BRITISH FARMER GROWER

Farming pin-ups

Your free four-page practical guide

EDITION 9: CARRYING OUT YOUR OWN FARM TRIAL

Pull me out and pin me to your noticeboard

THE WORLD is changing fast. Weather is becoming more variable and new techniques or equipment or varieties continually appear. Often you may not be certain whether doing something in a different way will make a real difference to your farm, in which case trying out a new way of doing something in a simple farm trial makes good sense.

This month's *Farming Pin-up* aims to talk you through the basics of on-farm experimental trials. This guide was inspired by the Duchy Originals Future Farming Programme – an initiative that helps farmers improve productivity in an environmentally responsible way.

Whether it's trialling peat-free propogation composts, testing a spray programme on arable crops or varying the nutrition fed to your herd, this guide should show the fundamentals of where to start and how to measure success and failure.

WITHOUT CAREFUL THOUGHT, IT BECOMES IMPOSSIBLE TO TELL WHETHER ANY DIFFERENCE YOU SEE IS REAL

BACK TO BASICS

So, you decide to set aside a small area of land and test your new management practice. Your idea is to compare your returns from that area with another area managed in the old way, or to compare your returns with those of last year on the same land.

Unfortunately, that's not as simple as it sounds. The reason is that the world is extremely variable from place to place and from year to year. This inevitably means that comparing one patch of land with another, or one year's performance against another, even if nothing had changed, would give you a different return. Without careful thought it becomes impossible to tell whether any difference you see is 'real' and genuinely due to the management, or due to chance factors such as weather, soil or a micro-climate creating differences.

You have to design your trial in such a way that you can tell whether any differences that you see are due to a change in management. There are a few straightforward principles – that don't need any training in statistics – to help you design your own trial and make firmer conclusions.

BREAKING DOWN THE PROCESS

WHAT IS YOUR QUESTION?

The most important thing to think about first is exactly what question you are asking. The easiest trials to design, conduct and interpret are ones in which you are varying a single factor, such as a management technique, or a crop variety. The ideal for such a trial is to keep everything as close to identical as possible and compare areas that vary only in this factor. For example, you may want to see whether a new variety does well on your farm, so your trial would consist of planting some plots of the new variety and comparing its yield to your normal variety. Keeping everything as close to identical would mean ensuring it was planted on similar land, at the same time, in the same way, and managed identically. If any of these factors vary, any difference in yield may not be due to the variety but to these other factors.

HAVE YOU GOT A CONTROL?

To see if a change in practice is creating a difference in results, you need a baseline for comparison. This is usually called a control. Typically, a control would involve keeping everything as normal, and the results from the control area are compared with the experimental area where the new thing is applied. It is possible to compare one year with another but this is more difficult as many things vary over time, so before and after trials are usually more difficult to interpret.

Some experimental designs require two controls. For example, a new method for weed control may require more tractor passes during the growing season. If you found a difference in yield between the old and new methods of weed control, you wouldn't know if was due to the method of weed control or the difference in the number of passes. To answer this, you would carry out one trial in which you compared results from your normal method to those from the new method and a second trial in which you compared the yield from plots which had received the same treatment, but differed in the number of passes over the ground.

REPEATING YOUR EXPERIMENT

The idea of replication is to repeat the experiment in many different places or over many years so that any differences in results are averaged out and you can tell whether the factor you are measuring is having an effect. As a rule of thumb, the more replication you do, the more likely you are to see a difference due to your treatment – if one exists.

Replication has nothing to do with the amount of data you collect. A replicate is an independent repeat of the control or the treatment. As a simple principle, the more you do control and experimental plots side-by-side (called a block), and the more blocks you have in different fields the better. You can carry out replicates within the same field, but you will have less confidence that your results will apply to another field, so the more fields you use, the better, as any factors specific to a particular field will be averaged out.

The size of a test depends on the question you have set. For example, if you are trialling a new 12m boom sprayer, then you will need 12m strips. If you are measuring the yield from different varieties of carrot you could hand sow a single metre square. In general, it is better to have more, but smaller, plots than fewer larger ones.

MAKE IT RANDOM

Randomisation goes hand in hand with replication because randomisation is a technique to make sure the replicates are independent from each other. To be independent means that each replicate is not connected and so can genuinely give you fresh information.

Often the simplest thing to do may not ensure independence. For example, it may be easier to place arable trials on the field edges, or livestock trials in a low-yielding field which isn't so valuable to your farm enterprise, in which case the results may not tell you about your fields in general. The best way is to pick plots at random. This needn't be too tricky: throw darts at the farm map and then draw your plots around where they stick! If you are doing control and treatment plots side-by-side, randomise them to left and right to ensure that the control plots are not always up-slope from the treatments.

THE BEST WAY IS TO PICK PLOTS AT RANDOM

WHAT ARE YOU MEASURING?

If you have designed your trial with a simple control and treatment block, replicated in a number of places around the farm, the rest should be easy. If you are interested in yields, but have no yield monitor, you can sample yields across your plot. With cereals, for example, 10 samples of ears from a 10cm x 10 cm square collected at random points across your plot, dried for a few days, and weighed will give you an accurate estimate.

Pull me out and pin me to your noticeboard

ANALYSING YOUR DATA

Data can be examined in many ways. A simple bar chart showing the yields of the control and treatment plots can be very informative. Often it is not necessary to carry out any statistical analysis as one can get a feel for whether the results show a real difference, and were not simply due to chance. Imagine you have five replicate pairs of control and treatment. If nothing was going on you'd expect that the results for treatment to be bigger than the results for control about half the time. However, if all the treatment results are larger than the control, it would be like tossing a coin five times and getting five heads. Statistically, you have a one-in-32 chance of getting five heads in a row. Therefore, you would be confident that the treatment was having an effect. If four of the five treatments were larger, you'd be less confident as this would be a one-in-six probability of this happening by chance. In statistics, the all important threshold is a probability of one-in-20, above which you have reasonable certainty that the treatment is significantly influencing the results.

WHAT CAN GO WRONG?

Things always go wrong, and very often the results from designed experiments are far from conclusive. The most likely reason for this is because the treatment is having a smaller effect than natural variability, for example, the treatment increases yields by five per cent, but yields fluctuate naturally by 15 per cent from place to place.

There are three ways round this:

- Make plots bigger, as a large plot will average out more small-scale variability
- Increase the number of plots
- Increase the size of the treatment, and with it, the chance of its effect being seen. For example, if investigating the impact of biochar, try adding more.

Another thing that can go wrong is that some trials don't provide data for reasons outside your control, for example livestock become sick, plots get flooded, or a crop gets trampled by a herd of cows. Very often the remaining data is harder to understand but it is not wasted. Seek advice because someone with knowledge of statistics may still be able to assess the 'signal to noise' ratio and give an indication of whether the results indicate something.

GETTING HELP

There is a huge amount of writing on the design of experiments, but much of it is written for people studying statistics, so seek out some the non-technical material that is available on agricultural teaching websites. The most useful statistical tests are pretty standard rather than specific to farming, so if you know someone working at your local university, they should be able to put you in touch with someone happy and able to help. There are also free apps too that can be installed on smart phones and used to record and analyse data. And, of course, spreadsheets will allow you to analyse your own data.

ON-FARM TRIAL CASE STUDY

HOW DO DIFFERENT LEGUME LEYS AFFECT YIELDS OF CEREAL CROPS?

Given the rapid leaching rate of the free-draining Cotswold Brash on which Barrington Park in Burford, Oxfordshire, is situated, farm manager Adrian Dolby wanted to know more about which ley species or species mixtures work best to build and sustain fertility in his fields. He carried out this trial on his farm between 2005 and 2008.

Which legume ley is best for sustained fertility?

Method

- In 2005, Adrian drilled one field with five-hectare strips of four different leys: white clover and chicory; white sweet clover; red clover; yellow trefoil. All leys were drilled on the same date and managed in exactly the same way.
- Leys were incorporated after 24 months.
- In 2008, Adrian drilled spring barley, spring oats and spring wheat in strips perpendicular to the ley strips such that each cereal was tested on each ley type.
- Adrian harvested the cereals in blocks corresponding to where each ley had been.

Measurement

- Yield of cereal crops from each block.

Findings

- Crops following the red clover ley tended to produce higher yields of each cereal

A word of caution

- The bar chart shows that for each crop, the red clover ley provided the highest yield. However, all the red clover data came from one strip of the field and the result could be down to chance if this was somehow better ground. Replication would give more confidence in the results – either in smaller plots within the field or in different fields.

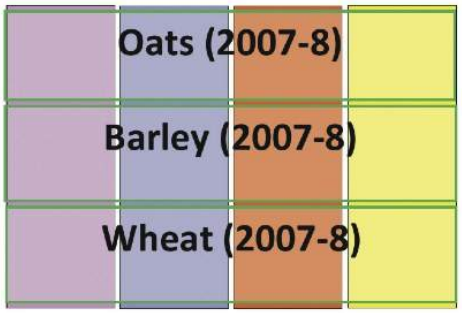

Ley species (2005-2007)

- White clover + chicory
- Red clover
- Yellow trefoil
- White sweet clover

Trial design

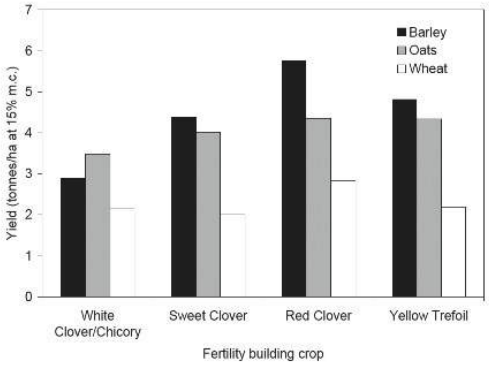

GRAPH: Grain yield of spring barley, spring oats and spring wheat following different fertility-building leys. Data from Parson's Piece, Barrington Park Estate, 2007.

WITH THANKS TO...

This *Farming Pin-up* was developed by Professor Tim Benton of the Global Food Security Programme and the University of Leeds, the Soil Association and the Organic Research Centre, with the support of the Duchy Originals Future Farming Programme. For more information, go to www.soilassociation.org/fieldlabs. Thanks also go to Adrian Dolby for sharing data from Barrington Park.