Statistics 2.

RCBD Review

2014. Prepared by Lauren Pincus

With input from Mark Bell and Richard Plant

Agriculture Innovation Program

Table of Contents

Questions for review

Assume you are doing a fertilizer experiment. You can either test different fertilizer TYPES or the RATE at which they are applied.

Answer the following:

1. What type of analysis is best to form a crop response curve to fertilizer application rates of 50, 100, 150, and 200 kg N ha^{-1}?

a) ANOVA
b) Linear regression
c) T-test
d) Factorial ANOVA

2. What type of analysis would you use if you tested crop response to four different types of fertilizer, such as urea, CAN, ammonium sulfate, and ammonium nitrate?

a) ANOVA
b) Linear regression
c) T-test
d) Factorial ANOVA

3. What type of analysis would you use if you tested crop response to two different fertilizer types at four different rates?

a) ANOVA
b) Linear regression
c) T-test
d) Factorial ANOVA

Answers

1. When you have continuous, or measureable, data such as a fertilizer rate it is better to use a regression to analyze your results. The rate of fertilizer is referred to as a continuous variable because there are many possible rates in addition to the ones the researcher selected to use in the experiment. If you use an ANOVA, you will be able to determine if one or more treatments is different than the others, but you will not be able to look at the full spectrum of possible fertilizer rates. What if 130 kg N ha^{-1} actually produces the highest yields?
2. If you are testing four different fertilizer TYPES, then you have four different categories, or groups, you are comparing and you need to use an ANOVA. If you only are testing two different fertilizers then you could use a t-test. Never use an ANOVA to analyze continuous variables.
3. In this instance you would use a factorial 2 x 4 ANOVA. Your factorial has two factors, “type” and “rate.” Type has 2 different levels and Rate has 4 levels. It is now possible to look at the interaction between “type” and “rate.” What if your crop has much higher yields at low rates of urea than CAN? A factorial analysis allows the researcher to test both the treatment main effects and their interaction. Fertilizer rate main effects can also be subjected to polynomial contrasts, a statistical method to determine if there are linear or quadratic components in the overall response.

Materials to Review

1. Randomized Complete Block Designs (RCBD)
2. RCBD across locations
3. RCBD with subsamples

1. Randomized Complete Block Designs (RCBD)

An RCBD is used to make sure treatments are compared under similar circumstances. For example, imagine the natural fertility of a field varies from one end to the other. If you are conducting a field experiment to compare different types of fertilizer, then clearly the natural fertility could affect your conclusions. For example, if you used a completely randomize design, then there is some chance that one fertilizer treatment may end up with more replicates on the more fertile part of the field. As a result, you might conclude that this fertilizer treatment was better than the others. In such a situation, it is said that the effects of the fertilizer was “confounded” with the natural fertility of the land (i.e., the natural soil fertility is interfering with the conclusions about the fertilizer effects).

As a result if there is known (or suspected) field variability, you want to “block” your field to remove any advantage or disadvantage one treatment might have because of its position in the field.

In a RCBD, a set of experimental treatments are grouped (blocked) in a way that minimizes the variability within each block. You then have other blocks and each contains a complete set of treatments. As a result, you reduce variability within a replicate. Meanwhile, you will actually increase variability between replicates which can be estimated as a separate source of variation. The removal of this source of within replicate variation *reduces experimental error* and improves the ability of the experiment to detect smaller treatment differences. The greater the variability among blocks, the more efficient the design becomes.

Note: If there are not large differences between blocks, then you could use a completely randomized design (CRD) which would actually be more efficient. This is because e CRD has more degrees of freedom for error and thus you have a smaller F value to detect significant differences among treatments.

In an RCBD:

- The field is divided into blocks (or replicates)
- Each block is then divided into a number of units (or plots) equal to the number of treatments (i.e., All treatments are observed within each block).
- Within each block, the treatments are assigned *at random* so that a different treatment is applied to each unit.

The defining feature of the Randomized (Complete) Block Design[1] is that each block sees each treatment at least once.

The ANOVA table contains two F tests: our main interest is to test the equality of treatment means, however an RCBD also tests for a significant block effect.

Source of variance	Degrees of Freedom	Sum of Squares (SS)	Mean square (MS)	Observed F
Total	tb-1	Total SS		
Treatment (T)	t-1	SS Treatment	MS Treatment	MST/MSE
Blocks (B)	b-1	SS Blocks	MS Blocks	MSB/MSE
Residual or experimental error	(t-1)(b-1)	SS Error	MS Error	

The ANOVA table for the RCBD has an additional row (Block) relative to that for a completely randomized design (CRD). The consequence of this is that there are FEWER degrees of freedom for error in the RCBD. In the RCBD, these degrees of freedom are said to have been "partitioned" from the error and assigned to the blocks.

Advantages of the RCBD

1. Generally more precise than the CRD.
2. No restriction on the number of treatments or replicates.
3. Some treatments may be replicated more times than others.
4. Missing plots are easily estimated.
5. Whole treatments or entire replicates may be deleted from the analysis.

[1] (For those interested in the greater statistics detail; The analysis assumes that there is no interaction between block and treatment, that is, it fits the model

$Y_{ijk} = \mu + \alpha_i + \beta_j + \varepsilon_{ijk}$

where the αs are the treatment effects and the βs are the block effects.

RCBD Model Assumptions

- Errors are independent
- Errors have homogeneous variance
- Errors have a normal distribution
- Means are additive—there is no interaction between block and treatment means. Each block affects all treatments equally.

Model Diagnostics

- Check that residuals (difference between measured value and the estimated mean value):
- Have approximately a normal distribution
- Do not have a pattern (showing unequal variance) across *blocks*
- Do not have a pattern (showing unequal variance) across *treatments*

Disadvantages of the RCBD

1. Error df is smaller than that for the CRD (somewhat of a problem if there are a small number of treatments).
2. If there is a large variation between experimental units within a block, a large error term may result (this may be due to too many treatments).
3. If there are missing data, a RCBD experiment may be less efficient than a CRD.
4. If there is no significant difference between blocks (blocks are not really different from each other) then a CRD is more powerful than an RCBD.

Sample RCBD layout

In this field there is a slope that causes a fertility gradient. The field is more fertile at the bottom than at the top. Different colors represent different treatments; each horizontal row represents a block. There are 4 blocks [I-IV] and 4 treatments [A (red); B (blue); C (yellow); D (green)] in this example.

Block	**Treatment**			
I	A	B	C	D
II	D	A	B	C
III	B	D	C	A
IV	C	A	B	D

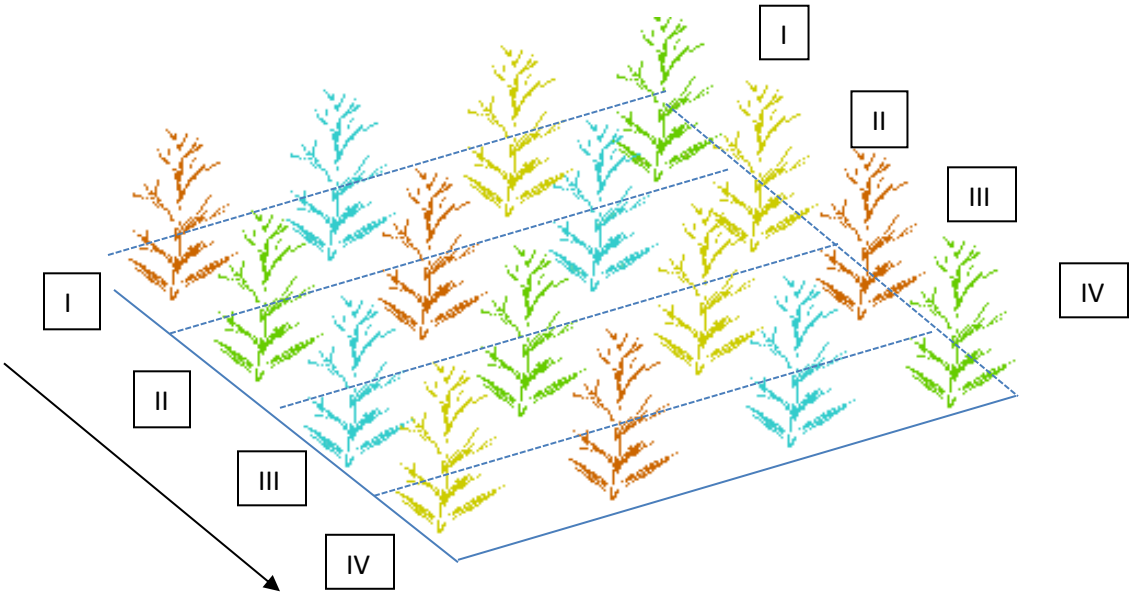

In this case, the df for the ANOVA would look like:

Source of variance	**Degrees of Freedom**
Total	tb-1 = 4 x 4 -1 = 15
Treatment (T)	t-1 = 4 -1 = 3
Blocks (B)	b-1 = 4 -1 = 3

Residual or experimental error	(t-1)(b-1) = 3 x 3 = 9

Why is this an appropriate layout?

Answer: Variation **within** a replicate is minimized – while variation **between** replicates is maximized.

A note on Degrees of Freedom for Error.

The figure below shows the general decrease in the F value (used to determine significance) versus the degrees of freedom for error (Edf). This graph helps us understand where we want to target for the minimum degrees of freedom for error.

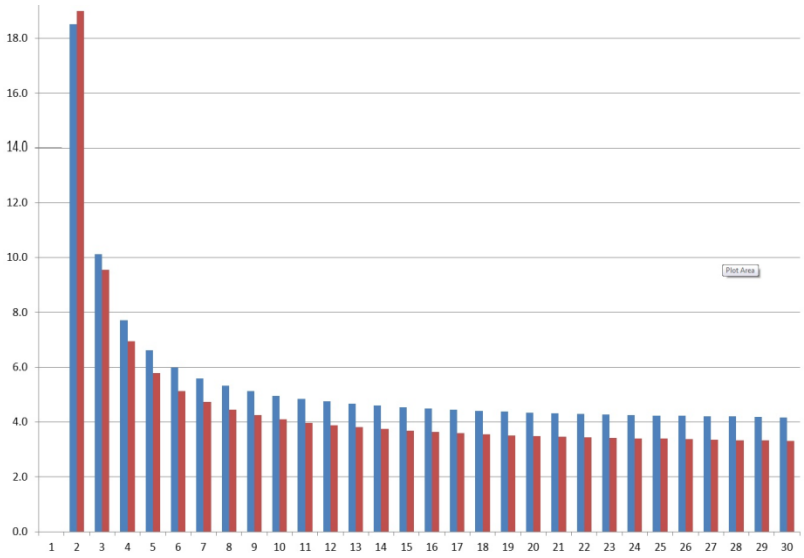

The value of F versus degrees of freedom for error.

This shows that around 12-20 df for error, F is beginning to flatten out. Thus, a rough rule of thumb is that in order to have the best opportunity to detect differences, you should have at least 12 df for error.

Test your understanding

Question 1. You are asked to do an experiment on a field that you are unfamiliar with and it is unclear if there is natural variability in soil fertility. The field is long and narrow. Is it better to a) create blocks that run the whole length of the field or b) cut the field into short blocks?

a) 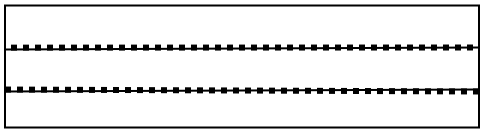

b) 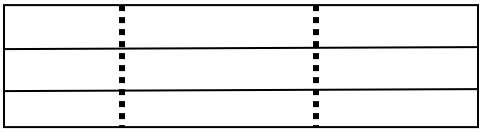

Answer. While there is no correct answer (since we don't know actual variability), it is more likely that design b) would be better.

Question 2: In your plot, fertility is higher at the north end of the plot. Which is a better layout for your blocks? a) or b) ?

a) 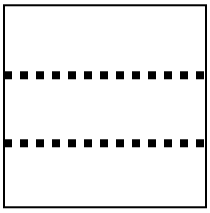

North 

b) 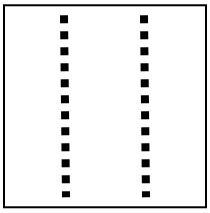

Answer. When blocking, each block should minimize within block variability and maximize differences between the other blocks. a) is correct.

Question 3. You set up your analysis as an RCBD by designating blocks and then randomizing your treatments so that each block contains each treatment. When you run your analysis you see that your blocks are not significant. Is it possible to reanalyze this experiment as a completely randomized design since blocks do not explain a significant portion of the variance?

Answer. NO. Because you have already randomized your treatments within blocks, it is not possible to change the analysis at this point. A CRD requires that each replicate of a treatment is randomly assigned to an experimental unit (a plot within the larger main field in this case), whereas a RCBD requires that treatments are randomized within each block individually. Multiple randomization events occur in an RCBD, but only one randomization occurs in a CRD. It is correct to report that the blocking factor is non-significant, but it is incorrect to change your analysis to a CRD if you have randomized according to an RCBD.

Question 4. Suppose you have 50 bean varieties to compare. Is it appropriate to use an RCBD to evaluate differences between varieties?

> Answer. It might be okay to use an RCBD. An RCBD design needs at least two distinct "block" areas that are different from each other, but have very little within-block variability. It might be hard to find a field large enough to have two blocks where all 50 varieties can be planted such that within variability is low and less likely that there are TWO plots that fit this criteria. For this reason many variety trials use an incomplete block design where each block does not include every treatment. Incomplete block designs are analyzed differently than randomized complete block designs.

Question 5. When you create blocks, is the objective of "blocking" to evaluate the difference between blocks?

> Answer: No. The objective of blocking is not to compare Block 1 to Block 2. The objective is to remove variability associated with the environment to more accurately assess differences between treatments. Blocks are part of the experimental design structure; they are NOT an experimental factor themselves. If you are interested in how a particular variety performs on different soil types, then "soil type" becomes a factor in the experiment and not a block.

More Advanced Questions

Question 6. You are doing a small variety trial and have randomized treatments according to an RCBD. Your blocks run east to west along the field. Each block is irrigated by the same sprinkler system. During a trip to the field you discover that the irrigation system was clogged for one row and the plants have died. Are you still able to analyze your results?

> Answer. Yes. In this example the RCBD is a single factor design and you have lost one replicate. You have an unbalanced design due to the missing data, but can still analyze this experiment. This would not be possible if your design was a split-plot or a factorial design because you would lose information necessary to create the interaction term.

Question 7. In your experiment, you decide to investigate the effect of estrogen on weight gain in sheep. You pick four ranches and at each ranch you inject two males and two females with two levels of estrogen (Est_0, Est_3). Your treatments are combinations of sex (M, F) and estrogen (Est_0, Est_3). The sheep are blocked by ranch, with one replication of each treatment level at each ranch. Can you include a treatment*block interaction in this model?

	Ranch			
Trtmt	**1**	**2**	**3**	**4**
M Est_0				
M Est_3				
F Est_0				
F Est_3				

Answer. NO. With only one replication per block, there are zero degrees of freedom for error. Remember that experimental error is defined as the variation among experimental units *that are treated alike*. The assumption of no interaction is referred to as the assumption of additivity of the main effects. If this assumption is violated, all F-tests will be inefficient and possibly misleading, particularly if the interaction effect is very large. To test for non-additive effects, also referred to as multiplicative effects, perform a Tukey's test for non-additivity. This test is only necessary when there is one observation per block*treatment combination. If there are two or more replications per block-treatment combination, the interaction can be tested directly in an exploratory model. A plot of the residuals versus the predicted values will show a curvilinear relationship rather than a random scatter if there is non-additivity between the block and treatment effects.

2. RCBD Repeated Across Locations

What if the RCBD is repeated across locations? In the example below there are three treatments per block, three blocks per location and three locations. Treatments are still assigned at random to each block, but now the number of blocks is the number of replications.

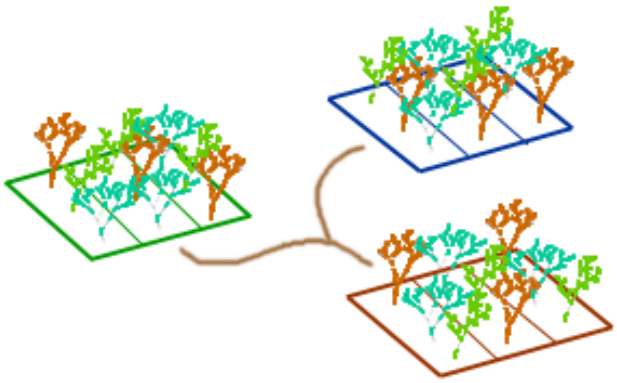

Location 1			Location 2			Location 3		
Block			Block			Block		
I	II	III	I	II	III	I	II	III
A	B	C	B	C	B	A	C	A
B	A	B	A	B	C	C	B	C
C	C	A	C	A	A	B	A	B

ANOVA for RCBD repeated across locations.

Source of variance	Degrees of Freedom	Sum of Squares (SS)	Mean square (MS)	Observed F
Locations (L)	L-1	SS Location	MSL = SSL/(L-1)	MSL/MSE
Error for Locations (E_L)	L*(b-1)	SS $Error_{Location}$	SSE_L/(1*(b-1))	
Treatment	t-1	SSTreatment	MST = SST/(t-1)	MST/MSE
Treatment* Location	(t-1)*(L-1)	$SS_{Treatment\ x\ Location}$	MS_{TxL}= SS_{TxL}/((t-1)*(L-1))	MS_{TxL}/MSE
Error (E)	L*(t-1)*(b-1)	SSError	MES = SSE/(L*(t-1)*(b-1))	
Total	L*t*b - 1	Total SS		

In our example, the df in the ANOVA are:

Source of variance	Degrees of Freedom	Degrees of Freedom
Locations (L)	L-1	3-1 = 2
Error for Locations (E_L)	L*(b-1)	3(3–1) = 6
Treatment	t-1	3-1 = 2
Treatment* Location	(t-1)*(L-1)	2*2 = 4
Error (E)	L*(t-1)*(b-1)	3*2*2 = 12
Total	L*t*b - 1	3*3*3 - 1 = 26

Test your understanding

Question 1. How many degrees of freedom does an experiment with 4 treatments and 4 blocks, repeated in 3 regions of the country have?

Answers

1. 3*4*4 – 1 = 47 df

3. RCBD with Subsamples

It is frequently the case that researchers want to collect multiple samples from within the same treatment. These measurements should be included in your analysis as subsamples, not replications. It is also important that these samples are independent from one another—for example, the height of two randomly sampled and independent plants is taken, NOT the height of the same plant twice during the growing season. If these measurements are not independent, then your design is NESTED.

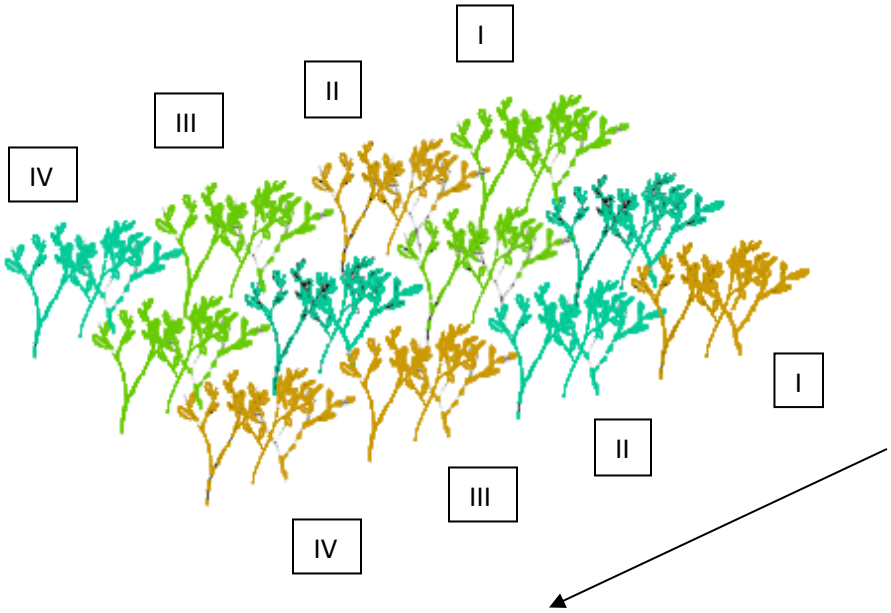

In this picture there are four blocks (I-IV) going down the slope, three treatments, and three subsamples (individual plants).

Source of variance	Degrees of Freedom	Sum of Squares (SS)	Mean square (MS)	Observed F
Blocks (B)	b-1	SSB	SSB/(b-1)	MSB/MSE
Treatment (T)	t-1	SST	SST/(t-1)	MST/MSE
Experimental Error (E)	(t-1)*(b-1)	SSE	SSE/((t-1)*(b-1))	MSE/MSS
Sampling Error (S)	t*b*(s-1)	SSS	SSS/(t*b*(s-1))	
Total	t*b*s-1	TSS		

3. RCBD may not always be the design to use

It is not always appropriate to conduct a RCBD when you have a factor you suspect has an effect on the response, but is not something you are interested in.

Nuisance Characteristic	Examples	How to treat
Known, moderately controllable (by choosing to block rather than adjust factor)	Temperature, location, time, batch, particular machine or operator	Blocking
Known, uncontrollable, measurable	Weight, gender, previous exposure	Analysis of Covariance
Unknown, uncontrollable	Experimenter bias, effect of order of treatments	Randomization

Test your understanding

You design a greenhouse pot experiment evaluating nitrogen uptake from three different rice varieties. You designate three different tables in the greenhouse as your blocks and plant four different seeds in each pot. Each pot will receive a different nitrogen source (urea, CAN, ammonium nitrate) as a fertilizer treatment.

Question 1. You decide that you will measure the foliar N content from two different plants per pot, but you will average them to create one N value per pot. What sort of design do you have?

a) RCBD with repeated measures
b) RCBD with two replicates per pot
c) RCBD with one replicate per pot

Answer 1. c) Because you averaged the values from each pot together, you only have one value per block*treatment combination.

Question 2. You change your mind and decide to measure N content from two different plants, but not average them. What sort of design do you have and is there any advantage to this method of data collection?

a) RCBD with repeated measures
b) RCBD with two replicates per pot
c) RCBD with one replicate per pot

Answer 2. b) You now have two values for each block*treatment combination and it is possible to have a block*treatment interaction*. If the interaction is significant, transform the data. If the interaction is not significant, then exclude it from the model and repeat the analysis.

Question 3. Say you want to evaluate changes in N uptake over time. You decide to sample one plant per pot five times over the course of the season. What sort of design do you have?

a) RCBD with repeated measures
b) RCBD with two replicates per pot
c) RCBD with one replicate per pot

3. a) This is a nested design because the measurement is coming from the same plant. The measures are therefore NOT independent.