COOPERATIVE EXTENSION SERVICE
University of Kentucky – College of Agriculture

LEXINGTON, KY 40546

Department of Plant and Soil Sciences

Soil Science News & Views

Vol. 26, No. 2, 2006

On-Farm Field Research: Replicating Your Valid Comparison

John H. Grove

On-farm research is motivated by a desire to learn more about a product/practice/system on land you manage. You may now have the tools (yield monitor, weighing grain cart, etc.) to accomplish on-farm research that generates information you can use in management decision-making. Your on-farm research should start with the design of a "valid comparison", according to your research objective (Soil Science News & Views 26:01).

Each comparison generates a single yield for each treatment, and thus providing a single observation of the yield difference between any two treatments. For example, you set up two strips comparing fertilizer nitrogen (N) rates for corn, treatment A being your usual rate and treatment B being your usual rate minus 40 pounds N per acre. At harvest you measure yields in the two strips and find a difference of 8 bushels per acre. Will this single observation give you enough confidence to make a management decision? Maybe yes, maybe no. You can answer this question if you understand your research objective and if you understand the purpose and value of **replication**.

First, check your objective for the comparison. You may be satisfied with the information from a single comparison, especially if your objective is to validate current practice. You may use such single replicate comparisons to simply validate product (seed, pesticide, and fertilizer) claims. Many products have a considerable amount of research information behind them, which you paid for when you purchased the product. Your on-farm research will cost you time and money, so you may reasonably conclude that a single comparison meets your needs. For example, your current fertilizer N rate for corn seems close to optimal (8 bu/A x \$2/bu $\approx$ 40 lb N/A x \$0.40/lb N $\approx$ \$16/A) for the field where the comparison was done.

Still, how much **confidence/uncertainty** do you have in your comparison? Watching your yield monitor as you harvest, you may notice yield varying by as much as 20 bushels per acre, which is greater than the yield difference between the comparison strips (8 bu/A). Knowing whether that 8 bushels per acre difference is "real", or just part of the "noise" in the data, is one of the things that "drive" replication (it also drives professional agronomists to small plots).

A single observation of the yield difference between any two treatments is one of a "population" of possible observations. Your yield difference (8 bu/A) may lie near the average value for this population of yield differences, but it might not. The "true mean difference" may be 2 bushels per acre, and 8 bushels per acre is closer to one end of the range in possible values (e.g. -8 bu/A to +12 bu/A). Replicating your comparison will give you more information about the variability in the population of yield differences, giving you more confidence in a management decision based upon your comparison.

The question of confidence in your comparison also often involves issues of "scale". If the field where you plan to conduct your comparison was large, containing several different soils, would the comparison be equally valid over the whole field? You grow corn in other fields on this farm, and on other farms. Would the "valid comparison" be appropriate (can it be extrapolated) to these other fields/farms? Additional replication of your valid comparison, within a large field, within other fields on the same farm, or within other fields on other farms, will give you more information about the spatial variability in the population of yield differences, but at different spatial scales. If you place additional replicates of your comparison in other fields and other farms, but not in the same field, you gain information at a larger scale, but lose information at the smaller, in-field, scale. You still might place all replicates of your valid comparison in one field to get a "feel" for the "noise" in the observed treatment difference(s), sometimes called "exploratory" research.

Having examined your research objectives, you decide to replicate. You will also need to **"randomize"** your treatments, to randomly allocate your two (or more) treatments within each "valid comparison" (replicate). Randomization is not required, but should probably be done, whenever a single comparison involves more than two treatments. However, whenever your on-farm experiment causes <u>more than one comparison</u> (one replicate) <u>to occur in the same field</u>, randomization is needed for full confidence in your results. Random allocation helps you avoid **systematic bias** that could occur if you simply alternated treatments (ex. 3 replicates of 2 treatments laid out as A, B; A, B; A, B), or if you simply put all replicates of one treatment on one side (ex. 3 replicates of 2 treatments laid out as A, A, A; B, B, B).

I assisted in the harvest of a Jackson Purchase strip trial, involving the use of a product (I will call it XYZ) on the corn crop. There were five replications of each of the two treatments. Corn was planted on the contour (a gentle slope was present), and the treatments were not randomized within each replicate. Instead, the treatments/replicates were planted side-by-side and alternated five times (A, B; A, B; A, B; A, B; A, B). Table 1 gives the results. The first thing you might note is that the yields generally declined as we harvested strips 1 to 10. In the field, being curious, I probed and found that the depth to the fragipan was shallower as we moved from replicate 1 to replicate 5. Second, making the "intended" comparison (see left side of Table 1), product XYZ looked "good" - there was a positive yield difference to the product's application, every single time (replication). Grower, county agent and I leave the field "enlightened".

Back at the office, I realize that the alternating strip layout gives me "another" comparison, if I leave out strips 1 and 10, to evaluate the experiment (see right side of Table 1). Product XYZ no longer looks "so good" and this causes a sober phone call from me to the agent and from the agent to the grower. What is the true impact of product XYZ on yield? Not sure, but probably lies somewhere between +13 and -11 bushels per acre, which means that we learned little after a lot of work. What could have been done differently, to improve our confidence in the results? First, we should have randomized our two treatments, either within each replicate

Table 1. Use of Product XYZ and Corn Grain Yield

"As Intended" Interpretation

strip #	replicate #	XYZ used?	grain yield	yield difference due to XYZ (yes-no)
			bu/A	bu/A
1	1	yes	226	
2	1	no	221	+5
3	2	yes	195	
4	2	no	172	+23
5	3	yes	204	
6	3	no	192	+12
7	4	yes	189	
8	4	no	178	+11
9	5	yes	132	
10	5	no	116	+16
average yield difference				+13

"Another" Interpretation

strip #	replicate #	XYZ used?	grain yield	yield difference due to XYZ (yes-no)
			bu/A	bu/A
2	1	no	221	
3	1	yes	195	-26
4	2	no	172	
5	2	yes	204	+32
6	3	no	192	
7	3	yes	189	-3
8	4	no	178	
9	4	yes	132	-46
average yield difference				-11

or within the entire study area. Also, we should probably have planted the test perpendicular to the trend in depth to fragipan.

How do you randomize? Randomization is done by flipping a coin (“heads” = treatment A; “tails” = treatment B), by putting treatment codes on individual papers and drawing these from an opaque container, or some similar approach. You can randomize across all replicates (a completely randomized design), but randomization within each replicate (randomized blocks) is usually recommended in on-farm research.

Finally, how many replications should you use? Remembering that each comparison takes time and effort, you might ask what you learn for the extra effort. Though the “Law of Unintended Consequences” can apply to those who use unreplicated information (and draw the wrong outcome from the “hat” (population) of possible outcomes), the “Law of Diminishing Returns” also applies to the gain in confidence from more and more replication.

Replication of a valid comparison accomplishes two tasks, relative to the “yield difference” between any two treatments within each comparison. Replication causes you to know more about the consistency of the “yield difference”. Replication also allows you to reduce the numerical value of the “yield difference” at which you have a certain “level of confidence” that the difference in yield is “true” (statistically significant). I should note an important distinction here. “Statistical significance” is not a measure of the “practical importance” of the observed difference. Statistics is used to “discipline” the validity of our conclusion (our claim) concerning an observed difference. Your professional experience is your guide to the practical importance of that same difference.

Your choice on the number of replicates you use in your experiment will interact with two other factors, one of which is under the control of the investigator (you). First, you decide the “level of confidence” you desire when you determine whether there is a “true” difference between the treatments. What you are choosing is the “probability” that you might make a wrong conclusion from the experiment - a one in five chance (80% level of confidence), a one in ten chance (90% level of confidence), or perhaps a one in twenty chance (95% level of confidence)? If you want to have a higher level of confidence, you will need more replication. The second factor, over which there is often little control, is the variation among individual treatment yields across your replicates, often expressed as the treatment "standard deviation". As the expected standard deviation rises, so does the need for replication, if you desire to "hold down" the numerical value of the statistically significant yield difference.

Table 2 takes these factors into account, together. I chose 80, 90 and 95% levels of confidence. To go any lower gets us close to 50%, for which we could flip a coin to make the management decision (and just skip on-farm research). I chose standard deviations of 3, 8 and 20 bushels per acre because these are 10% of the mean yield of good crops of double-crop soybean, wheat and corn, respectively. As the table illustrates, there are strong impacts from increasing your desired confidence level and from greater experimental variation (standard deviation), on the numerical value of the "true" yield difference. Greater replication reduces the numerical value, but this comes at a price.

Most on-farm research trials have at least three replications and many have four. Even with four replications, yield variation can cause professional agronomists (who often desire a 90% level of confidence) to "require" a relatively large numerical value to declare a "significant" yield difference.

As information at greater scales in space is desired it may be appropriate to compromise. Do some on-farm research on the land you farm, to inform the management questions you have. For the rest, consider cooperating with other growers to influence/get research done at a scale appropriate to the other questions you and they have (county extension; Kentucky commodity (corn, soybean, wheat, etc.) organizations; state and federal taxpayer funded research.

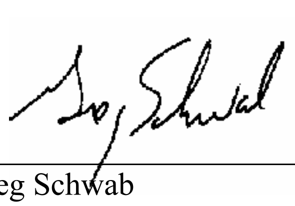

Greg Schwab
Extension Soils Specialist

Table 2. Level of Confidence, Yield Variation, Replication and the "True" Difference

Desired Level Of Confidence	treatment yield standard deviation	number of replicates	minimum yield difference for "true" difference*
(%)	bu/A		bu/A
80	3	2	9.2
80	8	2	24.6
80	20	2	61.6
80	3	3	4.6
80	8	3	12.3
80	20	3	30.8
90	3	3	7.1
90	8	3	19.1
90	20	3	47.7
95	3	3	10.5
95	8	3	28.1
95	20	3	70.2
80	3	4	3.5
80	8	4	9.3
80	20	4	23.2
90	3	4	5.0
90	8	4	13.3
90	20	4	33.3
95	3	4	6.8
95	8	4	18.0
95	20	4	45.0
90	3	6	3.5
90	8	6	9.3
90	20	6	23.3
95	3	10	3.0
95	8	10	8.1
95	20	10	20.2

* Least Significant Difference (LSD)

Cooperative Extension Service
U.S. Department Of Agriculture
University Of Kentucky
College Of Agriculture
Lexington, Kentucky 40546

Official Business
Penalty for Private Use, $300

An Equal Opportunity Employer

PRSRT STD
POSTAGE & FEES PAID
USDA
PERMIT NUMBER G268