Chapter 13

Experimental Design

P. Allen Hammer and Douglas A. Hopper

Contents

Introduction

Experimental design and statistical analysis of both field and greenhouse experiments have been discussed in texts (Anderson and McLean, 1974; Cochran and Cox, 1957; Steel and Torrie, 1980), but little attention has been given to statistical problems unique to growth chamber experiments. Went (1957) reported reduced plant phenotypic variability in growth chambers as compared with that in greenhouses. From these data, he proposed that less plant replication was needed in growth chamber experiments. This is perhaps a dangerous assumption; we will explore this recommendation and provide insight into some problems in the design and analysis of growth chamber experiments.

Growth chambers have provided the researcher precisely controlled environmental conditions in which to grow plants. Chambers serve two general purposes. First, the chamber produces a standard or defined environment in which plants are grown (commonly very important for the plant breeder or plant physiologist). Second, the chamber is used to study the effect of one or more controllable environmental parameters on growth (commonly very important for the horticulturist and agronomist).

Experimental design (blocking, replication, and randomization) is as important in growth chamber studies as it is in greenhouse or field studies. The growth chamber provides a known, researcher-determined environment for plant studies, but this does not mean that sources of "unwanted" variation have been eliminated. The use of appropriate experimental designs organizes the effect of unwanted variation so all treatments are equally affected by unwanted varia-

tion; consequently, such variation will have no effect on comparisons between applied treatments.

Standard Conditions

When the growth chamber provides a standard environment in which to grow plants, a given set of environmental parameters should be maintained. Problems such as decay of light, increase in plant size, and drift in temperature or humidity settings can be very troublesome in long-term studies. Accurate long-term measurements of the environmental conditions are necessary to compare results from experiments conducted at different times, or in different chambers set at the "same" environmental conditions. This measure of "sameness" should be precise enough to account for any environmental parameter that may interact with the applied treatments. Suggested methods of measurement are given in the appropriate sections of this handbook.

When experiments are repeated, control plants of the same species (nontreated plants), or a standard plant different from those being studied but with a known response, can be grown to monitor or confirm the correct setting of the environmental parameters. For example, if the researcher knows or suspects that a small difference in air temperature may affect the results of an applied treatment, then the standard plant should accurately measure a small difference in air temperature. The standard plant method is the simplest (requiring the least amount of equipment) and most accurate method when properly used.

Variable Environmental Conditions

Growth chambers are tools for studying the effects of environment on plant growth and development. Measured and controlled levels of environmental parameters can be maintained precisely, enabling observation of plant response at several settings over a desired range of study. The interaction of several environmental parameters can be properly investigated. Replication is necessary for effective experimental design, and replication must be done correctly in studies involving variable environments. Repeated independent applications of a treatment to designated experimental units enables a statistical calculation of *experimental error*. An estimate of *experimental error* allows valid statistical tests to be made between the different treatments.

Measurements made on several plants (or plant parts, such as several leaves on one plant) within the same treatment application, however, do not constitute true replication. Such measurements allow one to calculate what is called *sampling error*; however, *sampling error* cannot be used to make valid statistical tests of treatment differences. All plants in a growth chamber are considered to receive the same (not repeated) application of a treatment, that is, if specific settings of environmental parameters are the treatment. Thus, the use of multiple plants in a single chamber under a set of environmental parameters is not sufficient to properly replicate experimental treatments because multiple plants give only an estimate of sampling error.

True replication may be accomplished in two ways using growth chambers: (1) by repeating the same set of environmental parameters (a treatment combination) in two or more chambers (replication over space), or (2) by repeating the same set of environmental parameters in the same chamber at different times (replication over time). Either of these two methods enables the correct statistical calculation to estimate experimental error. Initial experiments may indicate a small experimental error relative to the sampling error. Our experience suggests, however, that this is not the case with most plant research investigations. Inherent variability within a chamber requires an estimation of experimental error

both between different chambers and within a chamber before using multiple samples from within a single chamber to statistically calculate an estimate of experimental error. These calculations must be done when the treatment is an environmental parameter and only one setting of a growth chamber is tested at a time.

Common Experimental Problems

Lack of uniformity among plants can result from unwanted variability. This has been determined for several cases: (1) among plants within a chamber (Carlson et al., 1964; Collip and Acock, 1967; Hammer and Langhans, 1972; Measures et al., 1973; Lee, 1977; Rawlings, 1979), (2) among plants grown in different chambers (Collip and Acock, 1967; Hammer et al., 1978; Lee, 1977; Rawlings, 1979), and (3) among plants grown in the same chamber at different times (Hammer et al., 1978; Lee, 1977; Rawlings, 1979). Differences in environmental conditions have been measured both within and between chambers (Kalbfleisch, 1963; Carlson et al., 1964; Gentner, 1967; Hammer and Langhans, 1972; Measures et al., 1973; Knievel, 1973; Tibbitts et al., 1976), and such differences are probably responsible for much of the unwanted variability and lack of reproducibility in plant growth both within and between chambers. Also, vibrations or handling of plants (Mitchell et al., 1975) and contaminants within buildings and chambers (Tibbitts et al., 1977) contribute unwanted sources of variation.

There certainly are enough data to suggest that time is an important variable in growth chamber studies (Lee, 1977; Hammer et al., 1978). In fact, it is so important that Lee (1977) and Rawlings (1979) suggested, from comprehensive studies of uniformity at the North Carolina State University Phytotron, that the between-trials (or runs over time) variation was more important (larger) than between-chamber variations. (Note: "variation" is usually referred to as a component of variance within the experiment.) They suggested blocking over trials (time) to account for this source of "unwanted" variability. When we use the term "block over time," we should be careful to clearly understand time as a block because time can index several things. Time can be chronological time (e.g., time of year, number of days), physiological time (maturity of a leaf, flowering), or time-related environmental variation. It is the time-related environmental variation we should block against in growth chamber studies.

Principles of Design

Fisher (1960) advanced three basic principles of experimental design: randomization, blocking, and replication. Since this initial work, many books have been devoted to this topic (Cochran and Cox, 1957; Federer, 1955; and Kempthorne, 1952). Consequently our comments only highlight the principles; the next section interprets them for specific growth chamber experimentation.

Variation

All biological material exhibits variation (such as genotypic variability), even when plant material of only one cultivar is grown under seemingly the same conditions. The previous section points out, however, that environments in growth chambers are not uniform throughout, regardless of the efforts made to control them. This lack of uniformity in environments adds to the already existing genotypic variability. The design principles of *blocking* and *randomization* are tools for minimizing the impact of biological (genotypic) and environmental variation on treatment comparisons of interests.

To clearly see the need for randomization, consider a hypothetical situation where treatments contain biological variation but no envi-

ronmental variation. Also, assume we have no way of predicting the biological variation before the study begins. Then suppose we unintentionally applied one treatment to all the smaller plants and another treatment to the larger plants. We had planned to record the effect of treatments on plant growth. Unfortunately, with this scenario our comparison of treatments would be confounded with (i.e., mixed up with) initial plant size. In other words, if we observe a different response between plants in two of the treatments, we do not know whether to ascribe this difference to the treatments themselves or to the initial size difference the plants had at the start of the treatments.

Randomization

Randomization virtually eliminates this problem. Randomization assures that each plant has the same chance of receiving any one of the treatments. Consequently, large and small plants have the same chance of influencing each treatment mean. In fact, randomization usually will do a better job of assigning plants to treatments than the investigator. This occurs because the investigator can perceive only a limited number of characteristics of the plant, whereas randomization assigns plants to treatments without regard to any specific characteristics. These comments apply only when the investigator has no specific information about the *organization of the biological variation*; if we have such information, we should use blocking, an idea discussed later.

Randomization has an important statistical benefit: It allows calculation of a valid estimate of the variation (variance) among similar plants. This quantity is required for making comparisons between treatment means. Rarely will two treatments have exactly the same means, so treatment means will almost always seem different by some amount. The estimate of variance is central to evaluating how much difference between treatment means might occur simply as a consequence of biological variation. The ratio of variation between treatments (*treatment error*) to variation within treatments (*experimental error*) will constitute the statistical test to determine if observed differences are indeed significant.

What could happen if the researcher assigned treatments to plants rather than randomized them? To answer this, suppose the researcher balanced the assignment of plants to two treatments so the average size of plants in each treatment was essentially equal before applying the treatments. This will have two effects that randomization would overcome: (1) the estimate of variance within treatments will be inflated by increasing the difference between individual plants within each of the treatments, while pretreatment means appear similar (if this does not make sense, take some numbers and try doing what is suggested, or see Federer, 1955, p. 14), and (2) removing all pretreatment difference in plant size will reduce the posttreatment difference. Consequently, such nonrandom assignment of treatments to plants reduces sensitivity in two ways: It reduces the observed treatment difference, and it increases the quantity used to measure the biological variation. Nonrandom assignment of plants to treatments can materially reduce a researcher's ability to detect real differences; other types of nonrandom assignment can also lead to biased results. *THUS: RANDOMIZE!*

How should an investigator actually carry out the randomization? Suppose there are 20 plants numbered 1 through 20 to assign to two treatments. Some random device or chance mechanism is needed to select 10 of these for the first treatment, with the rest going to the second treatment. For example, plastic discs numbered 1-20 can be put into a container, mixed thoroughly, and then 10 discs selected to give the plant numbers for the first treatment. (Slips of paper tend

to stick together so they do not work as well as plastic discs.) For a second method, an ordinary deck of (new) playing cards marked 1-52 may be used. The cards should be shuffled well, and cards taken from the top of the shuffled deck until 10 of the numbers 1-20 turn up. These plants go to the first treatment, and the rest go to the other treatment. For a third method, a table of random numbers (found in almost any statistics textbook) can be used. This is done by entering the random number table at some haphazardly chosen point and using the numbers found there to go to another part of the table. The investigator should find consecutive (non-overlapping) pairs of digits in the tables; he should ignore those over 20, but record the first 10 up to and including 20. These are the plants assigned to the first treatment.

In other contexts, a researcher may need only a few things; for two things, he should flip a coin. For 3, 4, 5, or 6 things, he should roll a die (one from a pair of dice) and ignore the higher number of dots for randomization of fewer than six things. A pair of dice should not be used because some numbers of dots (total of the two dice) occur with much higher frequency than others, thus yielding biased results instead of true randomization.

Randomization is not equivalent to haphazard or unplanned assignment. For example, an investigator could put the 20 plants of the previous paragraph in a row and select the first 10 for the first treatment. This would be a haphazard assignment. In a particular case, haphazard assignment may work as well as randomization, but there is no assurance of how good it is. In fact, randomization is much like insurance: You may not need it, but if you do need it, you need it badly! Thus randomize everything you can.

Blocking

The recommendations for randomization just given assumed no known environmental variation or recognized plant variation of a known sort. No matter how diligently we try, however, environments still vary somewhat. *BLOCKING* is a tool for dealing with known (measurable) sources of environmental and/or plant variation. A *block* is a set of *homogeneous* plants and/or microenvironments. *Homogeneous* means all plants are as similar as possible, and obvious differences are assigned to different blocks. At least, there should be more similarity of units within a block than between units in different blocks.

Typically, a block of environmentally similar units will be contiguous and usually quite close in space. In fact, it will often be a square area. Although square or nearly square blocks are conventional and typical, the essential feature is that the environment be relatively constant within the block. Data for the treatment should be recorded from the plants inside the block, and the guard-row plants would be discarded without any data recorded on those plants. A block could be irregularly shaped. Blocks should be made irregular in shape if enough is known about the total area of the experiment to indicate that the irregularly shaped area contains a more homogeneous environment than would any square or rectangular area. For example, extensive contour terracing has been done in the dryland wheat country of Kansas and Nebraska. A study of cultural practices or cultivars of wheat justifiably could have blocks laid out along the terraces. Such blocks would curve as they followed a contour across a field. Similar blocking is appropriate in a growth chamber if contours of constant performance, as with light levels, have been established by a relevant uniformity trial. (Note: the blocks defined here often are erroneously called replicates in the agronomic literature, whereas in most of the statistical literature, a replicate is an individual trial or plot.) Treatments are assigned to locations within each block so each treatment appears in each block

exactly the same number of times (usually once). A separate randomization must be performed to assign treatments to locations within each block. Each randomization should be performed as previously described.

Blocking would not be used if all plants grown in a controlled chamber were known to develop the same (say the same size). In this case, a single randomization should be performed over the whole growth chamber to assign treatments to locations. This is termed a *COMPLETELY RANDOMIZED DESIGN*. This type of design is actually most efficient because it allows more *degrees of freedom* to calculate the experimental error and thereby should yield the best estimate of this error (degrees of freedom are defined as a number that is one less than the number of units in a treatment, or n-1).

However, there are usually some detectable differences between plants grown in growth chamber experiments so the *RANDOMIZED COMPLETE BLOCK* should be used. Plants that exhibit variation of consequence (say a difference in size) should be grouped by that variation (grouped by size). Each constant-sized group should make up a block. In this case, blocks would minimize the impact of both microenvironment and plant variation on the comparisons of treatments. The experimental design described here is said to be complete because every treatment occurs in every block. (Because each treatment occurs the same number of times in every block, the design also is said to be balanced.) Incomplete designs, namely ones where not all treatments appear in each block, do exist, but they usually require a fairly large number of blocks. Such experiments must be set up and executed with care in controlled environments. A researcher should seek a statistician's counsel before deciding on the use of an incomplete block design. In the example above, each block consisted of constant-sized plants in a homogeneous environment. In other words, two sources of heterogeneity (differences in both environment and plant size) were held constant in each block. The effects of the two sources of heterogeneity are said to be confounded (inseparable). Confounding of sources of heterogeneity in construction of blocks is an acceptable experimental design practice. In contrast, it is not acceptable to confound treatments with the source of heterogeneity because then observed differences in plant growth may be a consequence of the known heterogeneity and not just of the treatments being studied.

Sometimes unwanted sources of heterogeneity cannot be confounded effectively in the blocks. For example, assume the variation among chambers was one source of heterogeneity and variation over several time periods was another. Both of these sources of variation need to be dealt with in the design of the experiment. Designs for two-way elimination of heterogeneity should be used. The simplest and most common one is called the *LATIN SQUARE DESIGN*, and a slightly generalized version is called a *LATIN RECTANGLE DESIGN*. For a study of three air temperatures in three chambers over three time periods, the latin square design would require that each air temperature appear once in each chamber and once in each time period. Such a design is restricted to having the same number of air temperatures, chambers, and time periods. The latin rectangle offers more flexibility by requiring that each air temperature appears the same number of times in each chamber and the same number of times in each time period. For example, a 2 x 4 latin rectangle could be used to compare two air temperatures in two chambers during four time periods.

Replication

The last major design component, *REPLICATION*, concerns observation of several *experimen-*

tal units under the same treatment and environment. Here an *experimental unit* is the amount of material to which a treatment was randomly allocated. Typically, an *experimental unit* is a plant, pot, or group of plants. The difference between replicate observations should reflect individual plant variation and micro-environmental variation not otherwise blocked out.

Replication has three major impacts: (1) it is required to obtain a valid estimate of *residual variance* (also called *error variance, experimental error, experimental variance,* or just *error*). This variance is the variation among experimental units treated alike; replication is required to have "alike" units to compare. Of course, randomization is just as critical for a valid estimation of variance as is replication. (2) The number of replications greatly influences the sensitivity of the experiment. A comparison between treatments is made by comparing the treatment means. The number of observations from which the mean is estimated greatly influences the closeness of the observed mean to the underlying true (population) mean. The more truly randomized observations that go into a mean, the closer the observed mean will be to the true mean. This occurs because observations above the mean average out with observations below the mean; more observations give a greater opportunity for this averaging out to function more completely. Specifically, if s^2 denotes the *variance* of individual plants and n is the number of plants in the treatment, then the *variance* of a treatment mean is s^2/n. In statistical terminology, s is referred to as the *standard deviation* of a treatment mean, and s/n is termed the *standard error of the treatment mean*. (3) More replication increases the precision with which the experimental (residual) variance (s^2) is estimated. Specifically, increased replication gives more degrees of freedom to estimate s^2. Even a cursory examination of a table of significant values of t or F shows how important it is to increase the degrees of freedom associated with the estimate of s^2, at least up to 30. Thirty or greater degrees of freedom normally yields a value of s^2 quite close to the true *population variance*, which is denoted by σ^2.

Two final thoughts on experimental design: Statistical procedures offer no substitute for careful planning and execution of an experiment. Instead, statistics offer the tools of experimental design to organize a study to minimize the impact of known environmental variability and provide accurate estimates of population parameters. Secondly, this merely highlights the principles of experimental design. Example designs are provided in the final section.

Conclusions

Variation is an important factor in growth chamber experimentation. We feel that Went's (1957) suggestion of using only a few replications may be wrong in many growth chamber experiments. One reason for using growth chambers is the much reduced experimental error, and everything possible should be done to minimize this error. Experimental error can be reduced by blocking, and we also suggest using guard rows or irregularly shaped blocks if the situation need it. Replication of the treatment in growth chamber studies is an important concern. When an environmental parameter is the treatment, the environmental treatment should be replicated in additional chambers. Multiple plants within a chamber provide only sampling error, at least until it is shown that no difference exists between runs or chambers. Each experiment conducted in a growth chamber should be considered unique, and the best statistical approach for that experiment should be used. A statistician should be consulted in the planning stages of each experiment.

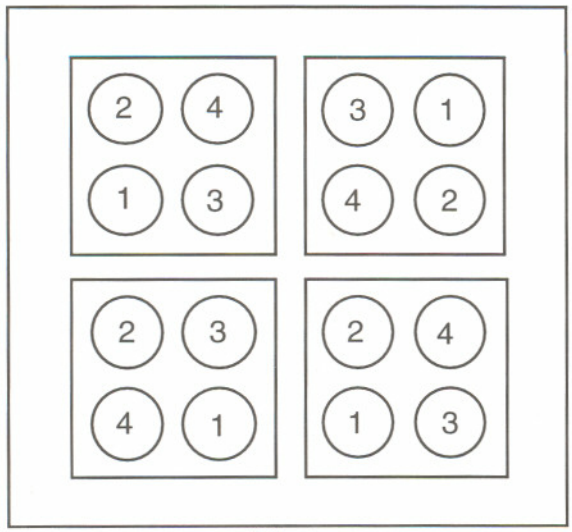

4 blocks
4 treatments (nitrogen level)

Source	df	E (MS)
Mean	1	
Blocks within chamber	3	$\sigma^2_C + 4\sigma^2_B$
Treatments (nitrogen)	3	σ^2_C + Nitrogen
Error	9	σ^2_C

Figure 1. The experimental units represented by the circles have been arranged in small square blocks to account for within chamber variation.

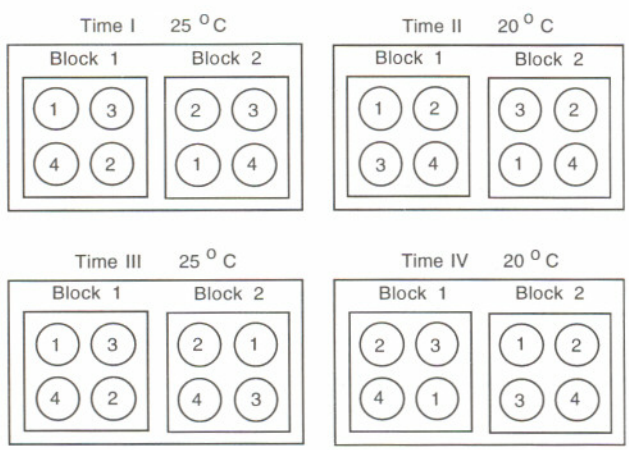

(The numbers in the circles represent 4 different nitrogen treatments)

Source	df	E (MS)
Mean	1	
Time (Blocks)	3	
Air temperature	1	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$ + Temperature
Error A	2	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$
Blocks within chamber	1	$\sigma^2_C + 4\sigma^2_B$ + Position
Error B	3	$\sigma^2_C + 4\sigma^2_B$
Subplots within blocks	24	
Nitrogen	3	σ^2_C + Nitrogen
Nitrogen X Air temperature	3	σ^2_C + Interaction
Error C	18	σ^2_C

Figure 2. When one chamber is used to study two temperatures blocked over time, differences in the response at the different temperatures will be unbiased comparisons. However, each response will be confounded with a chamber effect.

Example Designs and Analysis

The best way to discuss the general principles of design and analysis when the chamber is an experimental unit is to approach one illustration in several ways. The researcher may have a simple four-treatment study with one chamber or have one set of treatments using individual chambers and another set of treatments within each of the chambers. With a study in which all treatments can be maintained within one chamber (for example, a comparison of four nitrogen treatments), the study needs to be conducted only once, and the analysis is shown in Fig. 1. If, however, a researcher has treatments using separate chambers, the analysis is more complex. For this example, air temperatures of 20° and 25°C are the main plot (chamber) treatments and four nitrogen treatments are the within chamber treatments. If only one chamber is available to the investigator, the study should be run four times, two times at each temperature (Fig. 2). Blocks over time would be at the chamber level; similarly, Blocks within chambers would be within chamber variation. In this case, the investigator will get an unbiased look at the differences in response at the two temperatures. The temperature response would contain a treatment and chamber component; however, the chamber component is assumed to be nearly constant for each run and thus subtracts out from the temperature response. If such a design were used, the temperatures should be randomly allocated to runs, and the blocks should be located in the same position within the chamber for each run for the ANOVA table in Fig. 2 to be used. If chamber blocks are not in the same position, the variances or Blocks Within Chamber and Error B cannot be separated in the ANOVA table. Nevertheless blocking within chambers would still be important. There would be no change in the rest of the analysis.

When two chambers are available, the minimal design would be to block over time for four runs (Fig. 3). In this case, the analysis of main plots would be a 2 x 4 latin rectangle. Note the crossover of temperature and chamber each time and that the block effects within chambers A and B are different. Again, if the position of the blocks within a chamber changes each time, Blocks Within Chamber and Error B cannot be separated. In the ANOVA table of Fig. 3, several things become clear. Most of the power for testing differences is associated with the nitrogen treatments within chambers. Differences in response to temperature would need to be very large to be detected with a single degree of freedom. One may want to test Error A with Error B and pool them if they are not significantly different. With the errors pooled, there would be a slight increase in error degrees of freedom.

If, in the previous case, only two runs over time were used, the interpretation of results would be very difficult (Fig. 4). An estimate of Error A would not be available for testing a temperature effect. The mean square for chamber might serve as an error term in this case; however, with only 1 degree of freedom and probably a chamber effect present, the test would be so conservative that significance would be almost impossible. Again note the crossover of temperature and chamber (temperatures were changed between chambers for the runs) and the degrees of freedom associated with Error C.

And finally we get to the worst case of confounding, which is probably the one many investigators have faced. If two chambers are operated at two temperatures at one time, there is no way to test the temperature effect. It is confounded with chamber effect and has no appropriate error term. The analysis of variance in Fig. 5 shows why a misleading result may occur when comparing two chambers with different temperatures using within chamber variance.

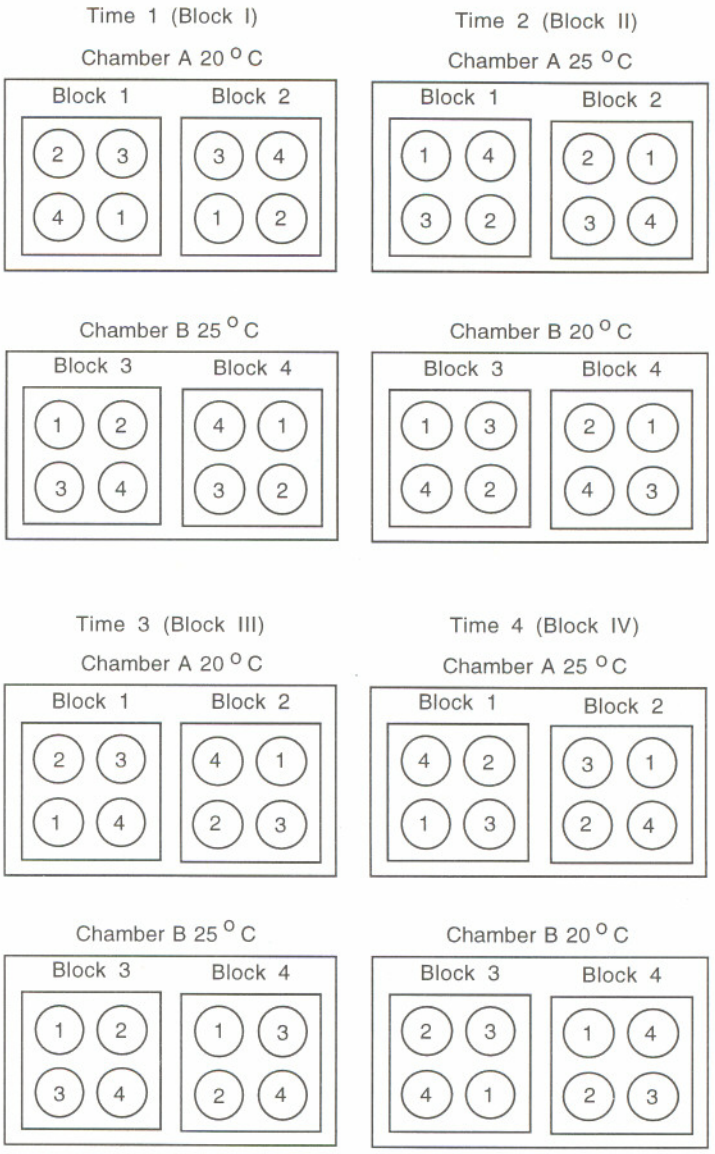

(The numbers in the circles represent 4 different nitrogen treatments)

Source	df	E (MS)
Mean	1	
Main Plot	7	
Time (Blocks)	3	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A + 16\sigma^2_T$
Chamber	1	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$ + Chamber
Air temperature	1	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$ + Temperature
Error A	2	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$
Blocks within a chamber	2	$\sigma^2_C + 4\sigma^2_B$ + Position
Error B	6	$\sigma^2_C + 4\sigma^2_B$
Subplots within blocks	48	
Nitrogen	3	σ^2_C + Nitrogen
Nitrogen X Air temperature	3	σ^2_C + Interaction
Error C	42	σ^2_C

Figure 3. This would be the minimal design needed for unconfounded, unbiased results for the example discussed. The main plots are a 2 x 4 latin square design.

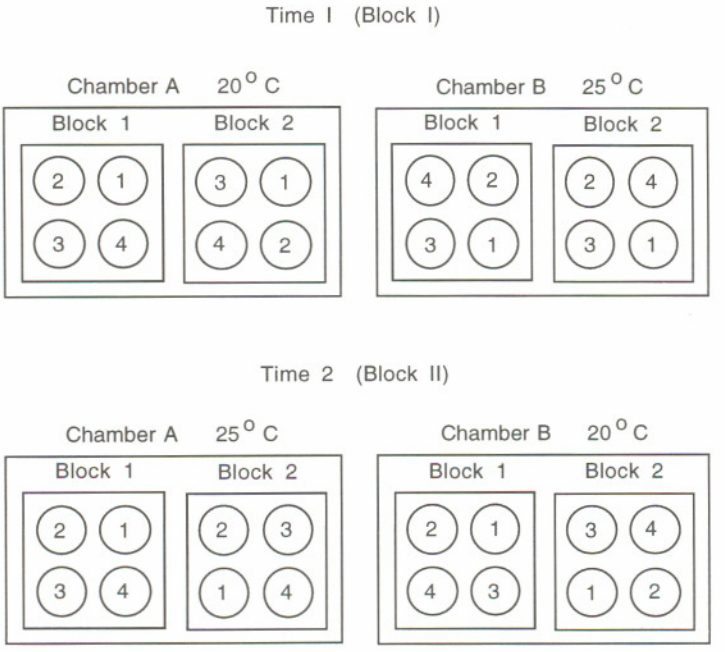

(The numbers in the circles represent 4 different nitrogen treatments)

Source	df	E (MS)
Mean	1	
Main Plot	3	
Time (Block)	1	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A + 16\sigma^2_T$
Chamber	1	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$ + Chamber
Air temperature	1	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$ + Temperature
Error A	0	Not available
Blocks within chamber	2	$\sigma^2_C + 4\sigma^2_B$ + Position
Error B	2	$\sigma^2_C + 4\sigma^2_B$
Subplots within blocks	24	
Nitrogen	3	σ^2_C + Nitrogen
Nitrogen X Air temperature	3	σ^2_C + Interaction
Error C	18	σ^2_C

Figure 4. In this case the main plot treatment is blocked twice over time. Although one can estimate the effect of time, chamber, and temperature, there is no estimate of error A with which to make a test.

The expected mean square, E(MS) for temperature, contains not only within chamber variance (σ^2_C) and the temperature effect, but it also contains $4\sigma^2_B + 8\sigma^2_A$ + chamber effects. Thus if Error C is used to test for a temperature effect, a significant F test may tell us nothing about temperature. Significance could be due to temperature effect, chamber effect, chamber to chamber variation (σ^2_A) and within chamber environmental variation (σ^2_B). Any combination of other factors could seem to be temperature effects without any temperature effect existing. This is a real example of "seeing" differences that may not be real.

In some cases, confounding cannot be eliminated (e.g., the requirement of special lamp fixtures in a chamber when comparing different lamp types), and the investigator has little choice except to be aware of the confounding of treatment and chamber and report results accordingly. However, it is important to repeat the treatments over time in this case.

These examples should show the importance of spending time with a statistician in the planning stages of each growth chamber experiment.

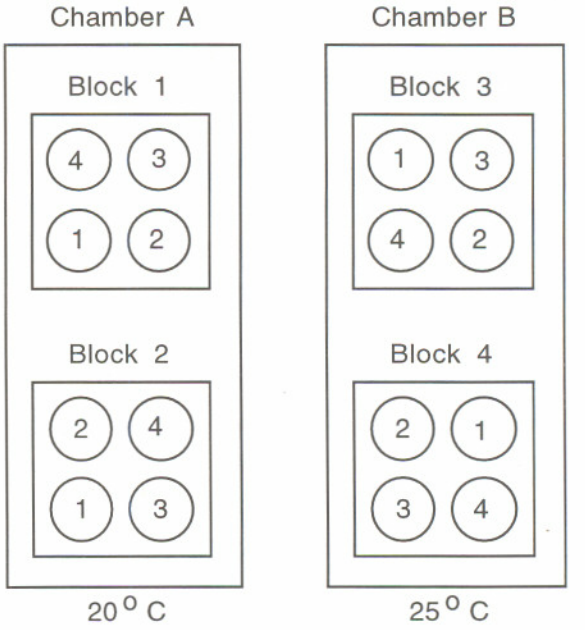

(The numbers in the circles represent 4 different nitrogen treatments)

Source	df	E (MS)
Mean	1	
Main Plot	1	
Temperature / Chamber	1	$\sigma^2_C + 4\sigma^2_B + 8\sigma^2_A$ + Temp/Chamber
Error A	0	Not available
Blocks within a chamber	2	$\sigma^2_C + 4\sigma^2_B$ + Position
Error B	0	Not available
Subplots within blocks	12	
Nitrogen	3	σ^2_C + Nitrogen
Nitrogen X Air temp.	3	σ^2_C + Interaction
Error C	6	σ^2_C

Figure 5. In this case the main plot has not been replicated. There is no estimate of error A. Temperature and chamber effect cannot be separated, and a test for temperature differences does not exist.

Some will argue that growth chamber space is much too expensive for the amount of replication and blocking we have suggested here. Is it not cheaper to conduct one well-planned and thus interpretable experiment than many poorly planned experiments that defy useful interpretation because of major confounding?

Literature Cited

Anderson, V. L., and R. A. McLean. 1974. Design of Experiments. Marcel Dekker, New York.

Bancroft, T. A. 1964. Analysis and inference for incompletely specified models involving the use of preliminary test(s) of significance. Biometrics 20:427-442.

Carlson, G. E., G. A.Motter, and V. G. Sprague. 1964. Uniformity of light distribution and plant growth in controlled-environment chambers. Agron. J. 56:212-243.

Cochran, W. G., and G. M. Cox. 1957. Experimental Designs. Wiley, New York.

Collip, H. F., and Acock, B. 1967. Variation in plant growth within and between growth cabinets. Report of the School of Agriculture, University of Nottingham. pp. 81-87.

Federer, W. T. 1955. Experimental Design. Macmillan, New York.

Fisher, R. A. 1960. The Design of Experiments. 7th ed. Oliver and Boyd, London.

Gentner, W. A. 1967. Maintenance and use of controlled environment chambers. Weeds 15:312-316.

Hammer, P. A., and R. W. Langhans. 1972. Experimental design consideration for growth chamber studies. HortScience 7:481-483.

Hammer, P. A., T. W. Tibbitts, R. W. Langhans, and J. C. McFarlane. 1978. Base-line growth studies of "Grand Rapids" lettuce in controlled environments. J. Am. Soc. Hortic. Sci. 103:649-655.

Hruschka, H. W., and E. J. Koch. 1964. A reason for randomization within controlled environment chambers. J. Am. Soc. Hortic. Sci. 85:677-684.

Kalbfleisch, W. 1963. Artificial light for plant growth. Pp. 159-174. In: Engineering Aspects of Environmental Control for Plant Growth. Proc. Symp., Sept. 1-5. 1962, CSIRO, Melbourne, Australia.

Kempthorne, O. 1952. The Design and Analysis of Experiments. Wiley, New York.

Knievel, D. P. 1973. Temperature variation within and between rooting media in plant growth chambers. Agron. J. 65:398-399.

Lee, C. S. 1977. Uniformity studies with soybeans at the North Carolina State University Phytotron. Thesis Series No. 1153. Institute of Statistics. North Carolina State University, Ra1eigh, North Carolina.

Measures, M., P. Weinberger, and H. Baer. 1973. Variability of plant growth within controlled environment chambers as related to temperature and light distribution. Can. J. Plant Sci. 53:215-220.

Mitchell, C. A., C. J. Severson, J. A. Wott, and P. A. Hammer. 1975. Seismomorphogenic regulation of plant growth. J. Am. Soc. Hortic. Sci. 100:161-165.

Rawlings, J. O. 1979. Personal Communications. Department of Statistics, North Carolina State University, Raleigh, North Carolina.

Snedecor, G. W., and W. G. Cochran. 1971. Statistical Methods. Iowa State University Press, Ames.

Steel, R. G. D., and J. H.Torrie. 1980. Principles and Procedures of Statistics. 2nd ed. McGraw-Hill Book Co., New York.

Tibbitts, T. W., J. C. Mc Farlane, D. T. Krizek, W. L. Berry, P. A. Hammer, R. H. Hodgson, and R. W. Langhans. 1977. Contaminants in plant growth chambers. HortScience 12:310-311.

Tibbitts, T. W., J. C. Mc Farlane, D. T. Krizek, W. L. Berry, P. A. Hammer, R. W. Langhans, R. A. Larson, and D. P. Ormrod. 1976. Radiation environment of growth chambers. J. Am. Soc. Hortic. Sci. 101: 164-170.

Went, F. W. 1957. Genotypic and phenotypic variability. Pp. 195-201. In: Experimental Control of Plant Growth. Chronica Botanica, Waltham, Massachusetts.