The Rothamsted School - its approach to analysis of experiments.

Research Staff Training Session - 2019

by Curt Lee

AGRO-TECH RESEARCH FARM

1

Rothamsted School

- Leading statisticians such as Fisher, Yates, Nelder, Bailey
- Developed analysis of variance not in terms of linear models but in terms of symmetry
- High point was John Nelder's theory of general balance (1965)

General Balance

1) Establish and define block structure
2) Establish and define treatment structure
3) Given randomisation the analysis then follows automatically

(c) Stephen Senn

2

Rothamsted is the longest running agricultural research institute in the world and has a long and rich history of famous statisticians.

3

ASA-CSSA-SSSA presentation

https://a-c-s.confex.com/crops/2007am/techprogram/P30406.HTM

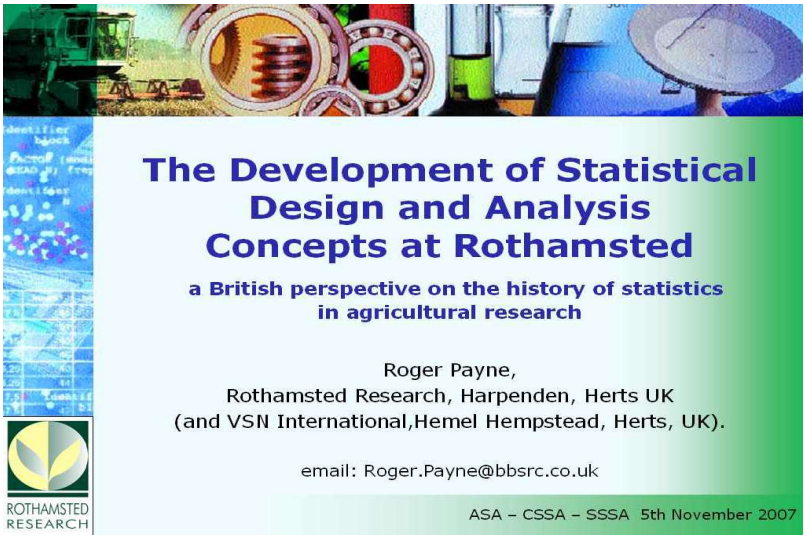

4

The ideas and methods developed at **Rothamstad** were formalized in a statistical software called **GenStat.**

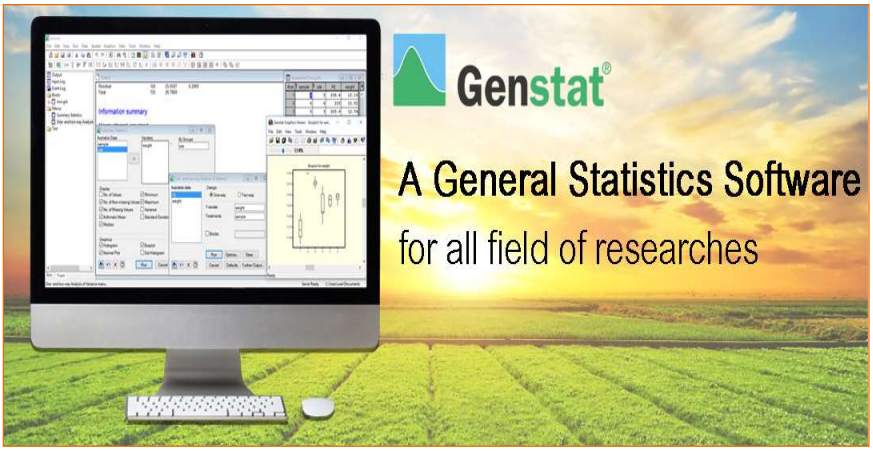

5

A History of the GenStat ANOVA by Roger Payne

Australasian Applied Statistics Conference 2018

50 Years of Genstat ANOVA

Graham Wilkinson's ANOVA algorithm has been a key part of Genstat ever since its early days. It was one of the motivations for the creation of Genstat. It was also the reason why I myself originally became involved with Genstat – initially to take up the responsibility for ANOVA ready for Graham's departure from Rothamsted in 1974. The algorithm provides a very efficient method of analysis that, even after more than 50 years, is unmatched elsewhere.

Analysis of variance and the associated design of experiments had been a Rothamsted speciality since the establishment of the Statistics Department under Sir Ronald Fisher in 1919. This was not a merely theoretical interest, but was motivated by the many experiments that needed to be designed, and then analysed, for the Rothamsted biologists. Fisher retired to Adelaide, and had a strong influence on Graham's statistical views. The Rothamsted connection was strengthened in 1965, when John Nelder visited the WAITE Institute in Adelaide, and began his collaboration with Graham. This laid the foundations for Genstat. Work on Genstat began in earnest, when John was appointed as head of the Rothamsted Statistics Department in 1968, and Graham joined the Department in 1971.

The original ANOVA algorithm was described by Wilkinson (1970), and its theoretical underpinnings by James & Wilkinson (1971). Payne & Wilkinson (1977) described the more efficient method for determining the structure of the design, that was my first task to get working when I took over. The relationship between the *first-order balanced* designs, that ANOVA analyses, and Nelder's (1965) *general balance* was explained by Payne & Tobias (1992), together with their algorithm that extended ANOVA, to estimate variance components and calculate estimates of treatment effects that combine information from all the strata in the design. Payne (2004) described how to obtain degrees of freedom for these combined effects.

The algorithm involves a sequence of *sweeps* in which effects are estimated, and then removed, from a working variate. There is also special sweep, known as a *pivot*, that projects the working variate into a specific stratum of the design. Matrix inversion is thus required only for the estimation of covariate regression coefficients and, as a result, the algorithm is very efficient in its use of workspace and computing time. Even 50 years on, this remains an important consideration.

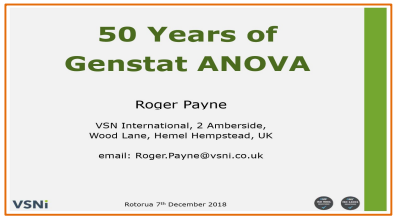

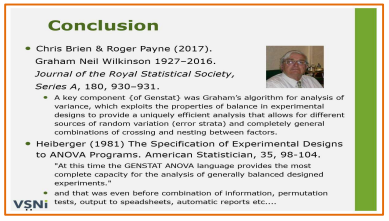

6

GenStat, is based on John Nelder's analyses of variance , which is a powerful formalization of the ideas in the work of others associated with Rothamsted (Senn, 2018).

- ✓ Developed analysis of variance in not in terms of linear models but in terms of symmetry.
- ✓ **How**? By incorporating Nelder's ideas of general balance into the Wilkinson ANOVA algorithm.
- ✓ Uses a very efficient sweep type algorithm.
- ✓ This is very intuitive for those analyzing designed experiments as it breaks the model down into two distinct parts - **treatment** and **block structure**.
- ✓ Made it possible to retain the conceptual simplicity of ANOVA type strata in the analysis.
- ✓ Will simply derive the correct analysis for almost any designed experiment.

These are unique and defining features of this software and make it extremely adept to the analysis of agricultural field experiments.

7

Statistical Science 2003, Vol. 18, No. 1, 118–131 © Institute of Mathematical Statistics, 2003 A Conversation with John Nelder by Stephen Senn

Nelder: During my first employment at Rothamsted, I was given the job of analyzing some relatively complex structured experiments on trace elements. There were crossed and nested classifications with confounding and all the rest of it, and I could produce analyses of variance for these designs. I then began to wonder how I knew what the proper analyses were and I thought that there must be some general principles that would allow one to deduce the form of the analysis from the structure of the design. The idea went underground for about 10 years. I finally resurrected it and constructed the theory of generally balanced designs, which took in virtually all the work of Fisher and Yates and Finney and put them into a single framework so that any design could be described in terms of two formulas. The first was for the block structure, which was the structure of the experimental units before you inserted the treatments. The second was the treatment structure—the treatments that were put on these units. The specification was completed by the data matrix showing which treatments went on to which unit. I published two papers in 1965, in the *Proceedings of the Royal Society of London*, at the end of which I postulated that it should be possible to write a general computer program, which, given these two structures, could deduce the form of the analysis of variance, the number of different kinds of contrasts of treatments in terms of their standard errors and so on.

Senn: Is this actually incorporated into Genstat, this particular approach of yours?

Nelder: Yes. As far as I know, Genstat is the only statistical package that does all these analyses by this single algorithm.

Senn: I must confess I have been very impressed, as a relatively recent Genstat user, by the way in which you are invited to declare the blocking structure and the treatment structure.

Nelder: Yes, this approach is almost unknown in the U.S., for example. It does seem to me to synthesize a lot of apparently different things and to put them into

8

The GenStat algorithm uses simple treatment and block structure.

Y-Variate = **treatment structure + block structure** + error structure

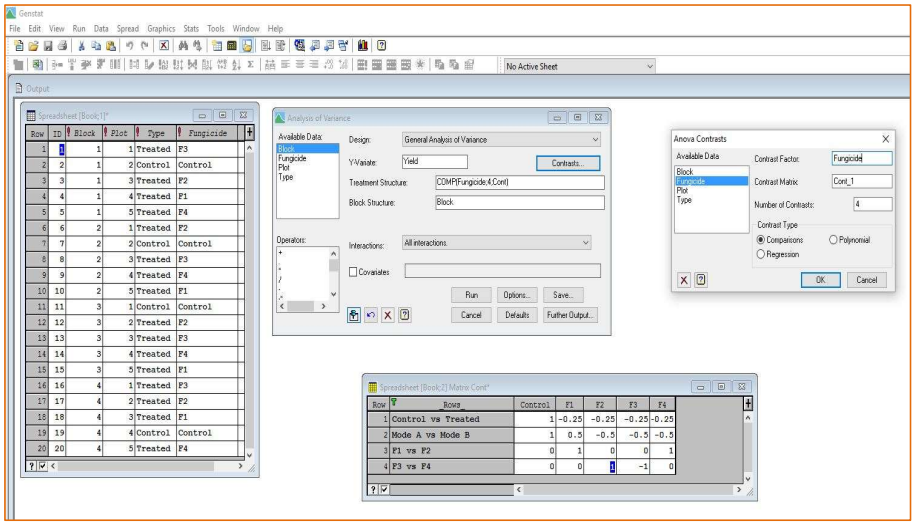

9

Treatment and block structure.

Y-Variate = **treatment structure + block structure** + error structure

Block structure – concerned with the structure (heterogeneity) of units.

Treatment structure– defines different treatment combinations applied to the units.

The ***structural component*** *(block)* allow to assess different sources of variation among experimental.

The ***explanatory component*** *(treatment)* provides information about the difference in response caused by treatments, and allows estimates of the sizes of these differences.

10

Deriving the Model

Treatment Structure specifies a *model formula* to define the model terms to be fitted.

In its simplest form, a model formula is a list of model terms, linked by the operator (**+**).

Example, **A + B** is a formula containing two terms, A and B, representing the main effects of factors **A** and **B** respectively.

The *crossing operator* (*****) is used to specify factorial structures.

Example, the treatment formula **Nitrogen * Sulphur** is expanded automatically by Genstat to become the formula **Nitrogen + Sulphur + Nitrogen.Sulphur**

The other commonly used operator is the nesting operator (**/**).

This occurs most often in block models an is specified by the **Block Structure**.

Example, the formula Block/Plot is expanded to become the formula Block + Block.Plot

11

Utilizes the concept of strata

Restrictions are imposed on layout of an experiment every time we design and conduct an experiment. These restrictions create different structural sources of variability among the experimental units called ***strata***. Each restriction in the structure of an experiment is called a **stratum**.

Stratum

In an analysis of variance calculated by `ANOVA`, a stratum is produced for every term in the formula specified by the `BLOCKSTRUCTURE` directive. There is thus one for each of the sources of random variation in the design. Genstat automatically discovers which treatment terms are estimated in each of the strata. (Strata is the plural of stratum.) The total sum of squares in each stratum is then partitioned into the sums of squares for the treatment terms estimated in the stratum, and a residual which provides an estimate of variability against which these treatment sums of squares should be compared. This ensures that each treatment line in the analysis of variance table is automatically compared with the correct residual.

12

RCB Example

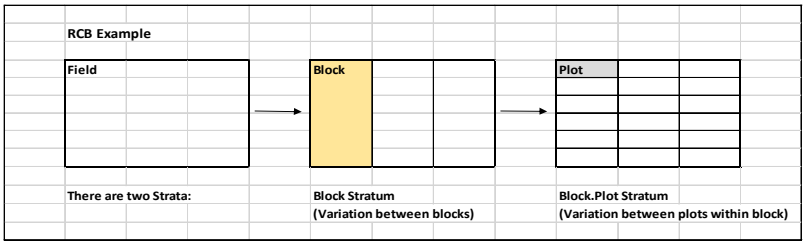

Explanatory component: TRT
Structural component: Block/Plot

Analysis of variance

Variate: Yield

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block stratum	3	174554.	58185.	3.04	
Block.PLOT stratum					
TRT	8	2616461.	327058.	17.10	<.001
Residual	24	459068.	19128.		
Total	35	3250084.			

13

Produces multi-stratum ANOVA tables

Recognition of the different roles played by the structural and explanatory components leads to the idea of the multi-stratum ANOVA - and makes explicate separation between them (Welham).

The multi-stratum ANOVA table results in an ANOVA table with separate components for each strata defined in the block structure

Variation with in each stratum is partitioned in sum of squares associated with the treatments that vary between units at that level of design.

If you use a multi-stratum ANOVA and you've specified your structure correctly then treatments terms get tested at the correct level of structure, **thus the correct analysis is completed**.

14

Comparison of ANOVA Tables

Analysis of variance

A **simple ANOVA** table does not make any distinction between describing the underlying structure of the data and those indicating the treatments applied.

Variate: Yield

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block	3	1944361.	648120.	5.86	
Treatment	5	1198331.	239666.	2.17	0.113
Residual	15	1658376.	110558.		
Total	23	4801068.			

The **multi-stratum ANOVA** table for the RCBD rearranges the simple ANOVA table to reflect the structure of the experiment. The RCBD has two distinct strata, a **Block stratum** and a **Block.Plot stratum**.

Variate: Yield

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block stratum	3	1944361.	648120.	5.86	
Block.Plot stratum					
Treatment	5	1198331.	239666.	2.17	0.113
Residual	15	1658376.	110558.		
Total	23	4801068.			

15

Split Plot Example

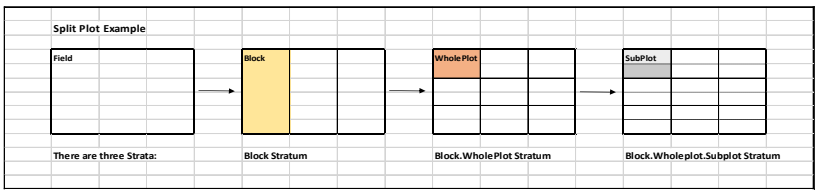

Explanatory component: **Variety*Harvest Method**
Structural component: **Block/WholePlot/SubPlot**

Expands to: Block + Block.WholePlot + Block.Wholeplot.SubPlot

Analysis of variance

A Split plot has three strata, a Block stratum and a Block.WholePlot stratum, and a Block.WholePlot.SubPlot stratum.

Variate: Yield

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block stratum	3	174554.	58185.	2.24	
Block.WholePlot stratum					
Variety	2	2138952.	1069476.	41.17	<.001
Residual	6	155847.	25974.	1.54	
Block.WholePlot.SubPlot stratum					
Harvest_Method	2	203119.	101560.	6.03	0.010
Variety.Harvest_Method	4	274390.	68598.	4.07	0.016
Residual	18	303221.	16846.		
Total	35	3250084.			

Uses the correct error terms so you get the correct F test

16

Easy to form contrasts & partition SOV

(Example from Welham et al, 2015)

Explanatory component: COMP(Fungicide;4;Cont_1)
Structural component: Block/Plot

Analysis of variance

Variate: Yield

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block stratum	3	14987.	4996.	1.43	
Block.Plot stratum					
Fungicide	4	133419.	33355.	9.58	0.001
Control vs Fungicide	1	125294.	125294.	35.97	<.001
Mode A vs. Mode B	1	5402.	5402.	1.55	0.237
F1 vs. F4	1	2178.	2178.	0.63	0.444
F2 vs. F3	1	544.	544.	0.16	0.699
Residual	12	41797.	3483.		
Total	19	190203.			

Explanatory component: POL(N;3)
Structural component: Block/Plot

Analysis of variance

Variate: Yield

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block stratum	2	2.8385	1.4192	4.40	
Block.Plot stratum					
N	3	6.1434	2.0478	6.35	0.027
Lin	1	5.9283	5.9283	18.37	0.005
Quad	1	0.0085	0.0085	0.03	0.876
Deviations	1	0.2065	0.2065	0.64	0.454
Residual	6	1.9359	0.3227		
Total	11	10.9178			

17

Easy to produce and plot residuals

Diagnostic plots, contour plot, shade plot, tables

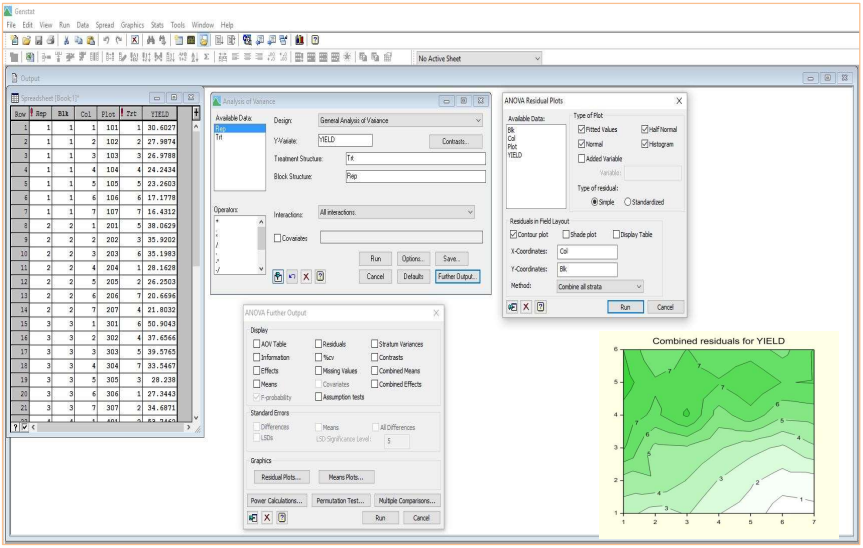

18

Comments on the GenStat ANOVA

In summary, the GenStat ANOVA fits the traditional approach to data modeling.

For example, in a linear model $Y_{ij} = \mu + \varepsilon_{ij}$, μ is the systemic part of the model and can be formulated by splitting this part into two components, treatment and block structure.

- ✓ *This does not, necessarily, require the formulation of a linear model in the traditional sense and transforming that model into computer code. The model can be derived by simple visualization of the design and use of simple notation in a menu system. The appropriateness of the analysis is easily confirmed by viewing the multi-stratum ANOVA (i.e. are all strata present and accounted for). This may be advantageous as it does not strictly require the learning and embedding of the tradition linear model mindset. According to Stroup (2013) this traditional mindset essentially has to be unlearned when transitioning to contemporary modeling (generalized linear mixed models).*

Can only be formed when the explanatory and structural components obey certain conditions of balance.

- ✓ The simplest case of which is when block and treatment factors are orthogonal, so that each term can be estimated independently of the other.

19

JOHN NELDER: FROM GENERAL BALANCE TO GENERALISED MODELS (BOTH LINEAR AND HIERARCHICAL). Stephen Senn Department of Statistics, University of Glasgow

(Published in Senn SJ. John Nelder: From general balance to generalised models (both linear and hierarchical). In Methods and Models in Statistics: In Honour of Professor John Nelder, FRS, Adams NM, Crowder MJ, Hand DJ, Stephens DA (eds). Imperial College Press: London, 2004.)

In fact, this extremely powerful general framework of John's[12 13] is the basis of the analysis of variance capabilities of Genstat®, a result of his further collaboration with Wilkinson already alluded to. A very wide class of designs, including completely randomised designs, randomised blocks, split plots, Latin and Graeco-Latin squares, split-split plots, balanced incomplete blocks, balanced lattices, Youden squares and many more[14], in fact all designs possessing the property of 'first-order balance', can be analysed using this approach. As far as I am aware, Genstat® is the only package that does this and although I am not going to attempt to explain the property of first-order balance I am going to draw attention to an explicit feature of this whole philosophy of analysis of variance that is lost in many modern approaches to data-analysis.

The feature is that a clear, and to my mind fundamental, distinction is drawn between blocking and treatment structures. Let me give an example from my own field, that of clinical trials. You could have a clinical trial in an indication in which you believed that the outcome would, other things being equal, differ strongly by sex. That being so you could decide to make sex a blocking factor by running two randomisation lists, one for men and one for women. Since, of course, you will have many patients of each sex under each treatment you have the structure of a two-way analysis of variance with replication. In a linear model you could have 'sex' as a main effect and 'treatment' as a main effect and also investigate the interaction between the two. Such a model makes no distinction of type between sex and treatment and in nearly all statistical packages there will be no way of distinguishing them. Not so with Genstat® in which you can declare 'sex' as a blocking factor and 'treatment' (appropriately) as a treatment factor. The point is that you have allocated the patients their treatments and these could have been different but you haven't allocated them their sexes and these could not and once you have declared one as a block and the other as a treatment Genstat® will go on to encourage you to make different sorts of inferences about them.

20

Comments on the GenStat ANOVA

Although the multi-stratum ANOVA implicitly identifies terms in the structural components of the model as a random term, it uses least squares estimates of their effects as if they were fixed terms.

The GenStat ANOVA is intuitive. It is the fundamental (starting point) approach to analysis of any agricultural field experiment.

It can be a useful tool to check and compare the output with more complex models and the correctness of output from other software where you do not understand how they derive F tests.

Fits a design based approach to data analysis. Analyze it as you randomize it!

21

Comments on the GenStat ANOVA

The GenStat ANOVA is not the best tool for all situations.

- ✓ In a modern day of computing is this mindset ($Y_{ij} = \mu + \varepsilon_{ij}$) correct or the best method, technically no. Is it useful, yes. If you keep in mind what it is for and its limitations (Stroup).

Linear mixed models (LMM) and generalized linear mixed models (GLMM) are often needed and preferred, these are available in GenStat.

LMM and GLMM are a powerful and flexible class of models that can be applied to may different situations found in agricultural experimentation.

- ✓ However, it is easy (especially for agronomists) to fit a wrong model and obtain misleading results with LMM. As the models become more complex (GLMM) so does the danger of misspecification of the model. This emphasizes that importance of analysis by the GenStat ANOVA, at least as a starting point for designed experiments.

22

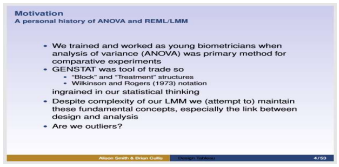

Design Tableau, Cullis & Smith, 2018

- ❑ We were fortunate enough to have trained and worked as young biometricians when analysis of variance (ANOVA) techniques were the primary method of analysis for comparative experiments.
- ❑ Our tool of trade was the GENSTAT package, so that the elegant notation of Wilkinson and Rogers and the framework of Block and Treatment structures became ingrained in our statistical thinking.
- ❑ So, despite the complexity of the LMM we now use, we appreciate the importance of maintaining these fundamental concepts, in particular the link between the analysis and the experimental design. We are concerned that this view is shared by only a few, as is evidenced by what we regard as a widespread mis-use of LMM for comparative experiments. This may either be due to an unintentional lapse in transitioning from ANOVA to LMM or a complete lack of exposure to traditional methods of analysis for comparative experiments.

23

Many of these ideas are summarized and taught in Welham's et al (2015) recent book

"**Statistical Methods in Biology**".

Contents:

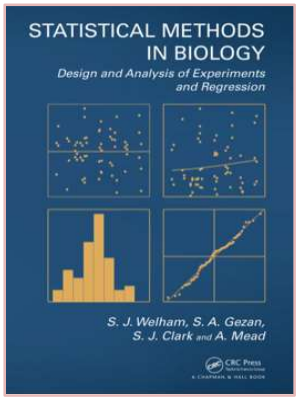

24

This approach is taught by **Mick O'Neill**

Statistics Manuals by Mick O'Neill

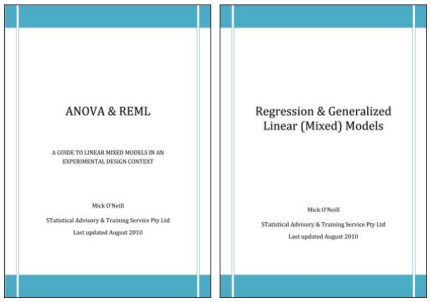

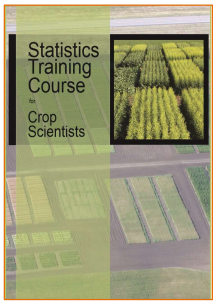

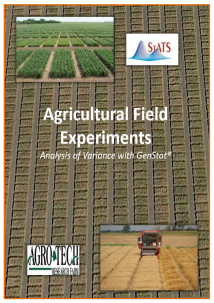

- Dataset for ANOVA & REML
- Dataset for Regression & Generalized Linear (Mixed) Model

25

Utilized by **Agro-Tech** to train agronomists and analyze data

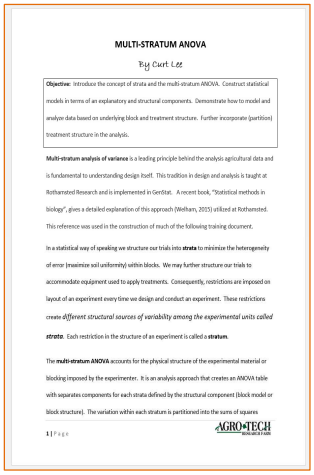

MULTI-STRATUM ANOVA

By Curt Lee

Objective: Introduce the concept of strata and the multi-stratum ANOVA. Construct statistical models in terms of an explanatory and structural components. Demonstrate how to model and analyze data based on underlying block and treatment structure. Further incorporate (partition) treatment structure in the analysis.

Multi-stratum analysis of variance is a leading principle behind the analysis agricultural data and is fundamental to understanding design itself. This tradition in design and analysis is taught at Rothamsted Research and is implemented in GenStat. A recent book, "Statistical methods in biology", gives a detailed explanation of this approach (Welham, 2015) utilized at Rothamsted. This reference was used in the construction of much of the following training document.

In a statistical way of speaking we structure our trials into **strata** to minimize the heterogeneity of error (maximize soil uniformity) within blocks. We may further structure our trials to accommodate equipment used to apply treatments. Consequently, restrictions are imposed on layout of an experiment every time we design and conduct an experiment. These restrictions create *different structural sources of variability among the experimental units called strata*. Each restriction in the structure of an experiment is called a **stratum**.

The **multi-stratum ANOVA** accounts for the physical structure of the experimental material or blocking imposed by the experimenter. It is an analysis approach that creates an ANOVA table with separates components for each strata defined by the structural component (block model or block structure). The variation within each stratum is partitioned into the sums of squares

1 | Page AGRO•TECH

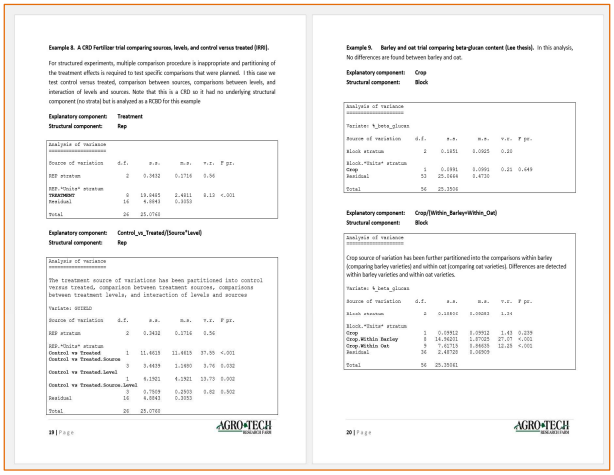

Example 8. A CRD Fertilizer trial comparing sources, levels, and control versus treated (999).

For structured experiments, multiple comparison procedure is inappropriate and partitioning of the treatment effects is required to test specific comparisons that were planned. I this case we test control versus treated, comparison between sources, comparisons between levels, and interaction of levels and sources. Note that this is a CRD so it had no underlying structural component (no strata) but is analyzed as a RCBD for this example

Explanatory component: Treatment
Structural component: Rep

Analysis of variance

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
REP stratum	2	0.0432	0.1716	0.56	
REP.*Units* stratum					
TREATMENT	6	19.0445	2.4811	8.13	<.001
Residual	16	4.6943	0.3053		
Total	24	25.0760			

Explanatory component: Control_vs_Treated/(Source*Level)
Structural component: Rep

Analysis of variance

The treatment source of variations has been partitioned into control versus treated, comparisn between treatment sources, comparisons between treatment levels, and interaction of levels and sources

Variate: YIELD

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
REP stratum	2	0.0432	0.1716	0.56	
REP.*Units* stratum					
Control vs Treated	1	11.4615	11.4615	37.55	<.001
Control vs Treated.Source	3	3.4439	1.1480	3.76	0.032
Control vs Treated.Level	1	4.1921	4.1921	13.73	0.002
Control vs Treated.Source.Level	3	0.7569	0.2503	0.82	0.502
Residual	16	4.6943	0.3053		
Total	24	25.0760			

19 | Page AGRO•TECH

Example 9. Barley and oat trial comparing beta-glucan content (Lee thesis). In this analysis, No differences are found between barley and oat.

Explanatory component: Crop
Structural component: Block

Analysis of variance

Variate: %_beta_glucan

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block stratum	2	0.1851	0.0925	0.20	
Block.*Units* stratum					
Crop	1	0.0991	0.0991	0.21	0.649
Residual	53	25.1644	0.4750		
Total	56	25.3936			

Explanatory component: Crop/(Within_Barley+Within_Oat)
Structural component: Block

Analysis of variance

Crop source of variation has been further partitioned into the comparisons within barley (comparing barley varieties) and within oat (comparing oat varieties). Differences are detected within barley varieties and within oat varieties.

Variate: %_beta_glucan

Source of variation	d.f.	s.s.	m.s.	v.r.	F pr.
Block stratum	2	0.18404	0.09202	1.24	
Block.*Units* stratum					
Crop	1	0.09912	0.09912	1.43	0.239
Crop.Within_Barley	8	14.96231	1.87029	27.07	<.001
Crop.Within_Oat	9	7.61715	0.84635	12.25	<.001
Residual	36	2.48720	0.06909		
Total	56	25.35041			

20 | Page AGRO•TECH

26

Call it what you want, the Rothamsted School, Nelder's Approach, the Strata Concept or just GenStat.

"Applying this approach", as per Stephen Senn https://errorstatistics.com

Here β is the multiple of the baseline difference that is used to correct the difference at outcome. However, at that time I had not appreciated the power of Nelder's approach to designed experiments. This, when applied, makes the issue crystal clear. What I did was apply Nelder's approach, which has the following key features

1. Recognising the distinction between blocking structure and treatment structure. The former reflect variation in the experimental material that exists logically prior to any experimentation and the latter variation that can in principle be affected by experimentation.
2. Defining the block structure.
3. Defining the treatment structure.
4. Mapping the treatment structure onto the block structure.
5. Analysing the results in terms of outcome, block structure and treatment structure.

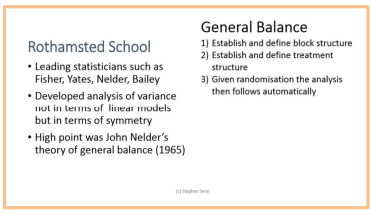

27

Summary

This Rothamsted approach to data analysis

- ✓ Implemented in the GenStat ANOVA algorithm which developed the the analysis of variance in terms of symmetry and general balance (works by an efficient sequence of sweeps).
- ✓ Makes a clear distinction between **Block** & **Treatment Structures**.
- ✓ Uses concept of **strata** (provides effects from each stratum where a term is estimated).
- ✓ Clear and comprehensive output (**multi-stratum ANOVA)**.
- ✓ Intuitive (given a randomization the correct analysis will follow.)

✓ Is a fundamental (must have) skill set for those who analyze agronomy experiments and, in my opinion, *should be required training for field research agronomists.*

28

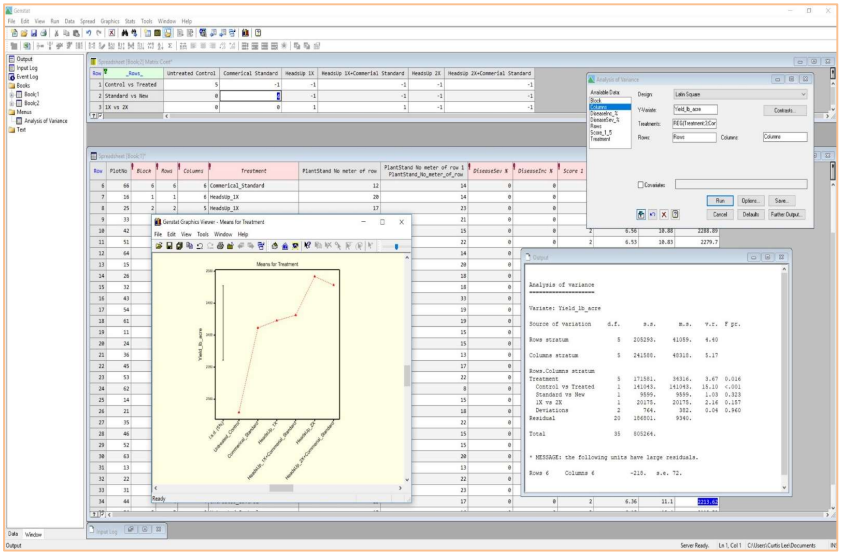

29

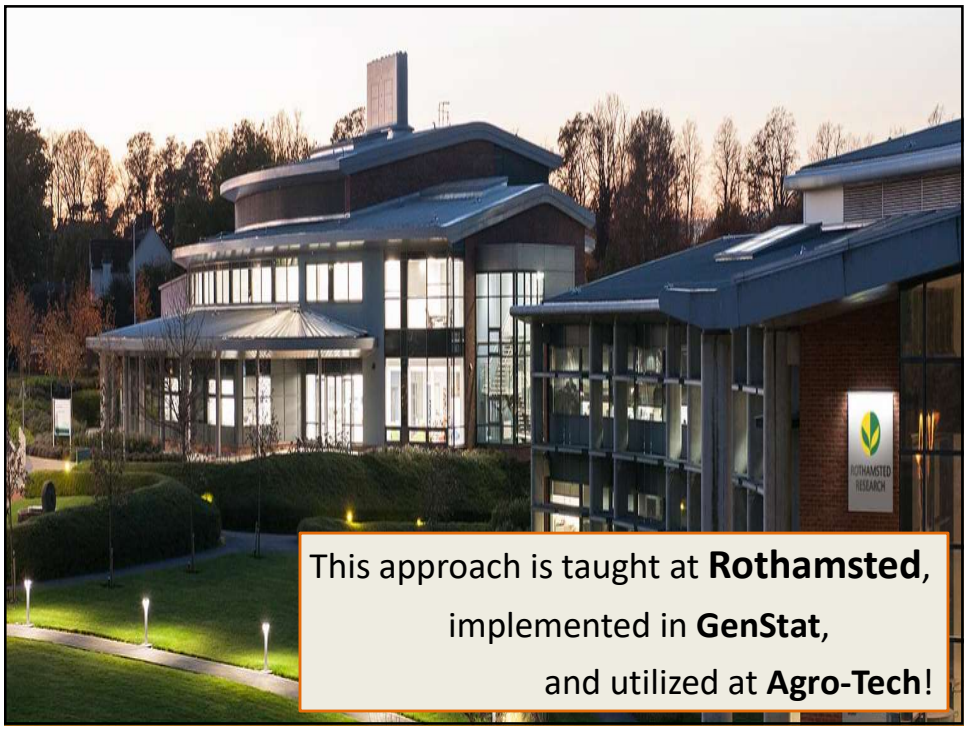
This approach is taught at **Rothamsted**,
implemented in **GenStat**,
and utilized at **Agro-Tech**!

30