Research Guide for Farmer Managed Trials on Organic Farms

J. Diane Knight, U. of S. Dept. of Soil Science

Steve Shirtliffe, U. of S. Dept of Plant Sciences

Farmers are constantly experimenting; trying out new products, new equipment and new management styles. You can answer many questions about your farm by following the basic scientific method; developing a hypothesis (the question), testing the hypothesis (the experiment) and drawing conclusions based on the data (was the practice economically viable?). On-farm research must be practical for the farmer by using field-scale plots that can be managed with standard farm machinery. It needs to be conducted in such a way that factors like field variation will not discredit the reliability of the results. Successful on-farm research begins with thorough planning before going to the field.

The purpose of this guide is to empower farmers to:

1) Quantitatively evaluate agricultural practices
2) Produce results that are considered credible by both agronomists and other farmers

1. Forming a hypothesis / Asking the question

The hypothesis is a testable statement that forms the basis of the experiment. The question needs to be simple, and have a definite quantifiable answer. For example, a farmer might be interested in using cattle manure for improving soil fertility. What (s)he might really be wanting to know is if manure increases crop yield. However, the question cannot be answered for every crop at the same time nor with every type of manure at the same time. The farmer must narrow it down to a specific crop or crops and a specific type of manure. We will use wheat and cattle manure. So the general interest in using manure to improve soil fertility, now becomes a specific question: "Does the application of cattle manure increase yields of a wheat crop?" Just formulating the question brings to mind a number of decisions that will have to be made when designing the experiment, such as the rate of application and how the manure will be applied.

2. Testing the hypothesis / Planning the experiment

Testing the hypothesis involves figuring out what treatment(s) you will apply, deciding what you will measure and how, and planning how you will set up the experiment in the field. Remember simple is better. The project must be manageable, both in size and the amount of time it will take to prepare the site, seed, take measurements, harvest and anything else that needs to be done to properly maintain the trial.

2.1 The experimental design

The first step in the experimental design is to determine if the experiment you are doing requires a yield comparison. Many experiments done by farmers are done to find out if something works in their particular system. For example, trying out a new piece of machinery to see if it works in their soil does not require a complicated experimental design; you just drag it around the field and check behind! Likewise, the effects of some treatments may be so obvious that the producer might think that (s)he can see the effect without using a complicated experimental design. Be careful here. Making comparisons between fields that have had different practices applied to them can lead to mistakes. Differences in crop rotation, soil types, rainfall, etc., all conspire to make one field different from the other. How often have you had one field of a particular crop yield much higher than another field with the same crop even though you farmed both fields exactly the same? Imagine if you would have put Brand X on one of the fields. Now assume that in reality Brand X has no affect on crop yield. However, depending on if you had put Brand X on the low yielding field or on the high yielding field, you might conclude that it ruined your crop or gave you much better yields. To properly compare treatments that affect yield you must follow a more complex experimental design.

There are three components that all experiments must have: a **control** (a strip that you do not apply the input or management practice to); **replication** (testing the treatment against the control more than once) and **randomization (**to account for natural variation). Without these three factors any data you gather is essentially meaningless. Setting up an experiment in a structured way helps to reduce the amount of variation due to soil texture, fertility, weed populations, precipitation etc. in your experiment. Replication, randomization and the use of controls also allow you to do statistics on your data.

There are a number of experimental designs that researchers use. We are going to discuss a **paired comparison** design because it is the most practical for farmers to use and analyze. The paired comparison involves making comparisons between an untreated control and a treatment. Only one treatment is compared at a time in a particular experiment. In our example examining cattle manure application on wheat yields, we would select one constant application rate of the manure to test, and one method of application (broadcast and incorporate for example). (**NOTE**: If more than one treatment is compared to a control in the same experiment, the design becomes a **randomized complete block** design, and is a bit more complicated to analyze statistically).

We will also discuss several different ways of replicating the treatments: 1) where all of the replication is done within a field (within field comparisons); 2) where treatments are replicated on different fields on a single farm (between field comparisons); and 3) where a number of farmers get together and put the same treatments on several farms (another type of between field comparison).

2.1.1 Paired comparisons within a field:

This is the most complex type of paired comparison. It will be discussed first to give you an idea of why treatments need to be replicated and randomized and how this can be accomplished on your own farm. Once these ideas are introduced to you we will talk about ways of simplifying the experiments using between field paired comparisons.

2.1.1.1 Replication:

As you know yields vary from year to year because of different amounts of rainfall, temperatures and many other factors. Yield also varies from location to location within a field because of hills and depressions, fertility, variations in soil type, etc. This natural variation is why replication is essential in experiments.

Figure 1 shows actual yields of wheat strips harvested from a field that was treated the same way all the way across. There were no treatments applied. We know this for a fact because it was Steve Shirtliffe who collected this data! Notice that there is a slight variation between the yield in each of the strips, however it doesn't look very different. There could be many reasons for this, differences in how strait the swather was driven (remember these were the days before auto-steer), but probably the differences are due to small differences in the soil nutrients, water or pests (weeds, diseases or insects).

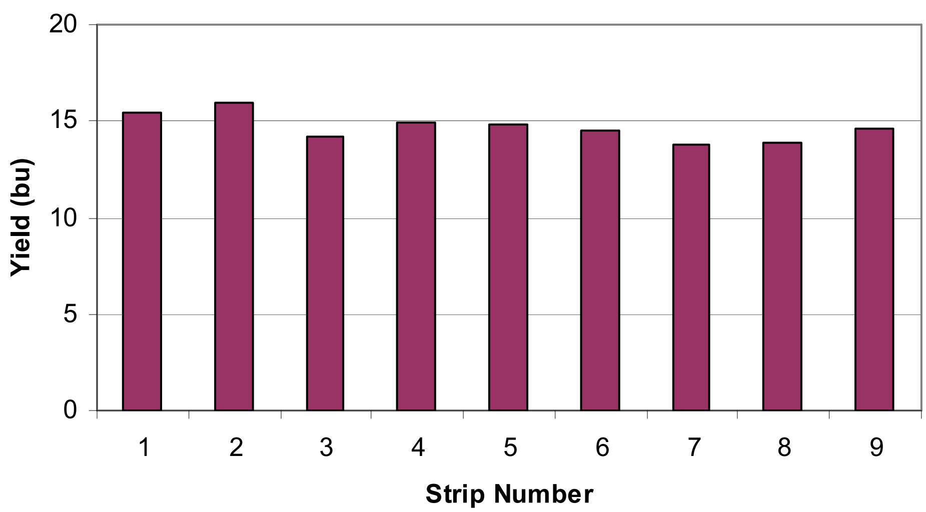

Figure 1. The yield of wheat strips in an uniform field. The strips were swathed then combined and weighed in a weigh wagon (Shirtliffe, 1997 unpublished results)

This all changes when you look at the percentage difference between one strip and its nearest neighbor strip in **Figure 2**. Notice that four of the eight pairs of strips have over a 5 % difference with their neighbor and that there is over a 10 % difference in one case. Now imagine if the ineffective Brand X had been applied to strip number 2 and the control had been applied to strip number 3. You would have incorrectly concluded that Brand X caused an 11 % yield increase. There are only two comparisons of the eight that you would have reached the correct conclusion that Brand X had no effect on yield. In order to account for the effect that this natural variation in the field has on your treatments we compare more than one strip and remove the random variation. This is what statistics do, they eliminate some of the uncertainty when comparing one treatment to another so that you can be confident that the difference is due to the treatment and not just chance variation.

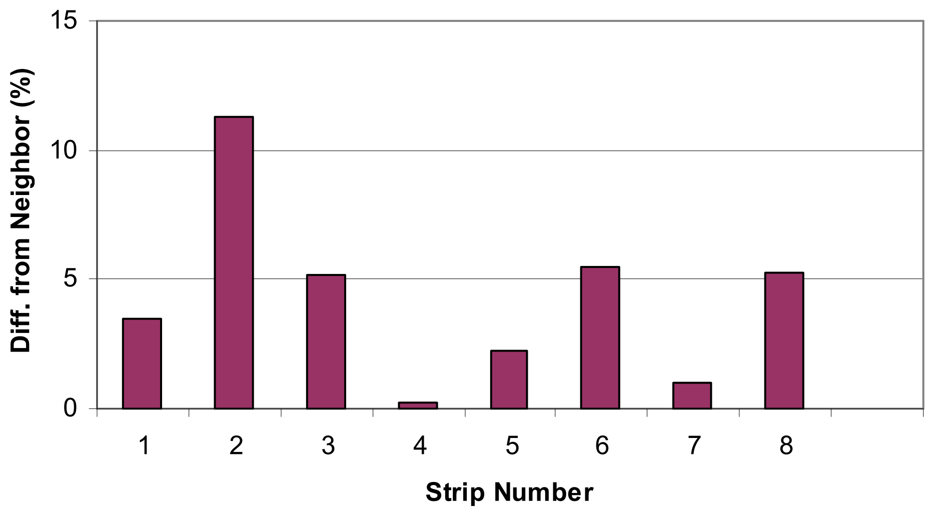

Figure 2. The absolute percentage difference in the yield of one strip compared to its neighbor. For example strip number 1 has a 3.5% lower yield than strip number 2. Note: strip 9 can not be compared to strip 1 as it is not a neighbor. (Shirtliffe, 1997 unpublished results).

The more times a treatment is repeated, the more likely it is that the measurements will reflect the effect of the treatment rather than the effect of natural variation in the field. However, there is a point of diminishing return after increasing replications doesn't give increased accuracy. As a rule of thumb for on-farm research, between four and six replication in long strip comparisons is a good number to use.

2.1.1.2 Randomization:

Randomization means that the order of the treatments cannot be the same in every replicate. **Figure 3** illustrates why.

Downhill/moist — Uphill/dry

Manure	No manure	Manure	No Manure	Manure	No manure	Manure	No manure
Rep. 1		Rep. 2		Rep. 3		Rep. 4	

Downhill/moist — Uphill/dry

No manure	Manure	Manure	No Manure	Manure	No manure	No manure	Manure
Rep. 1		Rep. 2		Rep. 3		Rep. 4	

Figure 3. An example of treatment randomization. Top example – not randomized Bottom example – one randomization pattern.

In the top example (not randomized) in every pair of treatments, the manure treatment will perform better than its no-manure partner because it is on a moister part of the field. As we move from left to right the yields for both treatments will probably be lower because the amount of water available drops, however the manure treatment will always yield higher than the adjacent control treatment in a particular rep.

In the bottom example (randomized) this problem is avoided, by randomly placing the manure and no manure plots. The manure plot is not always located on the moister land. A simple way to determine the placement of each treatment is to flip a coin or draw slips of paper from a hat, having assigned a specific treatment to one side of the coin or written on the slips.

Ideally, the field that the experiments are to be located on should be as flat, and uniform as possible. If you know that one field has a salt problem in one area, another field should be selected or at the very least the saline area avoided. However, selecting a uniform field does not substitute for randomization. It is impossible for us to know if a field is uniform for all factors (fertility levels for major and minor nutrients, soil texture, gravel patches, depth of the ground water table, to name just a few examples) that can affect the crop.

Sometimes, depending on the area your farm is located in, it is not possible to find even a relatively flat field, so that there will always be a water gradient across the field. In this situation, locating the plots parallel or following the slope rather than across the slope, will help reduce the bias associated with the gradient. **Figure 4** illustrates this. In the bottom figure all of the treatments are exposed to the range of water levels along the slope. All of the replicates do not have to be located side by side. If there is a small saline depression in the field you want to use, simple avoid this area. Locate the replicates on either side of it.

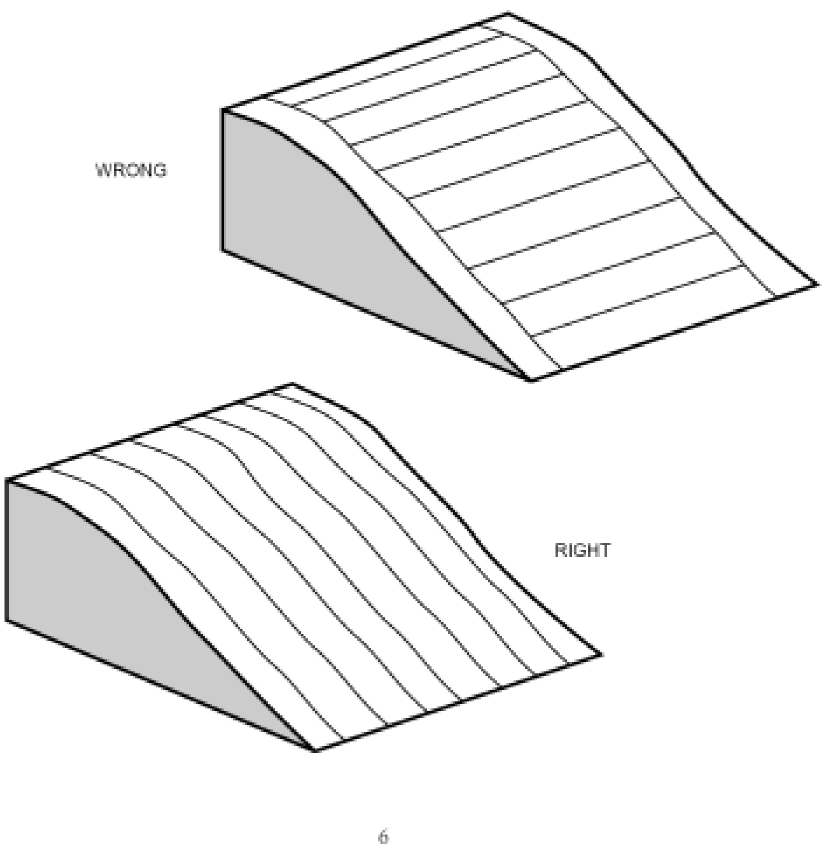

Figure 4. How to position plots on a slope (from D. Anderson, On-farm Research Guidebook, University of Illinois) .

A history of the field is also useful. Avoid areas that may have been treated differently than the rest of the field in the past such as old fence line, sod strips, building sites and manure piles. It is not uncommon for manure to be applied to only small areas in a field. Avoid these areas for putting experimental plots down for at least three years. In addition avoid the headlands of the field as this land is usually more uneven than the centre of the field. Past management can strongly influence present performance and can provide valuable clues as to why things turned out the way they did.

Drawing a Plot Plan:

Once the treatment and field is selected it is important to draw up a detailed plot plan. These should indicate the size of the plots relative to the field equipment that will be used to seed, harvest etc., the individual treatment(s) and their location in the field, as well as the number of replicates of the treatments and their location on the field.

A plot plan will make it possible to determine the number of stakes and markers required to adequately locate the treatment areas. It is very important that you are able to go back to the field and locate exactly where treatments and controls were applied.

Ideally, the width of plot strips should be several feet wider than harvesting equipment so that a continuous full width cut is possible. Enough distance should be left between treatment strips to allow efficient harvest of the border strips that remain after the main plots are cut. For example assume that strip plots will be harvested with a 20 foot swather. Leaving a total treatment width of 35 feet leaves about 15 feet between the harvested plots which is enough to allow efficient harvesting of the border strips. Indicate the plot widths and lengths and border strips that are appropriate for your equipment on the plot plan. If your treatment does not allow you to treat a large area you can use smaller plots. In this case you will have to go in and hand cut some squares of crop out to be hand threshed and weighed. The local Sask Ag employee can make arrangements for transporting these into Saskatoon for us to hand thresh and weigh them.

2.1.2 Between field comparisons:

Another way to deal with having non-uniform fields and to make an experiment easier to manage is to replicate the experiment on different fields on your farm or else to have a number of farmers in the area place the same treatments on their farms rather than replicating all of the treatments on the same field.

In this case you simply have one part of the field with the treatment and an adjacent area without the treatment (the control). The benefit of this method is its simplicity; it is relatively easy for a farmer to apply a strip of a treatment in a field. However, as always there is a catch. In order for you to be sure of the results, you must repeat the strip in at least three locations (preferably six). The more often the pairs that are compared are repeated the more confident you can be that the results are real. So if some organic farmers in a chapter get together and decide that everyone will run one strip on their farm it is fine. However, if you just run one strip you can't really say anything.

As mentioned, to accomplish this type of experiment you have to conduct it on at least three fields (preferably six). Run the treatment crosswise to the normal direction that the machinery travels in the field. Put stakes up in the ditch at each end so you know where the treatment ends and the control begins. By using the results from different fields or farms we will be able to calculate if there is a real difference in the treatments (remember the Brand X example?).

3. Conducting the experiment

3.1 Plot layout:

Once you have the plot plan developed it is time to go to the field and measure out the plot dimensions. Strip plots should be parallel to the field edge, but far enough from the edge to allow passage of the farmer's equipment. Measure the location with a surveyor's chain or a "walking stick" and record the measurements on the plot plan. Install flags to guide the tractor operator and stake to provide permanent markers. Pin flags (surveyor's flags) should be removed after the field is seeded or at least before harvesting, as pin flags can get caught up in other equipment. However, they are very useful at seeding because the tractor can drive over them.

The next step is to apply the treatments. This will vary depending on the actual treatment. In our manure example, the manure should be applied as uniformly and accurately as possible. Make sure you have enough manure to apply the same amount to each of the plots. The manure should all come from the same source, and be of the same age as manure from different sources will have different amounts of nutrients in them. This will introduce more unnecessary variation into the experiment that could mask the real results. The most important factor when applying the treatments is to treat all of the plots exactly the same except for the actual treatment. This means that all of the seed is from the same source; any other amendments are applied to all of the plots (controls and treatments) in the same manner, at the same rates; seeding rates and seeding depths are all the same. If after the crop has been seeded it is necessary to harrow for weed control, all of the plots must be harrowed. It is impossible to anticipate everything that will be done to the plots in terms of maintenance, but remember to treat everything the same. The only difference between your control plots and your treatment plots is the specific treatment. This makes it possible to attribute any yield increases in our example to the manure treatment, rather than some other factor.

3.2 Taking notes:

Organic producers are experienced record keepers, so this part should be easy. Keep a notebook dedicated to the research project. Record any disease, insect or weed problems. You might be able to see differences between the treatments. One concern with manure applications is that uncomposted manure can contain weed seeds. If your manure is not composted, or not properly composted an increase in weeds might cancel out any improvement in yield due to the extra nutrients applied. Meticulous note taking will make this easier to determine. It is surprising how quickly differences that stand out in the spring may not be noticeable, or remembered, even a few days later, so make sure you write your observations down right away. You should also document the weather: rainfall, temperatures (particularly if they are out of the ordinary) hail, windstorms etc.

3.3 Making measurements:

In research, a small part of the farm is controlled and measured in order to make projections (or predictions) about the how other parts of the farm will respond to a treatment. Once you have the experiment established there is not limit to the kinds of data you can collect. Often numbers of plants emerging and when they are emerging is of interest. It is impossible to count all of the plants emerging in all of the plots. Instead do a subsample. You could count two or more one metre long rows in each plot for example. These rows and locations within the rows should be located randomly for the same reasons discussed above. Be sure that what you are measuring will be useful in answering your research question. Most farmers are interested in yield and quality factors.

3.4 Harvest:

As with everything plan ahead. Measure yield from each plot separately. Do not combine all of the treatment samples and control samples into one bulk sample, weigh all plots individually. This can be done with a weigh wagon, yield monitor or by cutting, bagging and labeling individual plots yields and carrying them to a scale for weighing. A good way to hand harvest the plots is to harvest 10 one metre long rows located throughout each plot. You could harvest 5 rows at one spot and 5 rows at another. Put the samples in paper grocery bags that have been labeled with the replication number and the treatment. Don't pack a lot of green material in the bag or else it will start to rot. You will probably need more than one bag for each plot. Leave the bags in a dry place. We will make arrangements for researchers to thresh and weigh these sub-samples for you. Don't forget to write down how big the area was that you cut from each plot!

Keep track of which yields come from which plot. Note the harvest area from each plot so that plot yields can be converted to pounds per acre. Avoid including the border strips in your harvest sample.

If you don't want to go through all of this trouble you can send your yields to us and we will calculate if there is a statistical difference between them. We are all set up with computer programs to do the work for us.

4. Statistical analysis

An example of a paired comparison trial and its analysis are attached. It is actually fairly simple to perform.

If you don't want to go through all of this trouble you can send your yields to us and we will calculate if there is a statistical difference between them. We are all set up with computer programs to do the work for us.

5. SUMMARY

1. Decide on a question that you want an answer to. Keep the question simple as complex questions are difficult to address.
2. Design an experiment to address the question. Remember you must include an untreated control in order that you can compare the effect of your treatment. Be sure to stake out the field and record where your plots are so you can measure yield separately in these areas. This is especially important if the effects of the treatments may last more than one year.
3. Replicate your experiment. If you are using a unreplicated paired design, the test must be conducted on other fields or in other years so that you can be confident that the results are true. If you are using a replicated paired plot design it the test does not have to be done on other fields or years but this will still increase your confidence in the results.
4. Collect data carefully. Be sure to record data on the plots. This may include doing emergence counts and looking at the plots. Include the plots on your regular visits on the farm and watch them. When it comes time for harvest you must determine if there is a yield difference between the plots. The first way you can do this is by hand harvesting and drying a small portion of the plots (example cut off 10 rows of crop each 1 metre long). The bags should then left in a warm dry place to dry. The second way is to use a combine and record the weight of *each* plot.
5. Analyze the results. The simplest way to do this is to send us the yield data from each plot and we can calculate it for you. However if you are ambitious or good with computers you can follow the methods outlined here.
6. Congratulate yourself! You have become a farmer researcher and can now effectively evaluate treatments on your farm.

AN EXAMPLE OF A PAIRED-COMPARISON TRIAL AND ITS ANALYSIS

	Treatment				
Pair number (replicate)	Control	Manure	Difference (X)	(X-0)	$(X\text{-}0)^2$
1	117.9	118.9	-1.0	-6.62	43.78
2	125.2	112.0	+13.2	7.58	57.51
3	119.9	100.7	+19.2	13.58	184.51
4	114.9	110.8	+4.1	-1.52	2.30
5	116.9	119.9	-3.0	-8.62	74.25
6	119.7	118.5	+1.2	-4.42	19.51

n=6 0=5.62 (Mean difference) SS=381.85 (Sum of Squares)

Sample Variance: s_x^2=SS/(n-1) = 381.85/5 = 76.37

Variance of the Mean: $s_0^2 = s_x^2/n$ = 76.37/6 = 12.73

Standard Error of the Mean: $s_0 = (12.73)^{1/2}$ = 3.57 bu. Acre

(the ½ means take the square root of 12.73)

Student's $t_{0.05,\,5}$ Get the t from a t table. Almost all statistic books have a t table "0.05' is the chance of an error occurring at the 95% confidence level (i.e., 95% of the time the test will indicate a difference when there is a real difference). "5" is the number of pairs of treatments "n" minus 1, (6-1, in our example). The correct number for t is in the table at the intersection of the "0.05" column and the "5" row. You might want a higher level of confidence (99%) or you might settle for less stringent burden of proof (90% or 80%).

Least Significant Difference at 95% confidence level:

$LSD_{0.05}$ = s0 x $t_{0.05,\,5}$ = 3.57 x 2.57 = 9.17 bu./acre

In our example the observed difference between the treatment (x = 5.62 bu.) was less than the difference (9.17 bu.) that would occur simply by chance 5 times out of 100. That is, the difference observed was not statistically significant at the commonly used 95% confidence level. In this case there was no significant yield difference between the wheat that was manured and the wheat that was not manured. However, these results come from one field in one year. It's a good idea to repeat the trial elsewhere to see if the outcome is reproducible.