Research and extension services

A THOUGHT EXPERIMENT IN BENCHMARKING AND OPTIMISATION

Simon Ward
Increment Limited and The Policy Group

Increment Limited
11 Margett Street, Cambridge, UK CB24 8QY.
Email: Simon.Ward@increment.co.uk
+44 1954252859

Assistance and encouragement was provided by Mark Reader, Department of Land Economy, Cambridge University, Cambridge, UK
www.landecon.cam.ac.uk/directory/mark-reader

2448 words

Paper is applied (not academic)

A THOUGHT EXPERIMENT IN BENCHMARKING AND OPTIMISATION

Abstract

Agricultural benchmarking and input optimisation are ripe for further development using tools that are now freely available for manipulating larger volumes of data. These new systems will not replace existing benchmarking methods but, used carefully, add another dimension to help improve profitability. These techniques are beginning to merge systems used to derive agronomic optimums, such as random block replicated trials, and farm management analysis based on use of comparative data to relate inputs and outputs to target metrics, such as profit. However, mapped data while providing considerable scope for improving farm management practice is often badly analysed allowing incorrect conclusions to be reached. Data analysis is not a substitute for innovation.

Keywords: optimisation, benchmarking, farm, data, management

Introduction

There are multiple steps and levels to benchmarking and optimisation.

1. **Measure.** This may be physical or financial and must be related to a target objective. If an action cannot be measured it is almost impossible to improve.

Once an input or output has been measured, the data has to be related to a quantifiable objective.

2. **Examination.** In practice the largest benefit from benchmarking often occurs simply from questioning the data. The act of asking whether the job could be done more efficiently frequently results in improvement.
3. **Comparison.** The simplest benchmarks comparing prices received or paid can provide the most assured benefit and at least appear straightforward. However, even these measures need adjustment to take into account time of purchase or sale and other differences such as location. If judgement rather than luck is responsible

for the difference, it may be necessary to average over a number of years. Nonetheless there is a reasonable chance that having identified that too much was paid for the fertiliser, or too little received for the grain, action can be taken.

Beyond this, optimisation has been based either on replicated trials to produce a response curve, or by comparison of farm data.

Replicated trials can provide an extremely accurate estimate of the input optimum and provide a measure of certainty. However, they are expensive, provide information on a narrow range of conditions and statistical significance is (usually) based on a symmetrical distribution curve potentially resulting in the loss of valuable improvements. Replicated trials have been key to UK agronomic advancement but traditional randomised block trials with three or four replicates are inefficient in determining more complex optimums with multi-factor interactions or even identifying that there are multi-factor interactions.

Replicated trials are almost impossible to apply to some inputs. Increasing machinery cost with other variables held constant is not a realistic proposition. Expecting a meaningful assessment of the impact of increasing, say, accountancy cost for a specific business to find the optimum, is clearly ridiculous.

In farm management analysis comparison is with the most profitable businesses, often based on adoption of conventions such as cost (or use) per unit of the limiting factor. In the UK this is usually area but has also been per litre of milk where milk quotas limited production until March 2015 i.e. the cost of milk quota per litre was higher than the cost of land to produce that litre. The other inputs are cheaper to adjust. Other comparisons use the relationship of the input cost with output and cost per unit of production.

Again many businesses have made considerable strides through comparison. Comparison may be with the top 25% or with the most successful business (frontier analysis). But there are problems:

- Observed data relationships may not be causative. The average application of an input by the top 25% farmers may be low (or high) compared with the average of the whole sample but this may not be causative or an adjustment key to improving profitability. In practice there is rarely sufficient data to demonstrate the relationship between input and target metric and even more rarely any indication of the statistical strength of the relationship. Where the data does exist, it is often presented in such a way that it cannot easily be made use of by farmers or their advisers.

- Some constraints are implicit such as limits imposed by weather or soil type so however much adjustment is made the business still operates at a suboptimal level. In the UK it is sometimes found that the most efficient businesses are on the worst soil types since the precision has to be better to survive. For example, where profitability is low, it may not be possible to retain labour for sentimental reasons or in order to preserve a particular lifestyle. An Australian arable farmer or a New Zealand dairy farmer is likely to be far more efficient in physical terms than a UK farmer – despite higher profit in the UK. National factor productivity analysis tends to show the greater improvement in factor productivity in those countries with least subsidy and that there is even greater improvement in countries such as Argentina where export taxes actively reduce farm profitability relative to other sectors. It is not possible to prove a causal relationship.

- The quantification of response available through statistical trials (the point where marginal cost equals marginal return) is lost in the conventional comparative analysis.

- Averaging of data for some farming systems may result in the description of a system that does not exist. For example, in the UK both high yielding-high input dairy systems and low yielding-low input dairy systems may be profitable with the interim systems being less profitable. The average of the two profitable systems describes a less profitable business and, as importantly, so does the average of the top 25%.

- While effort is made to group farms into similar bands the grouping may not distinguish impacts of variables that are common to several of the farm groupings.

- Where frontier analysis is undertaken there is a good chance that either a unique factor is involved or more likely there is an error in the data. As much time can be spent understanding the comparator data as improving the target business.

Another key issue is the ability to identify how easy it is to change a sub optimum input or to establish the relationship between two or more variables.

The Analysis

Most benchmarking was developed when spreadsheets were rudimentary. Modern spreadsheets and databases allow considerably more data to be analysed, statistical analysis to be undertaken quickly and data to be graphed to provide a preliminary assessment. Larger data sets provide even better understanding, allowing assumed normal distribution curves to be replaced with the actual distribution curve for the relationship, removing a source of error in replicated trials.

The simple scatter graph allows the relationship of an input with the target output to be explored. Not only can the relationship be demonstrated but each input can be explored to determine whether the relationship is:

- Positive, negative or neutral (e.g. profit with crop yield (in general))
- Stepped (e.g. profit against labour or machinery (changes by whole units))
- Describes a production function with an optimum or minimum within the sample range. (e.g. fertiliser input)
- The statistical strength of the relationship
- The likely gain that the subject business could potentially make.

The shape of the relationship may also indicate whether there is a confounding interaction between two or more inputs. It also ensures that the averaging of two data points, or farm systems, to produce an average that fails to describe any system does not distort the action required to improve profit. The larger the data set the easier it is to test subsets to identify other variables historically excluded from both the farm management analysis and replicated trial.

A correlation of the database components may stimulate questions and allow relationships to be explored.

Linear multi-regression analysis is no longer beyond the average consultant and can be used to test concepts within the standard farm classifications. Use of crops, or crops grouped by similar characteristics, may allow costs to be allocated more accurately than was historically possible. An exercise carried out in the UK showed higher labour costs were associated with spring combinable crops than winter crops, despite a lower number of hours needed for their production. (The explanation was believed to be that spring cropping increased the requirement for full time labour in place of seasonal labour). In the case of farm scale, different inputs will show different rates of increase or decrease. A good linear regression fit may help estimate the impact of the stepped changes that can be made to the subject business.

Having established that there is a relationship the construction of frequency distribution graphs can also help prioritise management action. Where most farms may show inputs around the optimum, while the subject business is some distance from the optimum, the effort needed to adjust the business to achieve that optimum is likely to be small. If the deviation from the optimum is outside the operator's control (e.g. weather or soil type) it also indicates that more radical adjustment may be appropriate.

Conversely if most of those in the sample are a long way from the optimum, improvement is likely to be more difficult to achieve, suggesting effort should be put into easier targets.

Mapped data

The growth in mapped data has opened up as many questions as it has solved. For many it has been too easy to perceive relationships that do not exist and lack of rigorous analysis has held back progress. Analysis of the information produced for GIS/GPS systems provides a further link between farm management data and analysis and agronomic data and analysis. Ultimately there should be no reason to separate the two as far as analysis is concerned.

Large scale mapping data, linking physical and financial characteristics, has an important role in distinguishing between those factors within and outside management control.

Many of the big questions that should be soluble from these new data sets have still not been resolved. For example nutrient trials derived from a combination of conventional application trials, soil analysis and known crop requirements are inevitably crude. In practical terms it was never possible to carry out trials for all crops, on all soil types in all weather conditions but analysis of farm data should make this possible or prove lack of relationship. In fact the return on investment from phosphate application is often very low in the UK, despite widespread application. It has recently been reported that as little as 4% of the TSP fertiliser phosphate applied is taken up by the UK cereal crop so a low response in many situations is perhaps unsurprising[1].

[1] https://cereals.ahdb.org.uk/media/1362816/T2F-Mar-2018-Progress-in-the-phosphate-puzzle.pdf and https://pubs.acs.org/doi/pdf/10.1021/es501670j

In theory at least it should be relatively easy to solve. Soil samples are increasingly taken from precise areas of the field with coordinates provided by GPS. Several service providers collect 10,000 or more samples in a year. Many of these have records of application and crop yield data.

Using the techniques described above the data points can be aligned and graphed. A number of theses can be tested until the best linkage is identified. Thus the initial hypothesis might be:

Yield = x(soil nutrient level) + y (fertiliser application)

Subject to the maximum yield when another component becomes limiting

If the fertiliser application is constant (as has historically been the case) there should be sufficient variation of both yield and nutrient in most UK soils to explore the relationship. Providing other factors (such as weather) are constant or variation is insignificant compared to the variable under test, a scatter graph would quickly show whether there was a relationship and at what point other factors tended to limit. This assumes that the application is not so high relative to requirement, that it swamps soil differences. Double peaks or limits would suggest there might be secondary interactions.

Where farmers were prepared to work together three rates of application (with one at zero) would go a long way towards helping to obtain more information and improving future recommendations and more importantly allow the deviation in yield to be plotted against soil and fertiliser application. In contrast to a conventional trial it is possible to describe the response distribution for each change of fertiliser rate, soil phosphate level or combination of the two.

Lack of relationship is as important as finding one. In agriculture the extra profit very often does not represent sufficient return on technological investment to be worthwhile.

But care is needed. While maps may show a convincing picture, for many data sets the image is little more than an informed guess based on extrapolation from the data points. In some cases the additional boundaries presented to the producer create more over- and under- applications than would have occurred in a field-based analysis and possibly a lower return on investment. This does not occur for systems such as soil conductance measurement and may be reduced by better identification of infield soil boundaries.

Analysis is probably best broken down into the sample points providing greater accuracy. But it is highly unlikely that two measurements, such as soil nutrient and crop yield will have the same sample point so a decision has to be taken on the area that a data point from one layer may represent so that it can be compared with the data from a second layer. There are also likely to be dominant impacts from non-subject variables such as shading or rabbits.

Combine harvester yield is likely to be one measure and this creates a number of errors. Width is large averaging data over a large area, accuracy is low, there is a lag between crop entering the harvester and reporting of measurement.

Conclusion

Historic benchmarking and farm optimisation has provided considerable help to farm business managers but leaves a number of problems unresolved.

New techniques are available to improve farm optimisation and benchmarking. The study illustrates a possible framework using systems available to nearly all academics and consultants. The technique allows the convergence of farm management and agronomic methods to make a substantial improvement in overall business management.

Statistical analysis of mapped data sets is poor and many farmers and consultants are drawing incorrect or sub-optimal conclusions.

However, benchmarking is not a substitute for innovation to find a better management practice than others have achieved. To do the same as the top 25% should not be the objective although it may be a good start for those outside this banding.