled block-treatment interaction.
Effect Size is the ratio σm/σ
Alpha is the probability of rejecting a true null hypothesis.
Beta is the probability of accepting a false null hypothesis.

Summary Statements

A randomized-block design with one treatment factor at 3 levels has 2.0 blocks each with 3.0 treatment combinations. The square root of the block-treatment interaction is 1.067. This design achieves 42% power when an F test is used to test factor A at a 5% significance level and the actual standard deviation among the appropriate means is 1.940 (an effect size of 1.818).

This report shows the power for each of the four block counts. We see that adequate power of about 0.9 would have been achieved by three blocks.

It is important to emphasize that these power values are for the case when the effects associated with the alternative hypotheses are equal to those given by the data. It will often be informative to calculate the power for other values as well.

Term

This is the term (main effect or interaction) from the analysis of variance model being displayed on this line.

Power

This is the power of the *F* test for this term. Note that since adding and removing terms changes the denominator degrees of freedom (*df2*), the power depends on which other terms are included in the model.

Blocks

This is the number of blocks in the design.

Units

This is the number of sub-blocks (plots) in the design. It is the product of the number of treatment levels and the number of blocks.

df1

This is the numerator degrees of freedom of the *F* test.

df2

This is the denominator degrees of freedom of the *F* test. This value depends on which terms are included in the AOV model.

Std Dev of Means (σm)

This is the standard deviation of the means (or effects). It represents the size of the differences among the effects that is to be detected by the analysis. If you have entered hypothesized means, only their standard deviation is displayed here.

Block-Treatment Interaction

This is the pooled (averaged) block-treatment interaction mean square.

Effect Size

This is the standard deviation of the means divided by the block-treatment interaction. It provides an index of the magnitude of the difference among the means that can be detected by this design.

Alpha

This is the significance level of the *F* test. This is the probability of a type-I error given the null hypothesis of equal means and zero effects.

Beta

This is the probability of the type-II error for this test given the sample size, significance level, and effect size.

Example 2 – Validation using Prihoda

Prihoda (1983) presents details of an example that is given in Odeh and Fox (1991). In this example, *Alpha* is 0.025, *Sm* of *A* is 0.577, the number of treatments in factor *A* is 6, the number of treatments in factor *B* is 3, S is 1.0, and the Number of Blocks is 2, 3, 4, 5, 6, 7, and 8. Prihoda gives the power values for the *F* test on factor *A* as 0.477, 0.797, 0.935, 0.982, 0.995, 0.999, and 1.000.

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Randomized Block Analysis of Variance** procedure window by expanding **Means**, then clicking on **Multi-Factor Designs (ANOVA)**, and then clicking on **Randomized Block Analysis of Variance**. 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**.

Design Tab

Solve For .... **Power**
Alpha for All Terms **0.025**
Number of Factors **2**
A - Levels **6**
A - Means .... **As a Std Dev**
A - Std Dev **0.577**
B - Levels **3**
B - Means .... **As a Std Dev**
B - Std Dev **1**
A*B .... **Checked**
A*B - Effects .... **Std Dev of Effects**
A*B - Std Dev **1**
σ (Block-Treatment Interaction) **1.0**
Number of Blocks .... **2 to 8 by 1**

Output

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

Numeric Results

Term	Power	Number of Blocks	Units	df1	df2	Standard Deviation of Means (σm)	Block-Treatment Interaction (σ)	Effect Size (σm/σ)	Alpha	Beta
A	**0.4762**	2	36	5	17	0.577	1.000	0.577	0.03	0.52
B	0.9970	2	36	2	17	1.000	1.000	1.000	0.03	0.00
A*B	0.8534	2	36	10	17	1.000	1.000	1.000	0.03	0.15
A	**0.7952**	3	54	5	34	0.577	1.000	0.577	0.03	0.20
B	1.0000	3	54	2	34	1.000	1.000	1.000	0.03	0.00
A*B	0.9961	3	54	10	34	1.000	1.000	1.000	0.03	0.00
A	**0.9348**	4	72	5	51	0.577	1.000	0.577	0.03	0.07
B	1.0000	4	72	2	51	1.000	1.000	1.000	0.03	0.00
A*B	1.0000	4	72	10	51	1.000	1.000	1.000	0.03	0.00
A	**0.9823**	5	90	5	68	0.577	1.000	0.577	0.03	0.02
B	1.0000	5	90	2	68	1.000	1.000	1.000	0.03	0.00
A*B	1.0000	5	90	10	68	1.000	1.000	1.000	0.03	0.00
A	**0.9957**	6	108	5	85	0.577	1.000	0.577	0.03	0.00
B	1.0000	6	108	2	85	1.000	1.000	1.000	0.03	0.00
A*B	1.0000	6	108	10	85	1.000	1.000	1.000	0.03	0.00
A	**0.9991**	7	126	5	102	0.577	1.000	0.577	0.03	0.00
B	1.0000	7	126	2	102	1.000	1.000	1.000	0.03	0.00
A*B	1.0000	7	126	10	102	1.000	1.000	1.000	0.03	0.00
A	**0.9998**	8	144	5	119	0.577	1.000	0.577	0.03	0.00
B	1.0000	8	144	2	119	1.000	1.000	1.000	0.03	0.00
A*B	1.0000	8	144	10	119	1.000	1.000	1.000	0.03	0.00

We have bolded the power values on this report that should match Prihoda's results. You see that they do match.