

“checks” or “controls.” Without the proper controls, you will not be able to say that the new nematicide worked better than the currently used nematicide or even that the new nematicide worked better than no nematicide! The questions you wish the experiment to answer should indicate what treatments should be included as controls.

It is frequently desirable to have both a positive and a negative control in an experiment. The negative control helps you determine if the treatments being tested work better than some minimal treatment (or nothing) and positive controls help you determine if the treatments being tested work better than the current standard practice. You may have several control treatments in an experiment if you currently have several viable options from which to choose. For example, if you currently can choose either of two fungicides to control a leafspot problem, you may wish to include them both as controls in your experiment when you test new products. You do not have to include all currently available options as controls for the experiment to be useful, but you can.

REPLICATION

In an experiment, replication means that individual treatments (such as each of the five pesticides being tested in an experiment) have been applied to more than one plot. **Replication is necessary because all test plots are not identical**, and that leads to variation in the data you collect; you will not get exactly the same results from two plots that received the same treatment. You can take steps to minimize the effect of variation if it has an identifiable cause, but there will always be some variation among plots that cannot be controlled. In statistical terms, uncontrolled variation is called experimental error. The purpose of replication is to allow you to make a more accurate estimate of how each treatment performed even though there is uncontrolled variation in the experiment. This can best be shown in an example.

Suppose you have 10 rose bushes and you want to test if a new fungicide will protect the bushes from black spot, a fungal leaf disease. You could pick five plants to leave untreated as a control and spray the other five with the fungicide. Later, when black spot is evident on the leaves, you count the number of diseased spots on each plant and compare the two treatments. The five untreated plants have 26, 21, 19, 25, and 23 infected spots (a treatment mean, or average, of 22.8 spots per plant), and the fungicide treated plants have 20, 15, 18, 21, and 20 spots (a mean of 18.8). Statistical analysis indicates that the fungicide did in fact reduce the number of infected sites, but you would not be able to determine that if you only had one treated and one untreated plot. Think about the following examples based on the data above. If you only had two plants and the untreated plant had 26 spots and the treated plant had 15 spots, it may seem easy to determine that the fungicide reduced disease. But what if the untreated plant had 19 spots and the treated plant had 21 spots? You might conclude that the fungicide did not work or even that it increased disease! Adequate replication can minimize this problem.

It is common when you have several replications of each treatment to have data like that in the rose example above: The treatment means are different but individual measurements may overlap. In this example, the lowest measurement from the untreated plants was 19, and the highest measurement from the fungicide treated plants was 21, but the treatment means were 22.8 for the untreated plants and 18.8 for the treated plants. **Replication of treatments increases your ability to detect differences in treatment means.** Having more replications allows you to identify (statistically) smaller differences in treatment means than you could identify with fewer replications.

The number of replications that you need is influenced by the biology of what you are testing, how close together the treatment means are, and how much variation exists within a treatment. For field tests in plant pathology, nematology, weed science, soil fertility studies, and entomology, a minimum of four replications is suggested, but five or six replications are **much** better. If treatment means are close together or variation is relatively large among the plots that received the same treatment, then you may need more replications to detect differences among treatments.

Just as the data may vary within a replicated treatment, the results may vary among experiments if the whole experiment is repeated. This can happen because of different weather conditions, different disease or insect pressure, or many other factors beyond your control. This does not mean that the results of a single experiment are not valid, but it does make it dangerous to draw conclusions from a single experiment. The one set of results you have may indicate treatment differences, but if you repeated the test several times you might not see those treatment differences again. **If the test is repeated (and that means you cannot change any of the treatments) and you get similar results, then you can be much more confident that your conclusions are correct.**

RANDOMIZATION

Randomization in an experiment means that the treatments are assigned to plots with no discernable pattern to the assignments. The reason randomization is important is that the positioning of treatments within the block may affect their performance. One example of this is an experiment testing five corn hybrids (labeled 1 through 5) in which you plant the hybrids in the same order in each block: 1, 2, 3, 4, then 5 (see figure 1). If hybrid 2 is naturally much taller than the others, it can slightly shade the hybrids planted next to it (hybrids 1 and 3) and unfairly make them look a little bit worse than they would look if they were not planted next to hybrid 2. Another example is a field in which soil fertility gets progressively lower as you cross the field from east to west so that productivity is reduced as you go from one side of the field to the other. If two corn hybrids are planted side by side but within a block, hybrid 1 is always planted on the east side of hybrid

have more weeds, you can remove that source of variation from the statistical analysis by blocking and improve your chances of identifying differences among treatments.

The process of blocking follows a logical sequence. First, you determine that there is something (weeds, drainage, sun/shadow, water, soil type, etc.) that is not uniform throughout the experimental area (field, greenhouse, etc.) that may influence whatever you are measuring (yield, plant height, etc.). Then you can arrange your treatments into blocks so that the area within each block is as uniform as possible (see figure 2). Though the area within a block should be relatively uniform, there may be large differences among the blocks, but that is what makes blocking effective. **Your goal is to maximize the differences among blocks while minimizing the differences within a block.**

The shape of the blocks is not important as long as the plots within a block are as uniform as possible. Ideally, the only differences among plots within a block should be due to the treatments. Blocks in field experiments are usually square or rectangular, but they may be any shape. Blocks in the same experiment do not have to be the same shape; the shape of individual blocks will be determined by variation in the field that you are trying to minimize. If you are not sure what shape your blocks should be, square or nearly square blocks are usually a safe choice.

Blocks may be arranged through the field in many ways. If the field is wide enough, an easy way to arrange blocks is to place them side-by-side all the way down the field (see figure 3). But blocks do not have to be contiguous and may be scattered through the field in any way that is convenient for you.




Figure 3. An easy way to arrange blocks is to put them side by side across the field. Letters represent different treatments.

Factorial Arrangement of Treatments

A factorial arrangement of treatments is not an experimental design, though you will often hear it referred to as a factorial design or a factorial experiment. A factorial arrangement of treatments means that the experiment is testing two or more factors at the same time, and that the experiment includes all combinations of all factors. The term “factor” is used to describe a group of treatments that have something in common. Fungicides, sources of nitrogen, or corn hybrids could be considered factors in an experiment. Factors may be defined broadly or narrowly in different experiments. All herbicides may be grouped as a factor in one experiment, but pre-plant and post-plant herbicides may be treated as separate factors in another experiment. A single-factor experiment

tests one factor at a time; a two-factor experiment tests two factors at once.

Most simple on-farm experiments are single-factor experiments (in a Completely Randomized or Randomized Complete Block design) and compare things such as crop varieties or herbicides, but it is sometimes useful to test two or more factors at once. For example, a two-factor experiment would allow you to compare the yields five corn hybrids at three planting dates. This accomplishes three things at once:

1. It allows you to compare the corn hybrids to each other.
2. It allows you to evaluate the effect of planting date.
3. It allows you to determine if varying the planting date changes the relative performance of the hybrids (e.g. one hybrid may only perform well if planted early).

The first two could be done in separate single-factor experiments, but the third can only be achieved by having both factors in a single experiment. This becomes especially important if one factor can have a significant influence on the effect of the other factor. For example, you might test soybean varieties as one factor and nematicides as another factor. If a few varieties have good nematode resistance but others do not, they may appear equally good when effective nematicides are used but varieties with resistance would appear much better when nematicides are not used. In cases like this, the effect of one factor (variety) is strongly influenced by the other factor (nematicide). When one factor influences the effect of the other factor, there is said to be a significant interaction between the two factors. **It can be very important to know if there is an interaction between factors, because if there is an interaction, you can make predictions or recommendations based on the results of single-factor experiments ONLY when all other factors are at the same levels they were at in the experiment.** If you change some factor not included in the experiment, the results from your single-factor experiment may no longer be valid.

With a factorial arrangement of treatments, all values (or levels) of each factor must be paired with all levels of the other factors. If you have two nematicides and five soybean varieties, then your treatment list must include each variety with each nematicide for a total of 10 treatments. This would be referred to as a “two by five factorial” to denote how many factors were present in the experiment and how many levels of each factor were used. The number of treatments increases quickly when you add more levels for a factor (if you used three nematicides instead of two, you would have 15 treatments instead of 10), so choose your levels carefully or the experiment can get too large to manage.

A factorial arrangement of treatments can be a very powerful tool, but because the number of treatments can get very large it is best used when some reason exists to believe that the factors may influence each other and have a significant

interaction. **If there is no suspicion that the factors may influence each other, it is frequently easier and more thorough to test the factors in separate experiments.** A factorial arrangement of treatments can be used with a completely randomized experimental design or a randomized complete block design. The top half of figure 4 shows a factorial arrangement of treatments in a randomized complete block design.

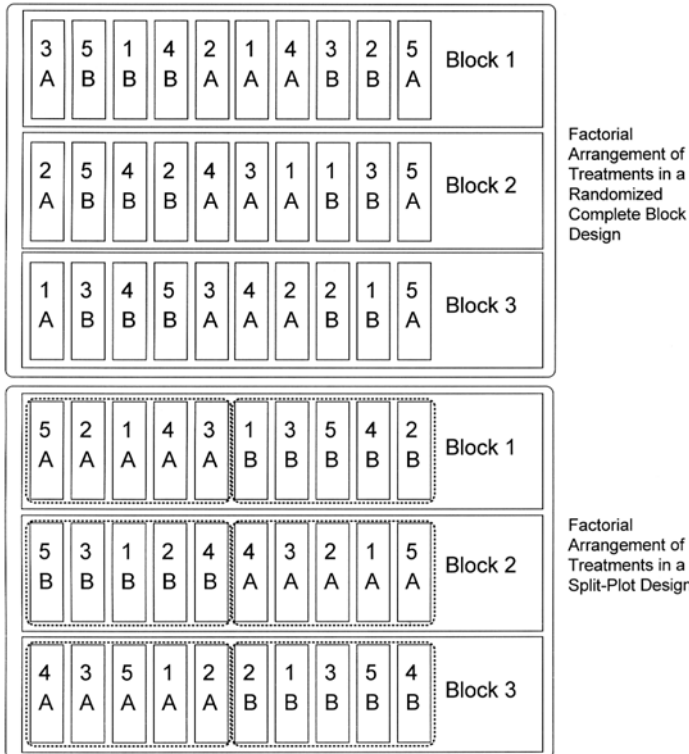


Figure 4. A 2x5 factorial arrangement of treatments in a randomized complete block design (above) and in a split-plot design (below). A and B represent two levels of one factor, and the numbers (1-5) represent five levels of a second factor. The combinations (e.g., 4A, 5B, etc.) denote individual treatment combinations. Either experimental design could be used, but the randomized complete block design is preferred unless the split-plot design is required by some limitation on randomization.

Split-Plot Experimental Design

A split-plot experimental design is a special design that is sometimes used with factorial arrangements of treatments. This design usually is used when an experiment has at least two factors and **some constraint prevents you from randomizing the treatments into a randomized complete block design.** Such a constraint may be based on equipment limitations or on biological considerations. For example, the equipment you have may make it difficult to put out a soil fumigant in randomized complete blocks, but you may be able to put out the fumigant so that all treatments within a block that get the fumigant will be clustered together rather than scattered throughout the block. You can use a split-plot experimental design to work around this limitation as long as you are able to randomize the other factors. There are other

situations when this design is appropriate, but a constraint on randomization is the most likely to occur.

Suppose you want to test the effect of five fungicides to control *Cylindrocladium Black Rot* on two varieties of peanut. In this test, you would have a 2x5 factorial arrangement of treatments: The two factors would be varieties (2 levels of this factor) and fungicides (5 levels of this factor). Because a factorial arrangement of treatments is not an experimental design, you still have to select an experimental design that best meets your needs. **If you are able to randomize varieties and fungicides within a block, then you should pick a randomized complete block design.** If there is some reason why you cannot completely randomize the treatments within each block, then you may be able to use a split-plot design to work around that limitation. For example, you may have a six-row planter but only enough space in the field to put out four-row plots. To resolve this dilemma, you could plant all of the plots that have the same peanut variety together within a block and then randomize the five fungicide treatments within each peanut variety.

In split-plot designs, the terms “whole plots” and “subplots” refer to the plots into which the factors are randomized. As the names imply, whole plots are subdivided into subplots. In figure 4, a whole plot would be the areas designated with A or B, and the subplots, the subdivisions within the whole plots, are designated 1, 2, 3, 4, or 5. In this example, A and B could represent two varieties (two levels of one factor) and the numbers could represent different fungicides (five levels of a second factor). Each whole plot serves as a block for the subplot treatments.

To assign treatments in a split-plot design, start by identifying where each block will be. Then randomize the whole plot treatments within each block. The whole plot treatments will be the treatment that you are unable to randomize into a randomized complete block design. The subplot treatments can then be randomized within each whole plot treatment (see figure 4).

DATA COLLECTION

You can collect an almost infinite amount of data in any experiment, but not all of it will be useful. Proper planning will ensure that you collect the right data to address your test’s objective. The “right” data to collect can usually be determined by examining the stated purposes of the experiment. For example, if the objective of a peanut leafspot fungicide trial is “to evaluate the ability of five fungicides to reduce leafspot incidence and severity,” then collecting data on leafspot incidence and severity and peanut yield should seem obvious. Collecting data on rainfall and temperature, which strongly influence leafspots on peanut, may be worthwhile because it can help you explain your results. But collecting data on soil physical properties does not seem to be related to the objective. It is useful to ask yourself, “How can this data be used?” If you have trouble answering that question, then collecting that data may be a waste of time. **It is much more common for people to collect too little data than to collect too much data.**

Deciding what data to collect is only part of the process. You also have to decide when to collect that data and if you need to collect the same type of data on more than one occasion. For example, in a nematocide trial, it is not sufficient to collect nematode population data at harvest; you must also collect data at planting to ensure that the plots started out equal. It is usually a good idea to collect nematode population data in the middle of the season because even in effective treatments nematode populations can sometimes increase to the level of the untreated control by the end of the season. The biology of the organisms involved will determine when and how frequently data should be collected.

So, how much data is enough? The answer is “enough data to fully address the test’s objective.” If you understand the biology of the organisms involved and how your data addresses the test objective, then you should be able to tell if you are collecting enough data.

You should take photographs of any differences among treatments that are easily visible. To most farmers, a picture is more convincing than a graph or data table.

COLLECTING UNBIASED DATA

It is critically important to collect unbiased data.

The only way to ensure this is to collect data without knowing what the treatment was in that plot. That would be difficult to do if the treatment were written on a stake in front of each plot. It is beneficial to use some type of code on the plot stakes so that you have to decode the stake number to determine what the treatment was. You can make up any code you like just so long as the person collecting the data cannot tell from the plot stake what the treatment was. For example, you can number the plots sequentially (1, 2, 3, etc.) and have a sheet of paper listing what treatment was applied to plot 1, plot 2, etc. When you collect the data, you write down your observation for plot 1 and later look at your list to see what treatment was in that plot.

If you know what treatment was in a plot, or which plots were the untreated controls, your evaluations (disease severity ratings, insect damage ratings, etc.) may inadvertently be influenced. Your subconscious may slightly increase the ratings for untreated plots and decrease it for the plots with treatments that you think should work well. You will probably not even be aware that it is happening, but these subtle influences can change the data enough to affect your ultimate conclusions from the test. **If you do not collect unbiased data, you cannot be certain that your conclusions are correct.**

STATISTICAL CALCULATIONS

After collecting data from a properly designed experiment, you will usually need to analyze the data with appropriate statistical calculations. Statistical analysis may not be necessary if treatment differences are very large and consistent; treatment means may then be sufficient. Statistical analysis is beyond the scope of this publication. Proper statistical analysis can be done if your experiment was designed according to the principles outlined in this publication, but proper analysis can be complicated greatly if these principles were not followed.

It is probably best for you to seek help in making statistical calculations. If your experiment was properly designed, Extension specialists and other scientists may be willing to help you with the statistics if you involve them early in the process. They can also check your proposed design for flaws and omissions. If you want to do the work yourself some simple statistics can be calculated by hand, but most people will make the calculations with the help of computer software. Specialized statistical software is available, but most spreadsheet software can calculate simple statistics.

SUMMARY

The following checklist can be used in designing an experiment. These items may be addressed in any order.

- Determine the objective of the test.
- Select treatments to address the objective. Consider including positive and negative controls.
- Determine what data should be collected, and when it should be collected, to address the objective.
- Select the number of replications to use. Consider four replications a minimum.
- Determine how big individual plots will be.
- Select an experimental design.
- Determine how blocks should be arranged in the field.
- Randomize treatments within blocks.

Properly designing and implementing a field trial may seem complex the first time, but it is really a logical process that should not be intimidating. You may need help the first time you design a trial to ensure that you are not overlooking something important, but if you learn the principles involved in the process, you should quickly gain confidence in your ability to conduct experiments on your own.

