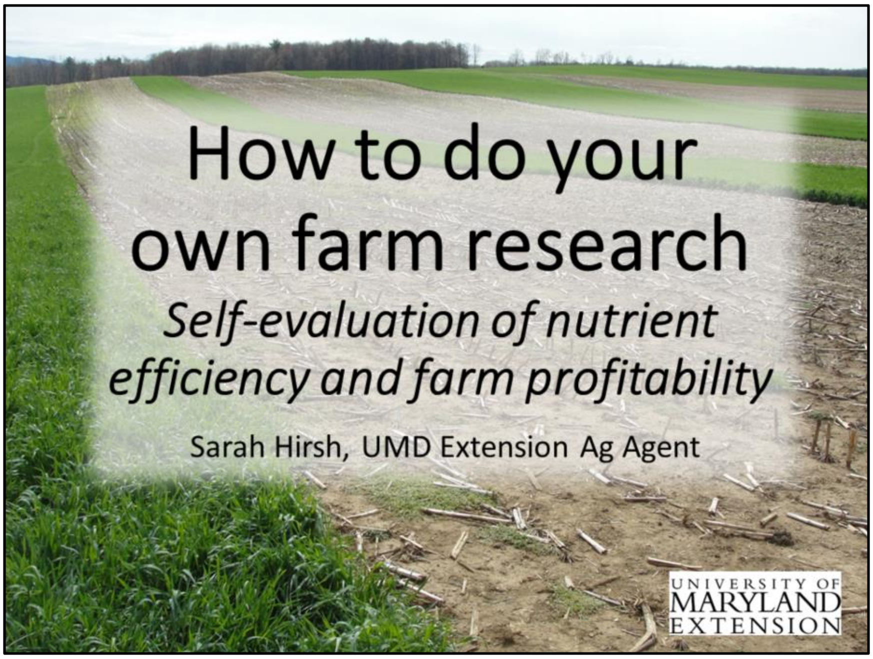

I will be discussing how to do your own on-farm research. I know that experimental design and statistics can be a dry subject, but I believe that understanding these concepts can be really helpful for getting the best information from practices that a lot of folks are already doing. I know that a lot of farmers are continually trying out practices or products to see if they improve yields, improve sustainability, improve profitability, etc.

Not this!

Mississippi State University annually conducts Variety Trials of public and private wheat and oat varieties (https://www.mississippi-crops.com/2019/10/03/2019-msu-short-list-of-suggested-wheat-varieties/)

So when I say “on-farm experiments”, I am not talking about this! These are variety trials run by Mississippi State University. Often when we think about experiments, we imagine something like this, but research trials can be much more simple.

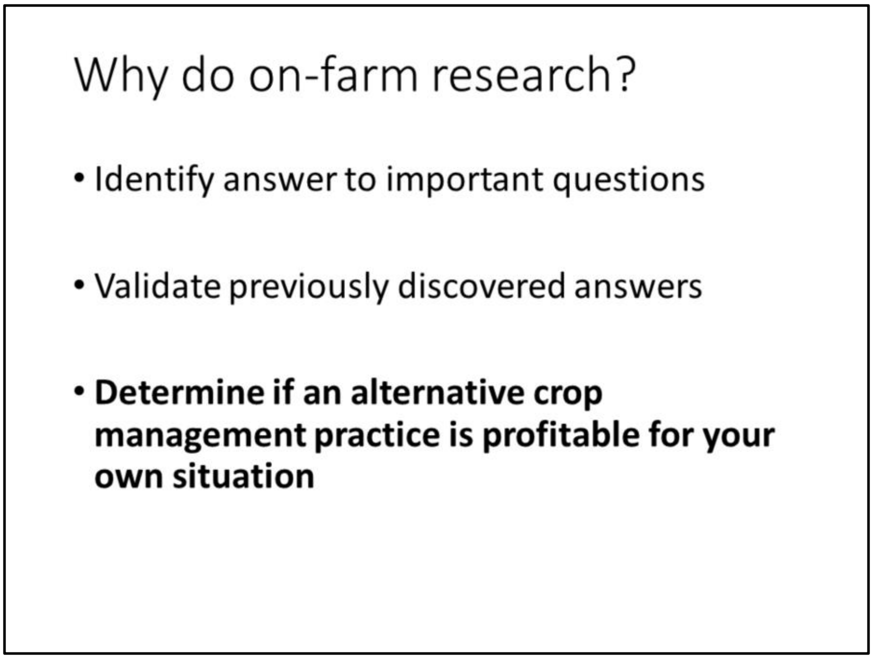

The primary reason to do on-farm research is to answer questions. These could be questions about whether a crop management practice or a pest control product will result in improved production or increased profit. Even if prior studies have found that a practice or product is worthwhile, on-farm research can be done to verify that a farmer gets consistent results on his or her own farm, with their unique climate, soil type, topography, management history, etc.

Why experiment and not just observe?

- Experiments are not documenting history
  - Variety 1 yielded higher than variety 2 last year
- Experiments predict future responses
  - Statistically probable that variety 1 will yield higher than variety 2 every time they are planted in the same field

So why would we want to run an experiment and not just make observations? Experiments actually enable us to predict future responses rather than just document history. Through making observations we can say, for example, that variety 1 yielded higher than variety 2 last year. But through running an experiment, we can say that it is statistically probable that variety 1 will yield higher than variety 2 every time they are planted in the same field.

Background noise

- Difference between evaluating historical performance and experimental data is because of background noise
- There are a lot of other "variables" in the environment (that you are not interested in evaluating)
  - Soil texture
  - Drainage
  - Compaction
  - Pest pressure (including wildlife)
- Designing an experiment with **replication** and **randomization** can control for this "noise" and allow us to evaluate differences from the variable we are interested in

The difference between making observations or evaluating historical performance and experimenting is due to what we call "background noise". Background noise is all of the other variables in the environment that we are not interested in evaluating. For example, if we are interested in evaluating two wheat varieties, background noise would be variables of soil texture, drainage, compaction, pest pressure (including wildlife), etc. When we design an experiment with replication and randomization (I will define these in a moment), we can control for this background noise and actually evaluate differences attributable to the variable we are interested in (e.g., the variety of wheat).

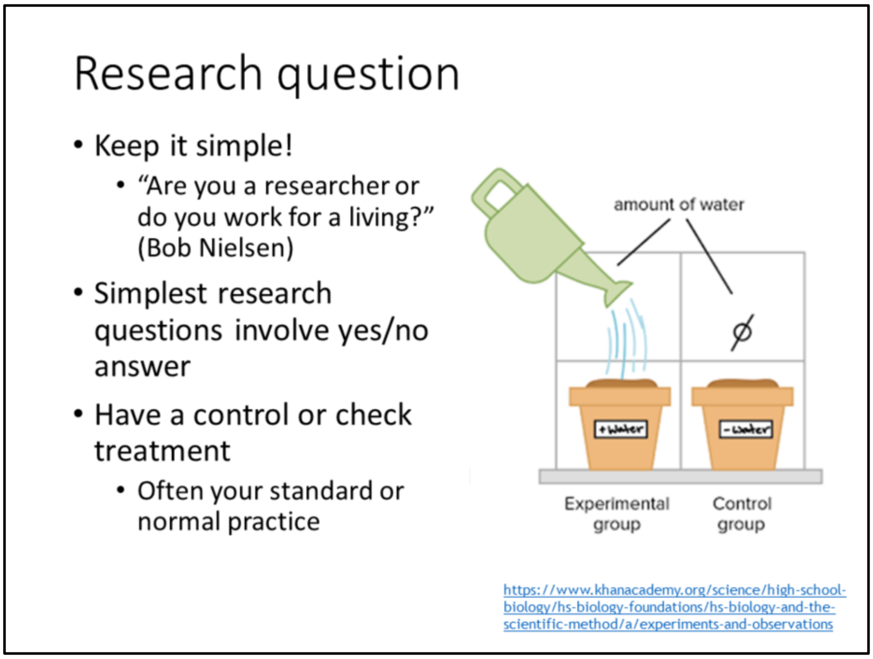

The first step to designing an experiment is having a clear research question. And probably the best advise someone could give is "keep it simple!". I like this quote by Bob Nielsen "Are you a researcher or do you work for a living?". It really drives home the point that experiments can quickly get out of hand when we are comparing a lot of variables or treatment levels, or when we are trying to measure too many things. The simplest research questions involve a simple yes or no answer. For example, does variety 1 yield higher than variety 2? Oftentimes, we will want to include control or check treatments. This will often be your standard or normal practice. So if you traditionally grow one variety of wheat, for the experiment, you may want to grow that variety and compare it to a new variety. Don't grow two new varieties, because then you aren't sure if the higher yields were because you had better varieties or for some other reason. For example, maybe you just had really good weather that year, in which case your traditional variety would have done just as well.

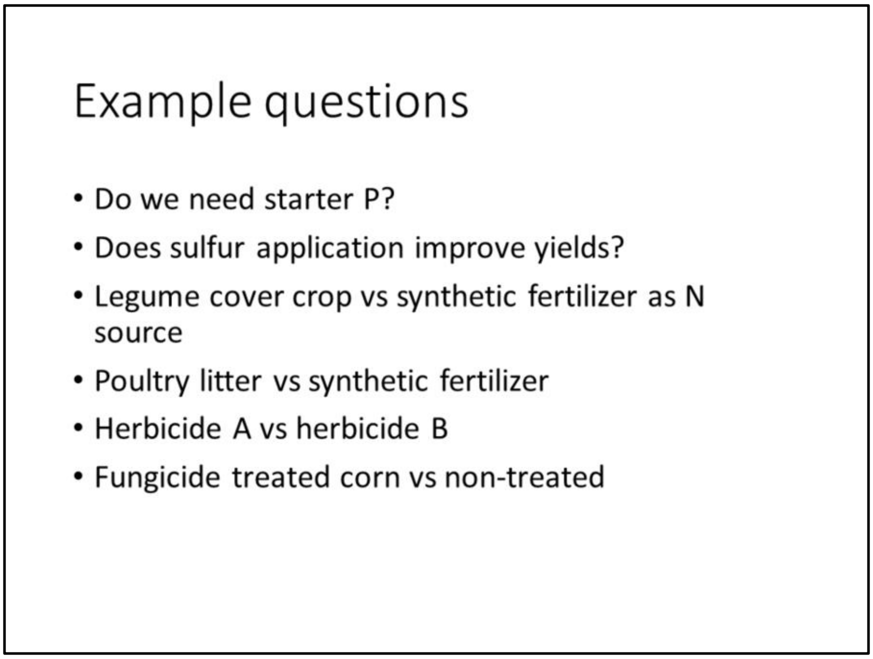

These are just a few example research questions. Do we need starter phosphorus? Does sulfur application improve yields? Does using a legume cover crop or synthetic fertilizer as a nitrogen source lead to higher corn yield? Is there a difference in yield following poultry litter versus synthetic fertilizer? Does herbicide A or B decrease weed pressure? Does using a fungicide on corn increase yield as compared to a non-fungicide treatment control? And you can think of many, many more!

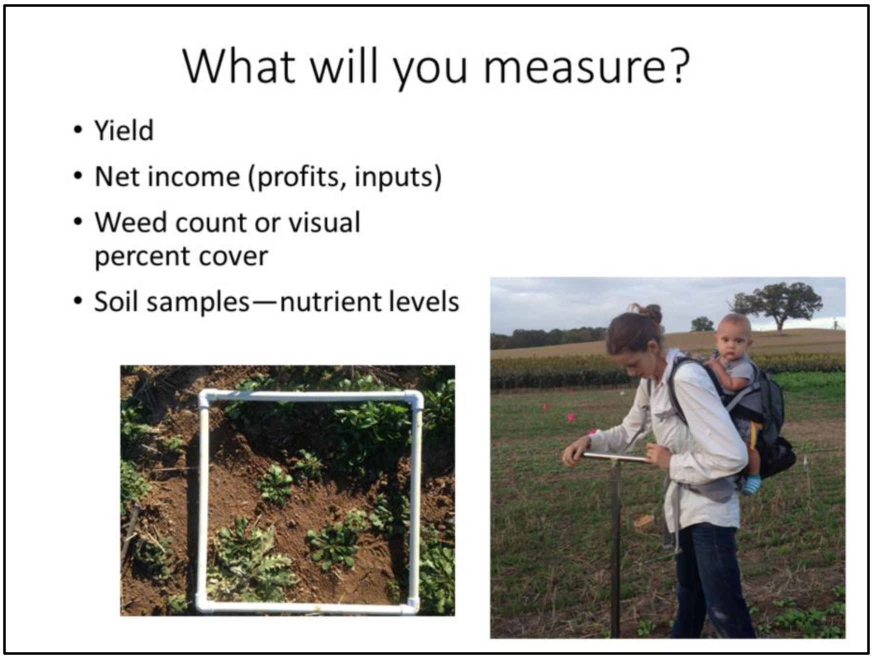

The next step is to decide what you will measure, because this will influence how the experiment is laid out and the size of your experimental areas. Yield is usually of great interest and can be relatively straightforward to measure (especially if the grower has a yield monitor). Farmers should consider doing economic analyses, because even if a practice or product results in increased crop yields, the cost of the practice or product may be inhibitive. Other things such as weed count or visual percent cover, or soil nutrient levels, can be measured, but these will be time intensive and could be expensive (e.g., you may have to pay soil lab fees). If you don't have the time (or the money to hire people) to collect the samples, train your children from a young age!

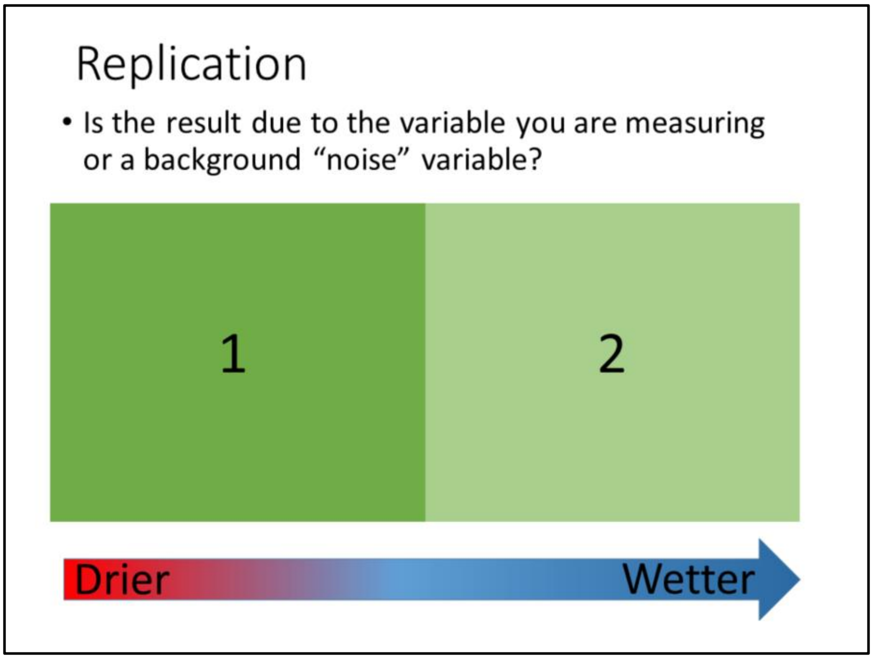

Next I will discuss replication and randomization, which are two key principles of experimental design. Replication is just repeating a treatment multiple times. With no replication, it is impossible to know if the result is really from the variable you are changing. For example, if we have treatment 1 on the left part of the field, and treatment 2 on the right, we can see that the left is also drier and the right wetter. Therefore, we don't know if the higher yield we found is a result of treatment 1 or a result of that part of the field being drier (or in other words, the result of background noise.

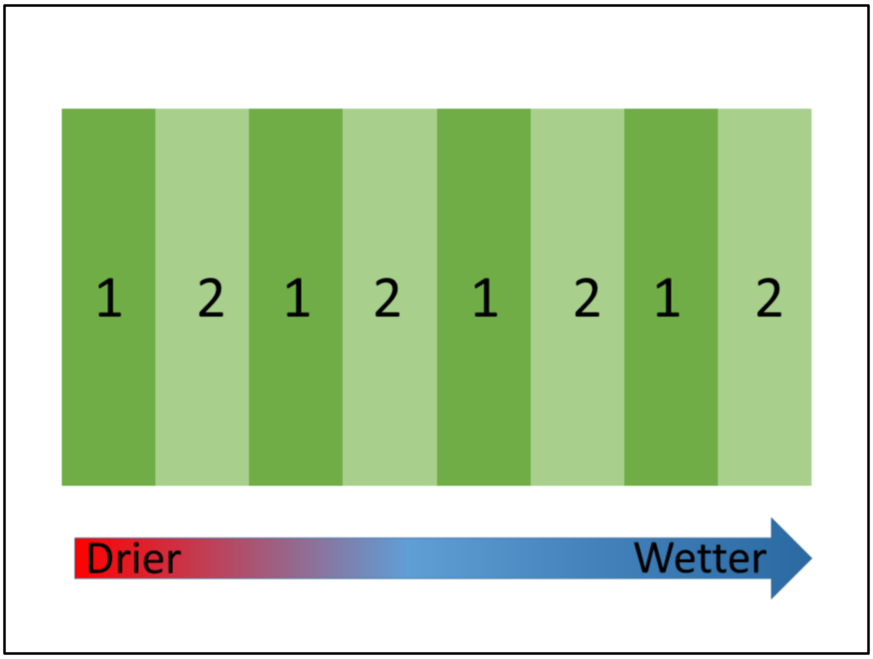

Now we can see that by replicating each treatment four times, we “spread out” the background noise, so it doesn’t influence results. Now we have treatment 1 and treatment 2 in both the wet and dry areas of the field.

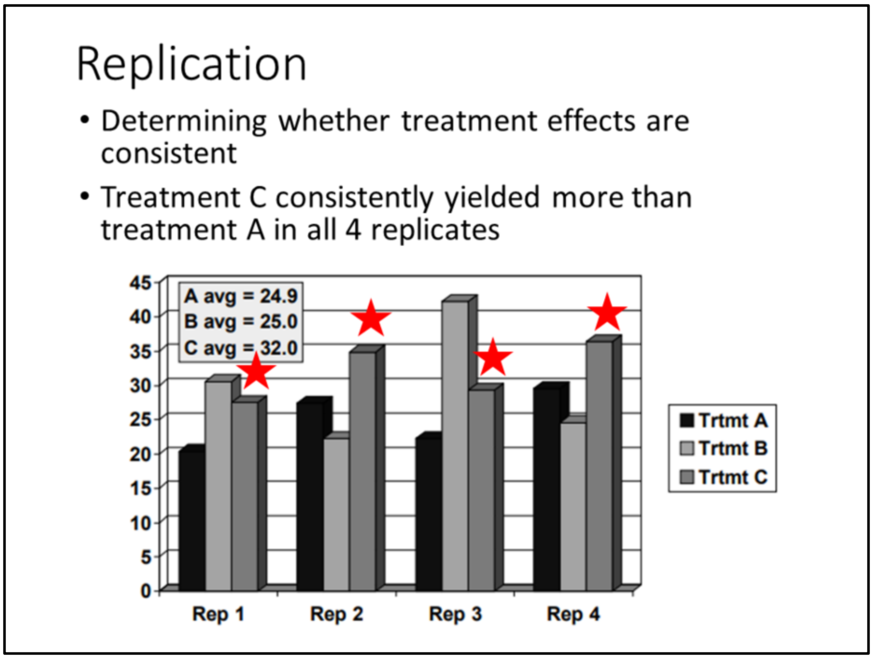

Replication also allows us to determine whether the treatment effects are consistent. In the bar chart, treatment C (indicated by stars) is always higher than treatment A (the black columns). So even without running the statistics, we know that treatment C likely has an advantage over treatment A. On the other hand, treatment B is sometimes higher and sometimes lower than treatment A, so it is not evident which treatment is better.

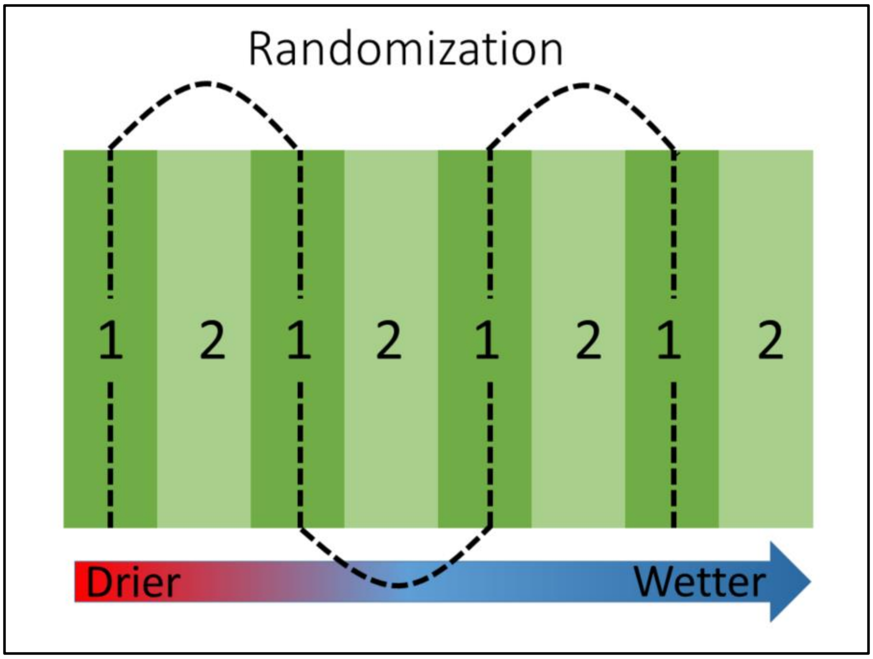

Randomization can further help to “spread out” background noise, so it doesn’t influence results. Here we can see that because we replicated our treatments, our variable of interest is spread out over the wet and dry areas of the field. However, imagine that a sprayer followed the pattern of the dotted line as it drove through the field, and it had a broken nozzle, resulting in a heavy spray on the center line. In this case, every “1” treatment could be influenced by that sprayed product, rather than by the variable of interest. Therefore, we want to include randomization of treatments as well as replication in experiments.

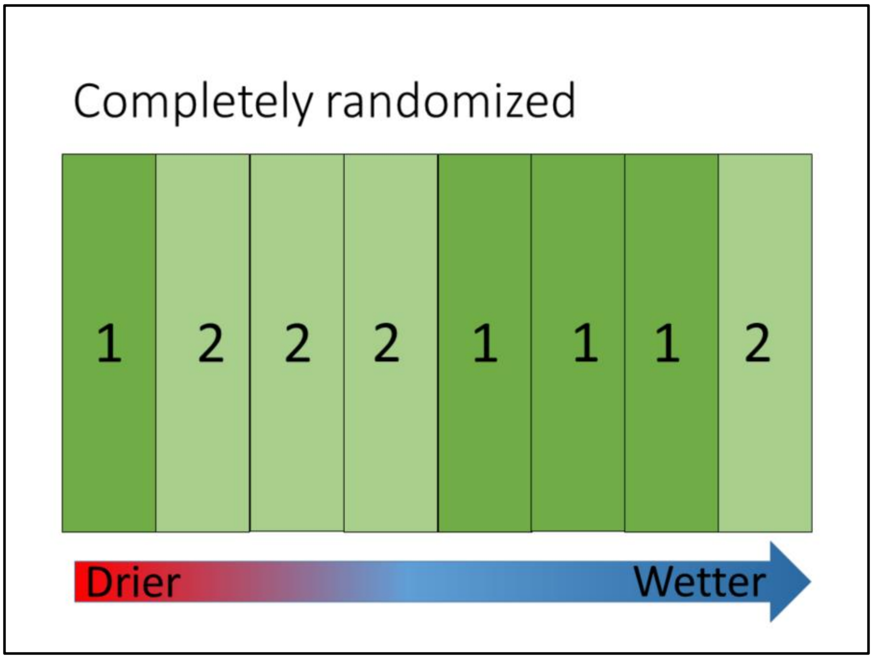

Complete randomization is when we put four “1s” into a hat and four “2s” into a hat and randomly draw the eight numbers. By chance, we could now, again, be grouping the “2s” into the drier end of the field and the “1s” into the wetter end of the field. This is why we often will choose to do randomization “in blocks” rather than complete randomization.

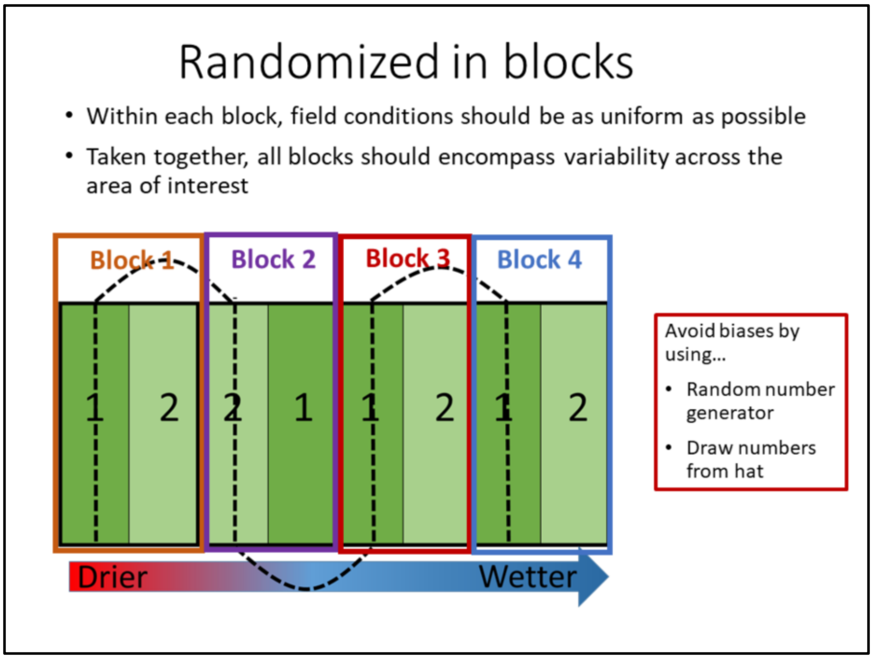

By randomizing in blocks, we are blocking off the field based on background variability (e.g., in our example the wetness of the field). Within a block, the field conditions should be as uniform as possible. Within a block, we will randomly draw numbers to decide the order of treatments. Taken together, the results from all of the blocks should encompass the variability across the area of interest, but we are not allowing that variability to influence our results. One other sidenote, it is a good idea to either flip a coin, use a random number generator (can google and find them), or draw numbers from a hat; avoid randomly label plots “from your head” because we tend to be biased.

Setting up the plots

- How wide and long
  - Think about your equipment
  - Yield monitor?

Before putting strips out in the field, it is very important to think about what data you want to collect from the strips and the equipment that you will be using. If you will collect yield, think about the width of your combine.

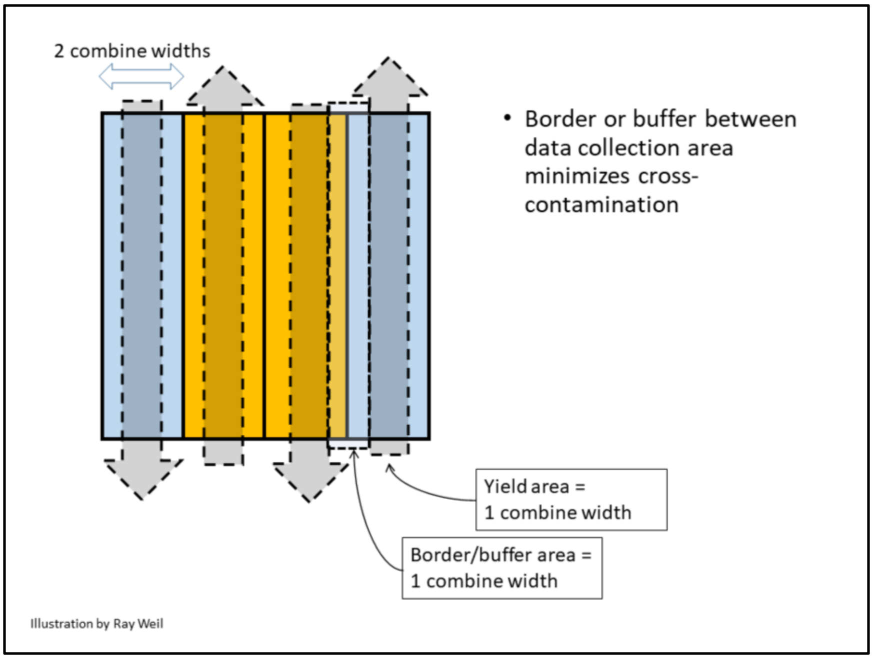

I would recommend making the strips twice as wide as the combine. This is because we want a border between treatment strips so we don't get drift from one treatment into the next. So it is convenient if we make a strip two times the width of a combine, so we can harvest the middle of the strip in one combine pass, and then the border of that strip and the neighboring strip in the following pass.

Record observations

- Any anomalies
  - Weather
  - Field conditions
  - Pests
- You can always decide to measure extra things based on observations of differences among treatments

During the study period, measure and record the data that you had planned to measure. Also, be aware of any additional observations. There will most likely be some unexpected influences on the experimental field. For example, we may have deer damage, as seen in the photo here. We would want to be sure to take a note about what plots that damage was in. Be aware of any weather related anomalies. For example, is there water pooling in a certain part of the field? Are there pests encroaching on the wooded edge of the field? Also, do you notice any patterns that you weren't expecting? You can always change what you are measuring or measure additional unplanned things. For example, if you notice one treatment has very few weeds, you could decide to measure weeds in both of the treatments to be able to run statistics and know if that treatment does indeed lead to lower weed pressure.

Analyzing the data

- Cleaning up the data (removing outliers) is allowed
  - E.g., removing yield data points from flooded area or area with heavy deer damage
- Statistical analysis
  - Ask someone with statistical software
  - Looking at averages and variability or "variance" and determining the likelihood that the results did not happen by chance
  - You can choose your confidence level, but typically 90% or 95%
    - 95% likely that herbicide A different than herbicide B, because of the herbicide, not as result of background noise

If you notice an anomaly in the field due to something like pests or weather, it is OK to throw that data point out! That is not cheating; it is called removing outliers or cleaning up the data. We do not want to include data that is telling us something that is clearly caused by another source, rather that the treatment of interest. Now is time to compare treatments through statistical analysis. Statistics are basically looking at the average measurements of a treatment and how much variation there is within the measurements of that treatment. You are determining the likelihood that the results did not just happen by chance. You can choose the confidence level, which is the level that you are comfortable with of the effect *not* being by chance. Typically people choose 90% or 95%. So, for example, if you choose 95% confidence level, you are saying that it is 95% likely that herbicide A resulted in lower weed pressure than herbicide B, because of the herbicide, not because of anything else or any "background noise".

Analyzing the data

- Averages and variability

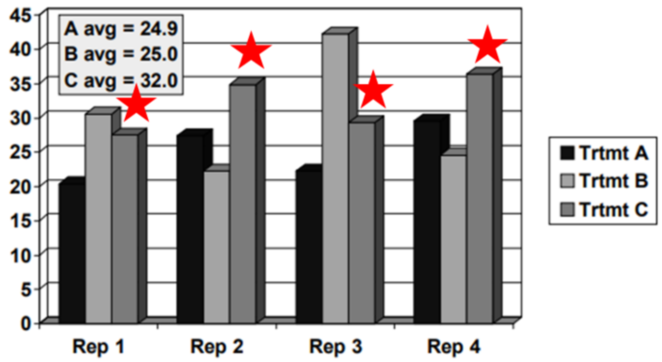

If we revisit this example, we can see that treatment C consistently leads to higher values than treatment A (in all of the four replications). The four measurements of treatment C do not vary much from each other, and the four measurements of treatment A do not vary much from each other. So if we ran the statistics, they would indicate that there is a high probability that treatment C leads to higher values than treatment A.

Research Question—
Clover as a green manure?

Chart 2 PERFORMANCE AND ROLES

Species	Legume N Source	Total N (lb./A)[1]	Dry Matter (lb./A/yr.)
Berseem clover *p. 118*	●	75–220	6,000–10,000
Cowpeas *p. 125*	●	100–150	2,500–4,500
Crimson clover *p. 130*	◕	70–130	3,500–5,500
Field peas *p. 135*	●	90–150	4,000–5,000

Managing Cover Crops Profitably, 3rd Ed (SARE)

Now let's put this all together and walk through designing an actual on-farm experiment. This farmer is interested in comparing yield of corn following a green manure versus following synthetic fertilizer. They did some background research and found that Crimson clover is a good green manure species for their region.

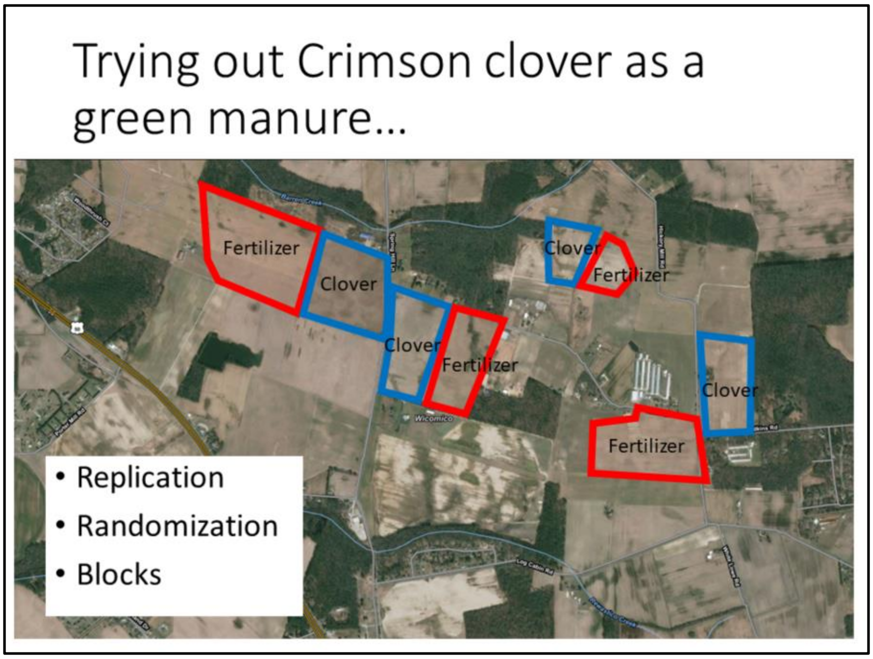

Rather than growing or applying Crimson clover and fertilizer in strips in the same field, they decided for convenience to put the clover and fertilizer in separate fields. This is fine! Treatment areas don't have to be uniform in size or all in the same field. This farmer had eight fields, which they split into four blocks, based on differences in soil type of the fields. Fields in the same block had similar soil types and other physical features. Each block contained two fields. Within those four sets of two fields, they chose which field would be clover and which fertilizer by flipping a coin. Therefore, this study has blocking, replication, and randomization.

Analyzing the data…

- Yield of following corn
- Economic analysis of inputs and profits

Corn yield was measured in each of the eight fields. You could also do a economic analysis, factoring in yield differences and input and labor costs.

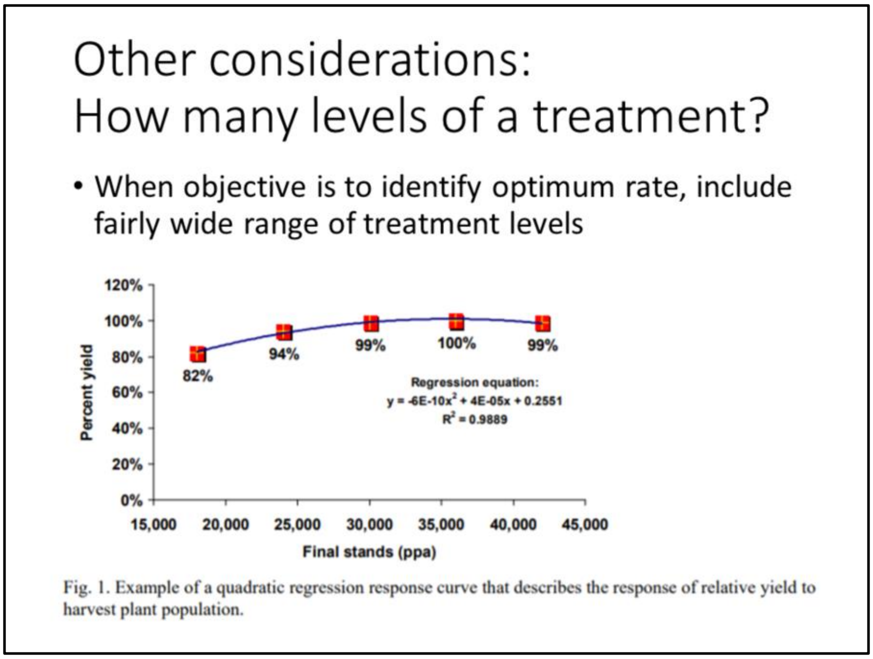

Now I will go through a few other considerations to be aware of when doing on-farm research. We should think about how many levels of a treatment make sense. When our objective is to identify an optimum rate, we want to include several and a fairly wide range of treatment levels, and have the lowest below what we think would be feasible and highest above what we think would be feasible. Remember, in this case, we want to be able to see the shape of a curve, including what happens with both the lower and upper extremes.

Other considerations: Economics

- Just because a certain practice yields more, doesn't mean it's cost effective!

And again, just because a certain practice yields more, that does not mean it's cost effective. We need to consider input costs and other factors.

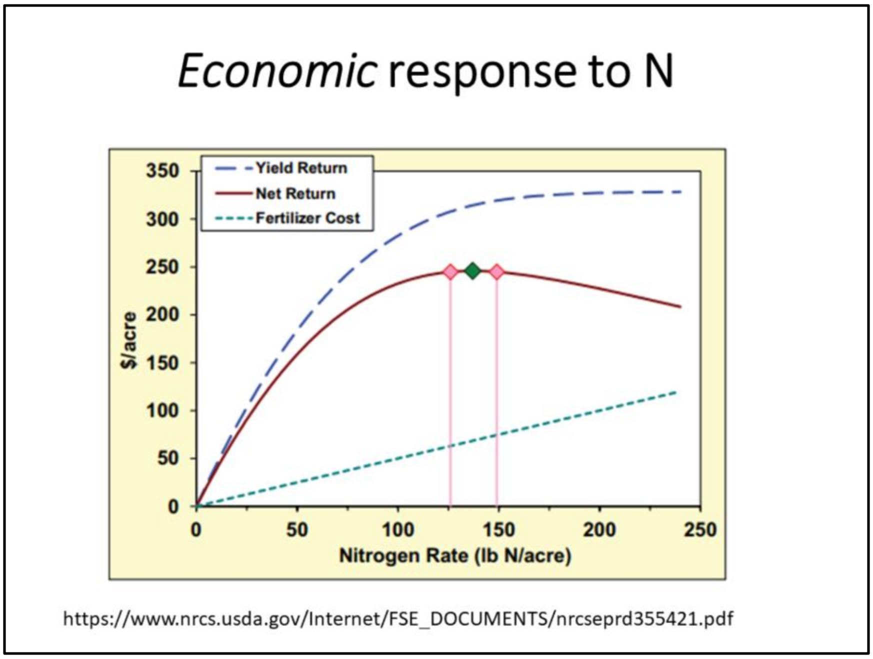

A good example of this is illustrated here with the economic response curve to nitrogen. Nitrogen fertilizer cost increases linearly with amount. However, the yield response to nitrogen diminishes as the nitrogen application amount increases. Therefore, there will be a point where it doesn't pay to add more nitrogen, even if yield has not quite peaked.

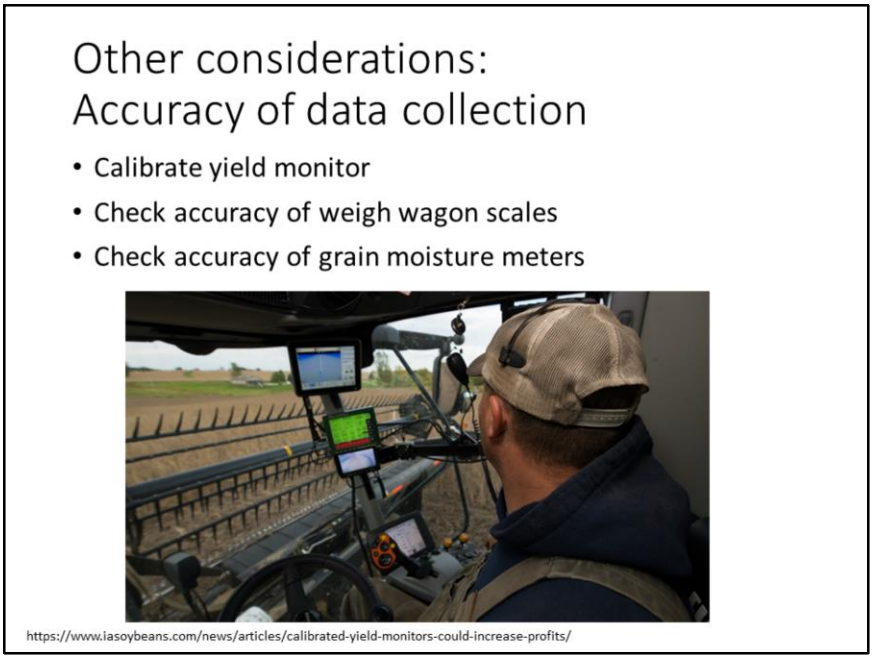

Another consideration when doing on-farm research is to make sure your equipment to measure yields is calibrated and accurate. You certainly don’t want to make the effort to do an on-farm trial and then have inaccurate data due to poor accuracy of equipment.

Finally, the more repetition the better! Repeating experiments over multiple sites and over multiple years is great. See if some of your neighbors are interested in trying the same thing as you. And don’t be over zealous in changing practices. I wouldn’t recommend changing practices on a large scale following one experiment, even if it was set-up using proper experimental design. You will at least want to repeat that experiment over multiple years. Especially with agricultural experiments, year by year variation in weather can change results. For example, one practice may result in higher yield during a dry yield, but not during a wet year.

Collaboration

- Other farmers
- University or industry professionals
  - Experimental design
  - Run statistics

Somerset Co Extension
30730 Park Drive
Princess Anne, MD

Hours: 8:00am - 4:30pm

Phone: (410) 651-1350

Research works best as a collaborative effort. Involve and get help from other farmers and University or industry professionals. For example, I am happy to be involved and I know my colleagues at other Extension offices would be interested as well.

In summary, farmers are already doing many of these things. They are continually trying new practices, making observations, taking measurements such as yield, and adapting management. Therefore, it would be relatively easy to make sure to incorporate replication and randomization when trying new things, in order to learn more from findings. Again, experimenting enables us to not just document what happened, but to predict what will happen in the future. And don't forget we are hear to help!

References

- Nielsen, R.L. 2010. A Practical Guide to On-Farm Research. Purdue University. https://www.agry.purdue.edu/ext/corn/news/timeless/OnFarmResearch.pdf

- Chaney, D. 2017. Technical Bulletin: How to Conduct Research on your Farm or Ranch. SARE. https://www.sare.org/Learning-Center/Bulletins/How-to-Conduct-Research-on-Your-Farm-or-Ranch

These are two great publications where much of the information from this PowerPoint was found.

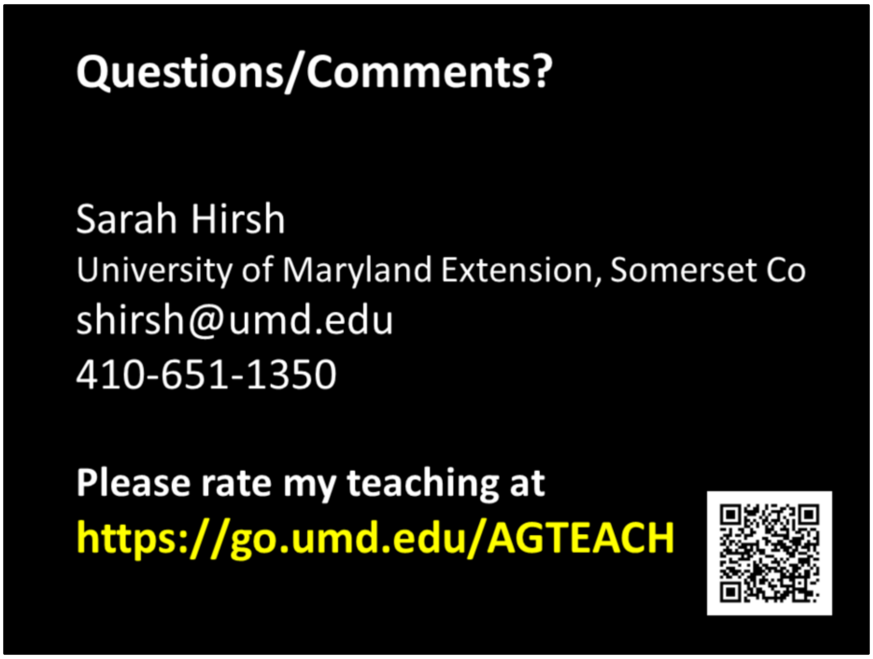

Feel free to reach out to me for comments or questions. Since we are funded by a public land-grant University, we ask our audiences to rate our teaching, so we can get feedback and improve. You can follow this link, and click my name from the dropdown. Thank you!