Using Yield Monitors for On-Farm Research

Randy Taylor, Scott Staggenborg, Kevin Dhuyvetter,
Terry Kastens, John Schmidt, Mark Schrock
K-State Research and Extension

One of the potential uses of yield monitors is to conduct on-farm research. This will allow producers to develop information that is specific to their farms or areas. However, to obtain valid results, a methodical approach must be taken. Producers should have a specific goal in mind when the research is planned. The goal could be to investigate yield responses to fertility, cultural practices, hybrids, or varieties. However, examining all these variables at one time could present a nightmare in terms of plot layout and data analysis. A better approach might be to pick one variable to research. Although yield monitors can be used for research, they have some limitations that should be understood. This publication is designed to assist the planning, management, and analysis of data for effective on-farm research.

Experiment Design

Designing a research experiment can be a fairly simple process if a few general rules are followed. Keep in mind, a poor design can lead to inaccurate results which in turn can lead to costly decisions. There are several items to remember when designing experiments. The first and foremost is the hypothesis. What is it that you want to know? This is more of a philosophical question, but one that must be answered at the onset of a research effort. Other items, like the number of treatments, plot layout, and replications are more mechanical in nature.

Testing a hypothesis

Any research effort begins with a hypothesis. The research is then designed to test this hypothesis. The simplest hypothesis is a basic statement that you believe to be true; for example, *using a starter fertilizer improves yield*. A research project (with and without starter fertilizer) can then be designed to test this hypothesis. The results will either support or reject the hypothesis. Statistical analysis of the data will be used to help make the decision by assessing the strength of the evidence against the hypothesis.

Selecting a hypothesis is the most crucial part of any research effort. Ideally, the hypothesis for on-farm research should be selected to supplement or validate other research results, either from universities or industry. The task of researching all of the variables of production agriculture would be mountainous. Trying to ‘fine tune’ university or industry research results for a particular farm or area is a very realistic goal. For instance, published research results may indicate that a higher seeding rate for corn will result in increased yields. Weighing the high cost of corn seed against the potential returns of higher yields is something many producers are doing. What is the optimum seeding rate for an individual farm? What is the most profitable seeding rate for a particular soil type? These are two researchable questions with testable hypotheses and will be examined in further detail in this publication.

Factors and Treatments

The main element of a hypothesis that is to be tested is called the **factor**. Different levels of a factor are called **treatments**. If we are studying seeding rates for corn, plant population is the factor and the different seeding rates are the treatments. There must be at least two treatments in any research project. However, research can be designed for multiple factors and treatments, but data analysis and treatment establishment becomes more complex and time consuming as the number of factors and treatments increase.

Though two treatments are a minimum, there are some instances when two are not enough. If the hypothesis involves an application rate of crop inputs such as seed, fertilizer, or lime, it is recommended that there be at least three treatments and preferably more to properly describe the relationship of the factor's influence on yield. As an example, consider some actual data from a single factor sorghum fertility study at the North Central Experiment Field (Figure 1). If the producer is currently applying 90 lbs N/acre (in a sorghum/soybean rotation) and wants to test if this is the proper amount, selecting three treatment amounts can make a difference in the results and management changes that could result from the work. If 60, 90, and 120 lbs N/acre are chosen (A's in Figure 1), then the results would indicate that the producer could reduce the amount of nitrogen fertilizer by 30 lbs/acre. However, if the producer believed that 90 lbs N/acre was the minimum that could be applied and chose 90, 120, and 150 lbs N/acre (B's in Figure 1), the results would be different. With these three fertilizer rates, the results would indicate that 90 lbs N/acre produced yields equivalent to 120 and 150 lbs N/acre. In both cases, the results indicated the sorghum yield response to nitrogen within the N fertilizer range tested, but the results lead to different fertilizer management decisions by the producer.

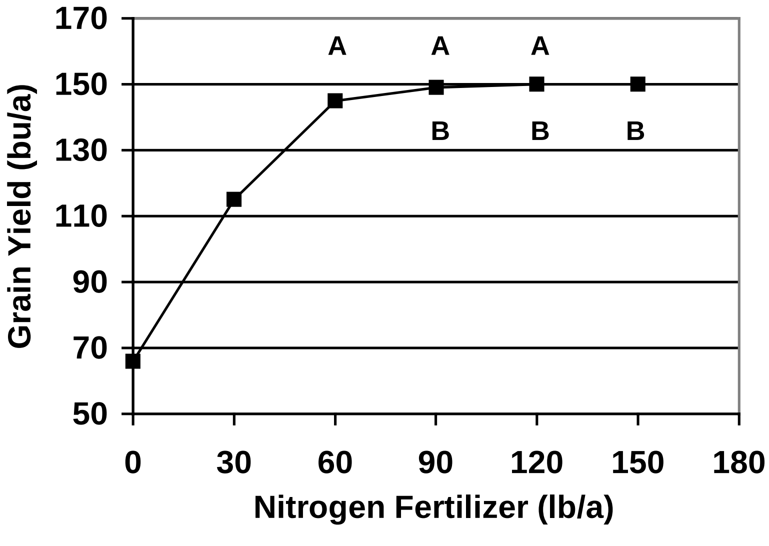

Figure 1. Results of a nitrogen fertilizer study on grain sorghum.

If three treatments are the limit in a research project, select what is perceived to be the level that results in optimum or maximum yields and then select one below and one above this optimum treatment so as to actually test this hypothesis. If four or more treatment levels are to be used, then selecting as wide a range as possible, within realistic levels, is recommended. Taking the fertility example above, if four treatments are to be used, 0, 20, 40, and 60 lbs N/acre would not show the optimum yield level as there is never a "leveling off" in yield. The results from such a study may lead one to believe that simply increasing nitrogen rates would result in increased yields without limits if the data are not carefully interpreted. However, 30, 60, 90, and 120 lbs N/acre would describe the yield response to nitrogen and lead one to conclude that 60 lbs N/acre results in the optimum yield.

An example of a two-factor experiment is a hybrid by population study that will be referenced throughout this publication. The field was planted with a twelve-row planter with the planter split between two hybrids. The first factor (hybrid) has two treatments (hybrids A and B). Thus, one pass through the field would plant six rows of each hybrid. This is typically called a split planter comparison or strip plots. The difference here is that each pass through the field was at a different population. Population is the second factor in the study. Each hybrid was planted at four different populations (4 treatments). Each population was replicated four times. While this is a well designed research experiment (2 hybrids x 4 populations x 4 replications = 32 plots), it might be more complex than most producers would want to try. Though analyzing this data will be more challenging than a single factor experiment, the real challenge will be establishing the plots. This can be time consuming during planting season when time may already be limited.

Producers should be wary of changing their entire farm into a research project. A base is needed to compare any new practices or varieties. The base should be something that is proven in your area. It may be a standard practice of blanket application of fertilizer or a variety that has performed well in the past. The best approach is to keep on-farm research simple and focused with hypotheses that are testable.

Plot layout

Producers should give some thought to plot layout during the early stages of planning research. Factors that are not a part of the experiment can influence research results. These should be eliminated or minimized. You should know as much about the plot area as possible. For instance, there is little value in conducting a phosphorus fertility study in an area where the soil test is high. Or, if there is a spot in the plot area that is historically low yielding, this spot should be avoided or divided equally among treatments.

Split planter

The split planter comparison is typically used for hybrid/variety comparisons. Each half of the planter would be filled with a different hybrid. Thus, one pass through the field would plant one treatment of each hybrid. For a split planter experiment, the planter must be twice the width of the combine's header to avoid harvesting across guess rows. These types of comparisons will typically be difficult to visualize on a yield map due to other items that affect yield across a field. However, one Pioneer Hybrid International, has supported development of software to analyze split planter comparisons by creating difference maps. This allows examining areas of the field where two hybrids may be different. The software is being distributed free.

Strips

Using strip plots is a reasonable alternative to split planter design when planter is not twice the harvesting width. Two passes of a 12-row planter (2 x 30’ = 60’) would equal three passes of an 8-row or 20’ header (3 x 20’ = 60’). Strip plots are fairly easy to establish when they run the full length of the field. The whole planter plants the predetermined number of passes for hybrid first, followed later by the same number of passes of another. Increasing replications or treatments will increase the workload of establishing strip plots. If you expect the difference between treatments to be large enough to see on a map, the strips should be wide. Ideally, strip plots would run the length of the field and each treatment would be at least four combine passes wide.

Plot size and yield monitor limitations

Though yield monitors can be an effective tool for measuring yield in on-farm research trials, they do have some limitations. The flow of grain in the combine does not start and stop abruptly. Crop cut at the edges of the header takes longer to reach the sensor than crop cut at the center of the header. The general dynamics of grain flow in the combine results in yield maps that only approximate crop yield at a given point. It is well documented that yield monitors tend to overestimate low yielding areas and underestimate high yielding areas.

It takes about 10-15 seconds for a combine to reach full capacity once it starts harvesting. A combine travels about 1.5 feet every second for each mph of ground speed. Harvesting at 5 mph, a combine travels about 100 feet in 15 seconds. Yield monitor data from plots that are shorter than this will have questionable accuracy. Plots should be at least 500 feet long. Also, if plots are perpendicular to combine travel, a buffer zone between plots is needed.

Response to crop inputs may be examined by applying a diagonal strip to a field. If you suspect that pH is affecting yields, lime can by applied across the field in a diagonal strip. The strip should be at least 200 feet wide to show up on a yield map. This is not really research, but it could provide some information about potential research.

Replications

Next to selecting a hypothesis, replication is probably one of the most important requirements for reliable research. Consider the hybrid/population example described earlier. There were four replications in this experiment. This means that each hybrid/population comparison was planted at four different locations in the field. More replications will allow smaller differences between treatments to be detected through statistical analysis. Sometimes research can be accomplished with two replications, but three to four replications are more typical. The adage that more is better certainly applies to replications in research.

Environmental Interaction

Multiple sites can be used as replications or site-years. This allows comparison in different environmental conditions. This can be done on a single farm or in cooperation with neighbors. The advantage of cooperative research is that it reduces the workload on each individual, but increases the confidence in any results uncovered in the research.

An example of this is the Riley County soybean and sorghum variety test that uses multiple locations as replications. This variety trial is designed to use 4 to 6 different locations within the county, each location with a different cooperator. Each cooperator plants a typical strip plot with each variety planted once as well as several strips which contain the tester variety. At the end of the season, each location is used as a replication, with the tester used to adjust yields across locations, and perhaps within a location if other factors substantially impact yield one side of the field to the other. This type of design is statistically valid and gives very consistent results while spreading the workload to several different producers with a common interest.

Data Collection

In Season Notes

Taking accurate notes and ratings during the season becomes extremely valuable after harvest and all that is left is plot yield data. When taking notes during the growing season, one must determine if patterns that exist in research plots are the result of the factors or treatments or the growing conditions. If increased vigor, emergence, or plant height occurs consistently in one plot compared to others, then it is reasonable to believe that the treatment may be responsible for the response in plant growth. However, be sure that the response is consistent across all of the replications in the study. If increased vigor, emergence, or plant height occurs in an area of the field that spans across several plots but does not encompass the entire plot area, careful notes should be taken so that it can be determined if the final yield of a given plot was adversely affected by this field problem. Georeferencing or marking all areas may help when analyzing data after harvest.

Emergence – It is important to take notes on differences or problems in corn, as reduced stands can impact yields. Stand establishment has less of an impact on sorghum, soybeans and wheat unless dramatic losses occur. Again, whether emergence problems span replications is important to document – to help interpret any yield differences later.

Vigor – Vigor is a good indication of a treatment's affect on seedling growth. You might arbitrarily assign a rating scale (1 to 5) for seedling vigor where 1 is low vigor and 5 is high. This results in a relative rating where treatments are compared to each other.

Weed/Disease Pressure – Researchers must determine if pest pressure is consistently different between treatments, suggesting that the treatment may have affected the pests as opposed to it being a field wide problem. Weed pressure may be significantly different between treatments if row spacing is being studied.

Plant Height – Plant height is easily measured or scored and is a good indicator of plant growth over time.

Plant Stress – If water stress is suspected, look for midday wilting. For nutrient studies, researchers need to be able to identify deficiency symptoms of the nutrient being studied.

Weather data – Weather has a significant impact on crop yields, so adequate weather information will help explain some responses. Rainfall (time, amount, intensity, and runoff potential) is probably the most critical piece of weather data. Noting the crop's growth stage at the time of rainfall events is also valuable. Temperature can also be valuable information during different growth stages. It has the possibility to stress the crop or provide excellent conditions for a bumper crop.

Anything that affects yield – Be on the look out for anything that may affect plant growth and ultimately yield. Such things as cultivator blight, compaction, and fertilizer or herbicide application problems may affect research results.

Harvest

Harvest is another crucial time in on-farm research. You typically feel pressured to keep the combine rolling. However, this is not the time to take short cuts. You have a full growing season invested in this research and poor data collection can ruin the entire experiment. Worse yet, it could lead to erroneous results that may lead to wrong conclusions.

Gathering data

The option of defining or marking loads with your yield monitor is a big advantage during harvest. This allows you to separate data electronically, although it is still a good idea to write down the values for each treatment along with any other observations. Be sure to use the wet weight from yield monitor to make comparisons. You can adjust for moisture later. The yield in the monitor is based partly on an estimated area. You're better off measuring the area of each plot if they are not the same size.

Combine speed is another item to consider while harvesting plots. Yield monitors are sensitive to large flow changes (some more so than others). Speed changes will affect grain flow in the combine. Proper calibration will eliminate most of the effect of minor flow changes. Figure 2 illustrates this concept. Though an attempt by the operator was made to harvest at a consistent speed, flow varied by ±20 percent from average. As flow rates (shown on the X-axis) increased, yield monitor accuracy (depicted as error on the Y-axis) decreased. The yield monitor was about 8 percent low on average for these plots.

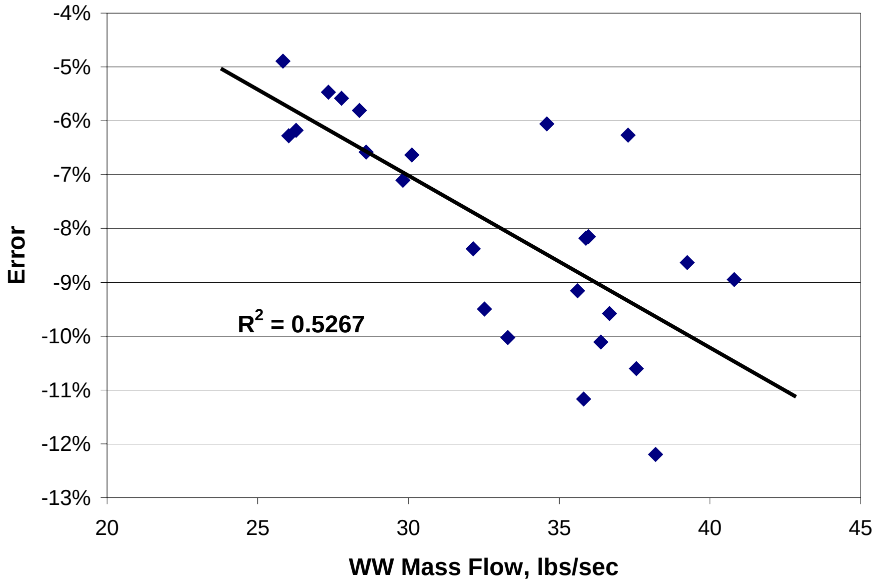

Figure 2. Influence of flow rate on yield monitor error. Flow rate is based on the weigh wagon total and error is the percent difference between the weigh wagon and yield monitor.

Data validation

Calibrating yield monitors is very important if you are planning to do research. If they are calibrated correctly, most yield monitors will be within 3-5 percent of the field total. You should expect similar accuracy on larger research plots. This alone may be enough to satisfy some producers, but a diligent researcher would need some type of validation. One method to validate results is to weigh the grain harvested from the entire plot area or one plot that is near the center of the plot area and compare it to the yield monitor total.

Another method of validation would be to weigh all or some of the treatments. Weighing all treatments defeats the purpose of using a yield monitor, but could instill some confidence that the yield monitor is ranking the treatments correctly. The data shown in Figure 2 were collected with a weigh wagon and yield monitor. Even though there was error with the yield monitor, the results of analyzing both data sets (yield monitor yield and weigh wagon yield) with regards to the research treatments were consistent. Regardless of which yield value was used, regression results led to the same conclusion.

Harvest Notes

As with in-season note taking, notes taken at harvest can also be valuable as the data are analyzed. Although it may seem time consuming, 15 to 20 minutes of note taking can make data analysis and interpretation easier and more meaningful.

Lodging **–** Lodging is often an indication of diseases and/or insect feeding. It is important to observe any difference that may occur between hybrids or varieties.

Plant Height **–** Plant height is easily measured or scored and is a good indicator of plant growth over time.

Maturity/Dry Down Differences **–** different hybrids mature and dry down at different rates. These notes may become valuable when selecting hybrids or varieties for future years.

Software and Data

Viewing differences between treatments on a yield map will be challenging. First, the difference between treatments needs to be large enough that the range selection will show it. Second, the size of the plots needs to be large enough that data from adjacent plots have minimal interference. If plots are small, data from an adjacent plot will be used during interpolation to determine the value for that plot. Large, replicated plots will be an asset during visual interpretation. If plots (factors and treatments) are georeferenced, they can be used in a GIS as part of a query along with other items that can affect yield.

Data Analysis

Once all the field research has been completed the data need to be analyzed to determine if the different factors or treatments had an effect on yield or some other variable (e.g., lodging, harvest moisture). When we refer to analyzing the data we typically are referring to a statistical analysis. Statistical significance is important for establishing confidence in the results. Using profit per acre as an example, consider farm A, which is assumed to have the following annual profit stream over 5 years: {–$80, $200, –$50, $300, –$270}. The average annual profit for farm A is $20/acre. What would you expect farm A's profit to be in year 6? Although your best guess is $20/acre, you would not have much confidence in that prediction. With the large variability displayed in farm A's profits it can easily be shown that its $20/acre average profit is not statistically different from 0. Now consider farm B, whose profit stream is {–$5, $30, $20, $25, $30}. Like farm A, farm B's average profit is also $20/acre. Now, however, it is much easier to have confidence in a $20 prediction for year 6. In this case, the $20 average is statistically different from 0. Thus, farm B's profits are said to be substantially more persistent than farm A's and it is much easier to believe that farm B's future profits will be $20/acre. On the other hand, it appears farm A's $20/acre profits might chiefly be due to chance and thus future profits may just as easily be more or less than $20/acre.

There are numerous computer software packages that can be used to statistically analyze data, however, most producers probably will not have a statistics package for routine use on their farms. Thus, all analyses in this paper are based on using the Excel[1] spreadsheet program. While this program has some limitations with regards to statistical

---

[1] Other spreadsheets such as Lotus or Quattro Pro are also capable of doing much of the same analyses. Excel is used here because it is cited as the most commonly used spreadsheet by producers we have worked with.

analysis, it is a very powerful tool and is sufficient for analyzing the data from most on-farm research projects.

This paper includes some basic statistical analyses of the on-farm research project previously discussed (i.e., planting population by hybrid with four replications). The yield monitor data for the two different hybrids and the four different populations for each of the four replications are given in Table 1. Prior to doing any statistical analyses of these data it is often helpful to graphically examine the data. Figure 3 is a scatter graph created in Excel showing the data from the four replications of the two different hybrids plotted against plant population. Based on an examination of this graph, it appears that hybrid B has had higher yields and that optimal population is probably between 26 and 30 thousand. The next step is to see if these observations are statistically significant.

As previously discussed, any research effort begins with an hypothesis. In this case two hypotheses come to mind: (1) there is no difference in yields between hybrid A and hybrid B (2) plant population has no effect on yield. The results of the statistical analyses will either support or reject these hypotheses. Before discussing any statistical analysis of this data a brief review of several statistical concepts is needed.

Table 1. Yield monitor data for 1996 hybrid by population study in Doniphan county, Kansas.

Replication	Planting Population (000)	Yield Hybrid A	Yield Hybrid B
1	22	183.4	176.6
2	22	178.2	182.3
3	22	158.9	174.6
4	22	162.9	147.6
1	26	181.2	199.7
2	26	187.4	187.8
3	26	174.1	187.0
4	26	164.9	174.2
1	30	188.4	180.7
2	30	186.1	196.3
3	30	176.0	189.7
4	30	173.5	186.7
1	34	189.9	211.1
2	34	164.9	184.0
3	34	156.4	168.5
4	34	143.4	165.8

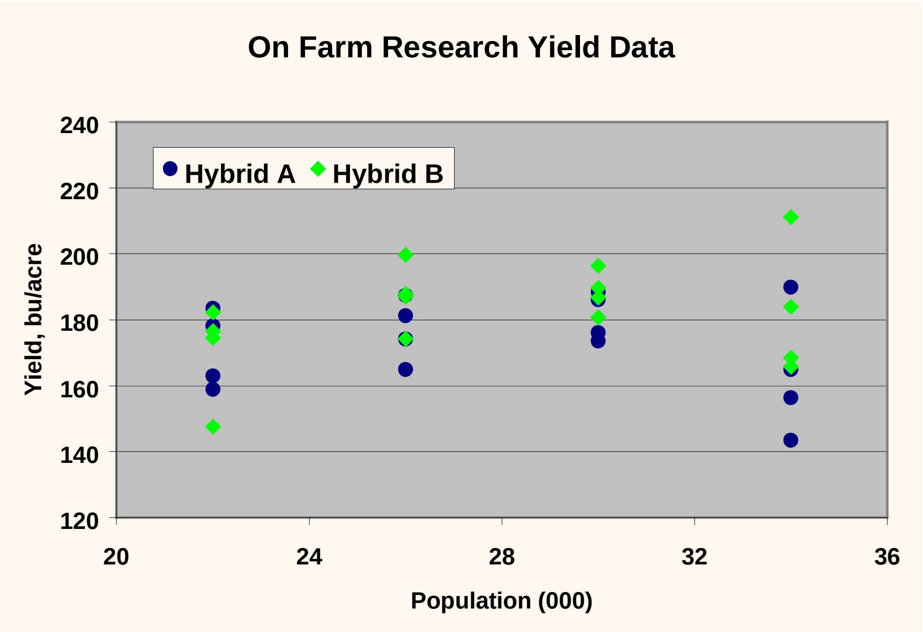

Figure 3. Yield monitor data from 1996 Doniphan County hybrid by population on-farm research project.

Normal Distribution

Many (but not all) statistical tests are based on the assumption that the data come from a normal distribution. Figure 4 shows an example of a normal distribution. If data follow a normal distribution, 68 percent of the observations will fall in the range of the mean (average) ± one standard deviation (Std) and 95 percent of the data will fall in the range of the mean ± two standard deviations. The assumption about data being normally distributed is often important for statistically testing results and making inferences about future values.

Standard Deviation

Standard deviation is a measure of the variability of the data. The greater the variability in the data, the greater the standard deviation.

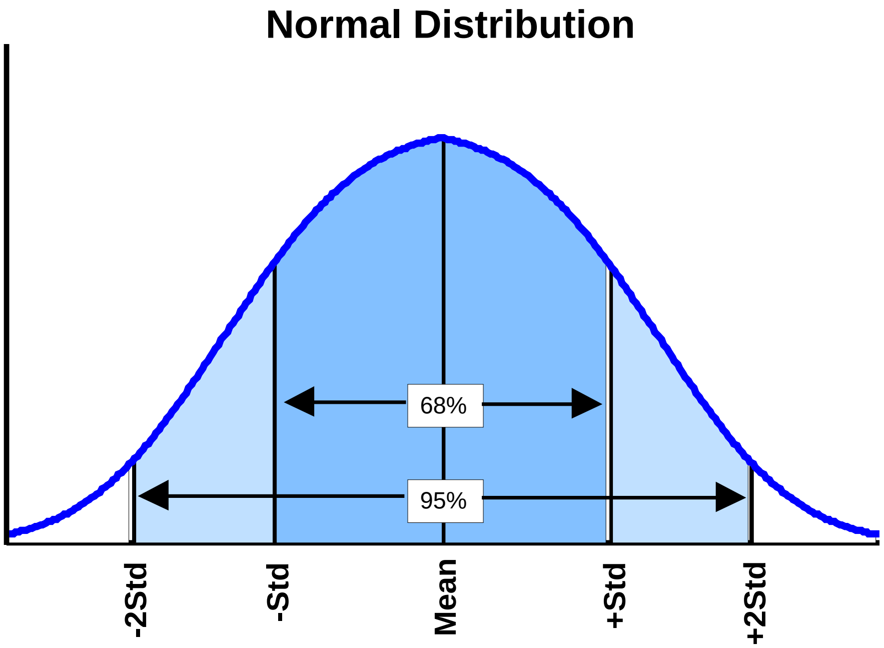

Figure 4. Example of a normal distribution.

P-Value

The *p*-value represents the probability of rejecting an hypothesis when it is actually true (i.e., we mistakenly reject a hypothesis) and [1 – *p*-value] is an indication of the confidence we can place on our results. For example, if we hypothesize that there is no difference between the yields of two hybrids and the *p*-value associated with testing this hypothesis is 0.05, we would conclude that there is a 5 percent chance that numerous repeated experiments of the same type would show that there actually is no difference. Put another way, we would be 95 percent confident that repeated experiments would reveal that there truly is a difference in yields between the two hybrids.

R-square (R^2)

A commonly reported statistic in regression analysis is the R^2. This value represents the amount of variability in the dependent variable (e.g., yield) that is explained by the explanatory variable(s) (e.g., hybrid, population). For example, if we fit a regression model with yield as a function of plant population and get an R^2 of 0.40, this indicates that 40 percent of the variability in yields is explained by plant population. The other 60 percent is due to something other than population.

Analyzing Data Using Excel

The Excel spreadsheet has a number of statistical analysis tools available for analyzing data, however, this paper will only discuss those that are deemed most applicable for the previously discussed research project. The data in table 1 will be analyzed with these tools. Though many methods will be discussed and compared, an appropriate analysis method should be chosen when the experiment is designed.

All of the statistical analysis tools used are accessed in Excel by clicking on "Tools" from the menu bar and then clicking on "Data Analysis…" The analysis tools considered in this paper, listed in alphabetical order, are: Anova: Single Factor, Anova: Two Factor With Replication, Descriptive Statistics, F-Test Two-Sample for Variances, Regression, t-Test: Paired Two Sample for Means, and t-Test: Two-Sample Assuming Unequal Variances. The results presented in the following tables come directly from Excel with only minor formatting changes made. Not all results of the analysis tools used will be discussed as some are beyond the scope of this paper.

Descriptive Statistics

The *Descriptive Statistics* tool in Excel is useful as it gives summary information of the data. The results of using this data analysis tool for all the data as well as for hybrids A and B separately are given in Table 2. The average yield for all populations and both hybrids is 177.6 with a standard deviation of 14.6. Thus, we would expect 68 percent of the yields to fall between 163.0 and 192.2 (177.6 ± 14.6). It can also be seen that, of the 32 observations, there were 67.7 bushels between the highest yield (211.1) and the lowest yield (143.4).

Hybrid B averaged 182.0 bushels per acre compared to 173.1 for hybrid A. However, the standard deviation of hybrid B was slightly higher than that of hybrid A, indicating hybrid B had more variation in yields. While the descriptive statistics tell us that the yield of hybrid B is higher, they do not tell us if it is significantly higher. That is, they do not tell us how much confidence we can place in the observation that "B yields more than A" – whether we expect it to yield more in repeated experiments in the future. More simply, is B's higher yield due to chance or can we be fairly confident that B will out yield A in the future? To answer this question we need to do some further statistical analysis. It should be noted that the information in Table 2 could also be obtained for the different population levels if desired.

Table 2. Results of using Descriptive Statistics.

	A & B	*A*	*B*
Mean	177.6	173.1	182.0
Standard Error	2.6	3.4	3.7
Median	179.5	175.1	183.2
Mode	164.9	164.9	N/A
Standard Deviation	14.6	13.4	14.8
Sample Variance	213.6	179.9	218.9
Kurtosis	0.4	-0.2	1.2
Skewness	-0.3	-0.7	-0.3
Range	67.7	46.5	63.5
Minimum	143.4	143.4	147.6
Maximum	211.1	189.9	211.1
Sum	5682.2	2769.6	2912.6
Count	32	16	16

t-Test: Two-Sample Assuming Unequal Variances

Given two sets of data, we can use the *t-Test: Two-Sample Assuming Unequal Variances* data analysis tool in Excel to test the hypothesis that there is no statistical difference between the means of the two data sets. In our example, we can use this tool to test our hypothesis that there is no statistical difference between hybrids A and B. Results of using this procedure for the yield data are given in Table 3. It can be seen that our hypothesized difference in yields is 0 and that there is only an eight percent chance we are wrong if we reject this hypothesis (*p*-value = 0.0835). In other words, we are approximately 92 percent confident that the yields of these two hybrids would not be equal in repeated experiments. Therefore, if we are comfortable with a 92 percent confidence level, we would conclude that the higher yield of hybrid B is not due to chance because it is statistically higher than the yield of hybrid A.

Table 3. Results of using t-Test: Two-Sample Assuming Unequal Variances.

	A	*B*
Mean	173.1	182.0
Variance	179.9	218.9
Observations	16	16
Hypothesized Mean Difference	0	
Df	30	
t Stat	-1.7903	
P(T<=t) one-tail	0.0418	
t Critical one-tail	1.6973	
P(T<=t) two-tail	0.0835	
t Critical two-tail	2.0423	

t-Test: Paired Two Sample for Means

The data from this research came from a split planter comparison. Therefore, the yields of the two hybrids at each population level are actually “paired” together. Because of this, we can use the *t-Test: Paired Two Sample for Means* data analysis tool to test our hypothesis that there is no statistical difference between the hybrids. The results of using this procedure for the yield data are given in Table 4. As before, our hypothesized difference in yields is 0, but now there is less than a one percent chance we are wrong if we reject this hypothesis (*p*-value = 0.0057). Because the data from our two samples are “paired,” with extraneous forces (everything but the hybrid treatment) impacting each of the two yields in a “pair” similarly, the *t-Test: Paired Two Sample for Means* is more appropriate than the *t-Test: Two-Sample Assuming Unequal Variances*. In this case, we are more than 99 percent confident that the yields of these two hybrids are not equal.

Table 4. Results of using t-Test: Paired Two Sample for Means.

	A	*B*
Mean	173.1	182.0
Variance	179.9	218.9
Observations	16	16
Pearson Correlation	0.6943	
Hypothesized Mean Difference	0	
Df	15	
t Stat	-3.2204	
P(T<=t) one-tail	0.0029	
t Critical one-tail	1.7531	
P(T<=t) two-tail	0.0057	
t Critical two-tail	2.1315	

F-Test: Two Sample for Variances

In addition to testing the means of the two data sets (i.e., differences in average yields for hybrids A and B), we can test whether or not the variability of yields for hybrid A is different than that of hybrid B. In Table 2 we saw that the variance and standard deviation of hybrid B were higher than those of hybrid A (variance = std^2). Similar to tests of average yields, we can test whether the observed difference in yield variability (variance or std) for hybrid A versus hybrid B can be expected in future plantings of these hybrids, or whether it was merely due to chance. To test whether or not the variances of two data sets are statistically different we use the *F-Test: Sample for Variances* data analysis tool. The results of using this procedure for the yield data are given in Table 5. As before, our hypothesis is that there is no difference in the variability of yields (i.e., variance or std). In this case, the *p*-value is 0.3544, indicating that we are only 64 percent confident that the yield variances of the two hybrids would not be equal in repeated experiments. Thus, we are very confident that the yields of hybrid B are greater than those of hybrid A and we are somewhat confident that hybrid B yields are more variable than those of hybrid A.

Table 5. Results of using F-Test: Paired Two Sample for Variances.

	A	*B*
Mean	173.1	182.0
Variance	179.9	218.9
Observations	16	16
Df	15	15
F	0.8218	
P(F<=f) one-tail	0.3544	
F Critical one-tail	0.4161	

Anova: Single Factor

Another way we can statistically test if there is a difference between factors is to use the ANOVA (analysis of variance) procedures in Excel. If we want to test the differences in yields between the hybrids (similar to the t-test procedures), the *Anova: Single Factor* analysis tool can be used. The results of using this procedure for the yield data are given in Table 6. It can be seen that the *p*-value associated with our hypothesis that there is no difference in mean yields of the two hybrids is exactly the same as that reported in Table 3. This implies that if we want to know if the means of two data sets are statistically different from each other we can use either a t-test or the ANOVA procedure and our conclusions will be the same. However, as previously noted, in this case we should use the *t-Test: Paired Two Sample for Means* analysis tool because we are working with paired data.

Table 6. Results of using ANOVA: Single Factor

SUMMARY						
Groups	*Count*	*Sum*	*Average*	*Variance*		
Hybrid A	16	2769.6	173.1	179.9		
Hybrid B	16	2912.6	182.0	218.9		
ANOVA						
Source of Variation	*SS*	*Df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	639.0	1	639.03	3.21	0.0835	4.17
Within Groups	5981.4	30	199.38			
Total	6620.5	31				

Anova: Two-Factor With Replication

The design of this particular research was that of two factors (i.e., hybrid and population) and each treatment was replicated four times. Therefore, the data from this study can be analyzed using the *Anova: Two-Factor With Replication* analysis tool. The results of using this procedure are reported in Table 7. Many of the results of using this procedure are summary statistics. For example, the output shows the average yield and the variance associated with each population-hybrid combination. The primary statistics of interest to us are the reported *p*-values. Examining the *p*-values allows us to reach the following conclusions: (1) we are 86 percent confident that there are differences in yields due to population (i.e., *p*-value = 0.1398), (2) we are 92 percent confident that the yields of the two hybrids are significantly different (i.e., *p*-value = 0.0766), and (3) we are only 43 percent confident that the effect population has on yield is different between the two hybrids, or put another way, the effect hybrid has on yields is different at different populations (i.e., *p*-value = 0.5733). The *p*-value associated with difference in hybrids is slightly different than before because we are now also including the effect of population, which was previously ignored, however, our conclusion has not changed.

Based on using the *Anova: Two-Factor With Replication* analysis tool, we know that the population level does have a statistically significant impact on yields. This suggests that we would reject our second hypothesis (i.e., plant population has no effect on yield). However, we do not know what the optimal yield level is yet. We also know that the effect population has on yields might be different for the two different hybrids, but we are less confident in this result due to the relatively high *p*-value.

Table 7. Results of using ANOVA: Two-Factor With Replication

SUMMARY	A	B	Total			
22						
Count	4	4	8			
Sum	683.4	681.1	1364.5			
Average	170.9	170.3	170.6			
Variance	139.2	239.2	162.2			
26						
Count	4	4	8			
Sum	707.6	748.7	1456.3			
Average	176.9	187.2	182.0			
Variance	93.5	108.6	116.8			
30						
Count	4	4	8			
Sum	724.0	753.4	1477.4			
Average	181.0	188.4	184.7			
Variance	54.0	42.1	56.6			
34						
Count	4	4	8			
Sum	654.6	729.4	1384.0			
Average	163.7	182.4	173.0			
Variance	384.4	431.7	449.7			
Total						
Count	16	16				
Sum	2769.6	2912.6				
Average	173.1	182.0				
Variance	179.9	218.9				
ANOVA						
Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample (population)	1123.44	3	374.48	2.01	0.1398	3.01
Columns (hybrid)	639.03	1	639.03	3.43	0.0766	4.26
Interaction (H x P)	380.21	3	126.74	0.68	0.5733	3.01
Within	4477.77	24	186.57			
Total	6620.45	31				

Regression

Because population is a continuous variable, as opposed to a discrete variable such as hybrid, we might benefit from moving into a regression analysis framework. Regression analysis is a method of estimating the relationship that exists between a dependent variable (e.g., yield) and one or more independent variables (e.g., hybrid and population). Based on our analysis up to this point, we might consider estimating the following model:

$$\text{Yield} = \beta_0 + \beta_1(\text{hybrid A}) + \beta_2(\text{population}) + \beta_3(\text{population})^2 + \varepsilon \qquad (1)$$

where β_0, β_1, β_2, and β_3 are the coefficients we want to estimate using regression analysis and ε depicts an estimation error. That is, β_0 - β_3 are placeholders – much as we use x in algebra – for numerals we need to derive. Hybrid A is a binary variable (also referred to as a dummy variable) and is equal to 1 if the yield came from hybrid A and equal to 0 if the yield came from hybrid B. Once we know the numerical values for β_0 - β_3 , we can plug in values for population, population^2, and hybrid A (1 if yield is from hybrid A, 0 if from B) to determine the expected yield.

Hybrid A is included in our model because we have shown that the yields between the two hybrids are significantly different from each other. Because hybrid A had lower yields than hybrid B, β_1 is expected to be negative. That way, when we plug in a 1 for hybrid A, it will result in a lower estimate of yield than when we plug in a 0 (for hybrid B). Population is included in our model because we have shown that yields are a function of population. Because we typically expect yields to increase as we increase plant population, β_2 is expected to be positive. Population squared is included because we expect the relationship between population and yield to be nonlinear (i.e., yield increases at a decreasing rate as population increases). This hypothesis is based on agronomic theory and is somewhat confirmed by visually examining the data in Figure 3. Accordingly, we expect β_3 to be negative. The error term (ε) is included in the model to recognize that yield will not be perfectly explained by the independent variables (hybrid A, population, and population^2).

To use the *Regression* analysis tool in Excel to estimate equation (1) we are asked to identify the Y and X ranges. The Y range is the dependent variable and would be a single column including all 32 yield observations. The X range is the independent variables and would be three columns including a binary variable for hybrid A, population, and population squared. As noted, the binary (dummy) variable for hybrid A is equal to 1 if the yield came from hybrid A and equal to 0 if the yield came from hybrid B. Results of using the *Regression* procedure to estimate equation (1) are given in Table 8. The R^2 of 0.26 indicates that hybrid, population and population^2 only explain about one-fourth of the variability in yields. This means that much of the variability in yields had to do with other factors -- possibly soil type, fertility, pest and weed pressures, etc.

Other values of primary interest in Table 8 are the coefficients and their *p*-values. The *p*-values are for the hypothesis that the coefficient estimate for a particular variable is equal to zero. The coefficient on hybrid A (i.e., β_1 in equation 1) is –8.94 indicating that hybrid A is expected to have yields almost 9 bushels per acre less than hybrid B. Notice that this is close to the overall observed yield difference of 8.9 depicted in Table 2. The *p*-value associated with this variable is 0.0655 indicating we are about 93 percent confident that there is truly a difference in yields (i.e., the coefficient is not equal to 0). This result should not surprise us given the previous t-test and ANOVA results. The coefficients on population and population squared both have the signs we expected (population is positive and population squared is negative). Further, we are 98 percent confident they are not equal to zero based on their *p*-values of about 0.02.

Table 8. Results of using Regression to estimate equation (1).

SUMMARY OUTPUT						
Regression Statistics						
Multiple R	0.51					
R Square	0.26					
Adj R Square	0.19					
Standard Error	13.19					
Observations	32					
ANOVA						
	df	*SS*	*MS*	*F*	*Significance F*	
Regression	3	1750.48	583.49	3.35	0.03	
Residual	28	4869.97	173.93			
Total	31	6620.45				
	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	-101.3	112.3	-0.90	0.3749	-331.4	128.8
Hybrid A	-8.9	4.7	-1.92	0.0655	-18.5	0.6
Pop. (000)	20.50	8.18	2.51	0.0182	3.76	37.25
Pop sq.	-0.3617	0.1457	-2.48	0.0193	-0.6602	-0.0632

Because it is often difficult to “visualize” the impact a nonlinear variable has on the dependent variable, it may be helpful to use the estimated coefficients to predict data that can be graphed. Furthermore, the regression framework makes it convenient to make inferences about populations other than those tested. But, inferences far beyond the minimum or maximum populations tested would not be appropriate. For example, in this case we can use the information in Table 8 to predict yields at various population levels for both hybrids and then create a graph of this data to identify optimal population. Figure 5 shows the actual yields of hybrids A and B as well as predicted yields based on the regression results reported in Table 8. Yields are maximized at a population of about 28,000. It can be seen that the predicted yield for hybrid B is about 9 bushels per acre greater than hybrid A at all population levels. It is important to note that this result is strictly due to the model that was estimated.

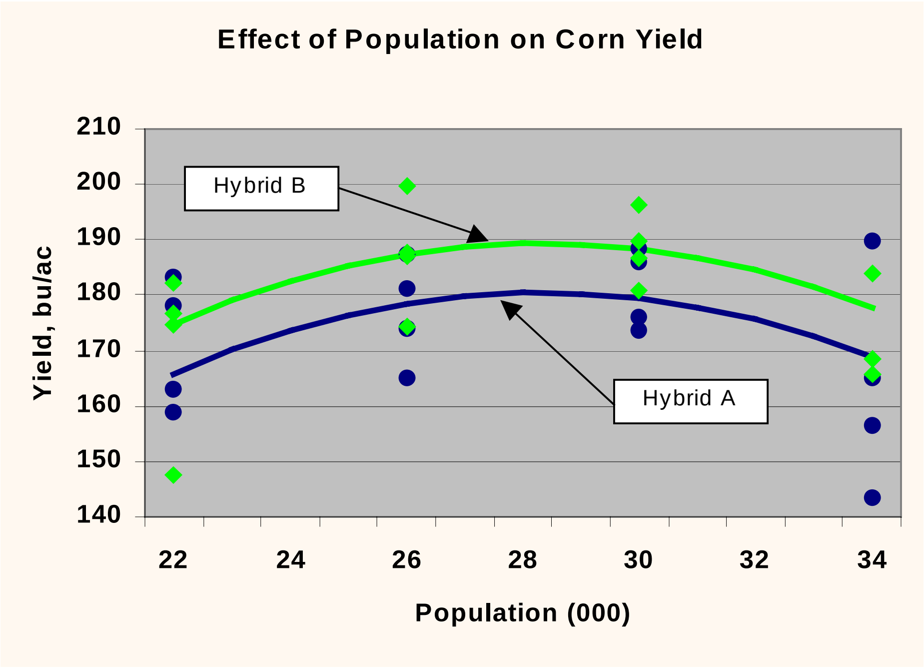

Figure 5. Actual and predicted yields for both hybrids (yields predicted using regression results reported in Table 8).

The ANOVA results presented in Table 7 indicated the possibility of an interaction between hybrid and population, which would imply a different optimal population for each hybrid. Remember though that we had considerably less confidence in this result. If we want to allow each hybrid to have a different optimal population, the models we need to estimate are the following:

$$\text{Hybrid A yield} = \beta_0 + \beta_1(\text{population}) + \beta_2(\text{population})^2 + \varepsilon \qquad (2)$$

$$\text{Hybrid B yield} = \beta_0 + \beta_1(\text{population}) + \beta_2(\text{population})^2 + \varepsilon \qquad (3)$$

where, all variables are the same as previously defined. The difference is that now we estimate a model for the yield data of each hybrid separately. As before, to use the *Regression* analysis tool in Excel to estimate equations (2) and (3) we have to identify the Y and X ranges. The Y range is the dependent variable and would be a single column including the yield data. The X range is the independent variables and would be two columns including population and population squared. Results of using the *Regression* procedure to estimate equations (2) and (3) are given in Tables 9 and 10, respectively. Notice that there are now only 16 observations in each analysis, since there are only 16 observations of yield for each hybrid. The population and population squared coefficients are significantly different than 0 at about the 10 percent level for both hybrids. In other words, we are about 90 percent confident these coefficients are not equal to the hypothesized value of zero. Thus, we are confident that yields increase at a decreasing rate as plant population increases for both hybrids, given the positive and negative signs on population and population squared, respectively.

In order to determine if the optimal population is expected to be different for the two hybrids, yields can be predicted based on the regression results in Tables 9 and 10. Figure 6 shows the actual yields of hybrids A and B as well as predicted yields based on the regression results reported in Tables 9 and 10. Yields are maximized at a population of about 27,000 for hybrid A, compared to 29,000 for hybrid B. Based on these results, the yields of the two hybrids are comparable at low populations but hybrid B appears to have considerably greater yields at the higher populations. While these are the yield relationships that we expect, we have not said anything about the statistical significance of this graphical difference in expected yield.[2]

Table 9. Results of using Regression to estimate equation (2) – Hybrid A.

SUMMARY OUTPUT						
Regression Statistics						
Multiple R	0.48					
R Square	0.23					
Adj R Square	0.11					
Standard Error	12.68					
Observations	16					
ANOVA						
	df	*SS*	*MS*	*F*	*Significance F*	
Regression	2	608.81	304.41	1.89	0.19	
Residual	13	2089.43	160.73			
Total	15	2698.24				
	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	-93.99	152.67	-0.62	0.5488	-423.81	235.83
Pop. (000)	20.04	11.12	1.80	0.0947	-3.98	44.05
Pop sq.	-0.3656	0.1981	-1.85	0.0878	-0.7936	0.0623

---

[2] A more complicated regression model was constructed to statistically test the difference in optimal population and it was found that we have less than a 10 percent confidence that optimal populations actually differ between the hybrids. This model has not been included as it is beyond the scope of this paper, for further information regarding this test contact the authors.

Table 10. Results of using Regression to estimate equation (3) – Hybrid B.

SUMMARY OUTPUT						
Regression Statistics						
Multiple R	0.49					
R Square	0.24					
Adj R Square	0.13					
Standard Error	13.81					
Observations	16					
ANOVA						
	df	*SS*	*MS*	*F*	*Significance F*	
Regression	2	804.16	402.08	2.11	0.16	
Residual	13	2479.02	190.69			
Total	15	3283.18				
	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	-117.51	166.29	-0.71	0.4923	-476.77	241.74
Pop. (000)	20.97	12.11	1.73	0.1069	-5.18	47.13
Pop sq.	-0.3578	0.2158	-1.66	0.1212	-0.8240	0.1083

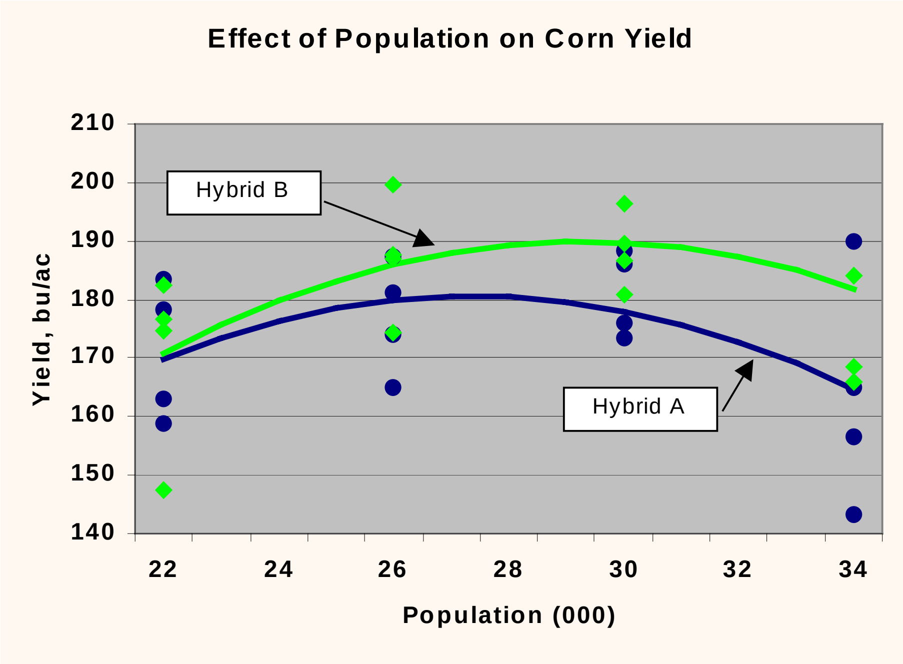

Figure 6. Actual and predicted yields for both hybrids (yields predicted using regression results reported in Tables 9 and 10).

Summary of Data Analysis

Research data is statistically analyzed to support, or reject, a specific hypothesis. For the on-farm research project considered in this paper the hypotheses we wanted to test were (1) yields are equal between hybrids and (2) population has no effect on yield. Based on the results of the statistical analyses presented, we reject both hypotheses. In other words, we found that the yields of hybrids A and B were statistically different from each other. Likewise, we found that yields were significantly affected by plant population with optimal population found to be about 28,000. When we allow population to have a different effect for each hybrid, we found graphically that hybrid A maximized yield at a population of about 27,000 and hybrid B at about 29,000. However, when this result was tested for statistical significance it was found that we have little confidence in this result (results of this test are not reported here). It is important to recognize that these results are all based on data from one year. Ideally, we would like to have more years of data to either confirm or refute these results.

All of the data analyses reported in this paper were done using the “Tools/Data Analysis” procedure in the Excel spreadsheet. Producers interested in analyzing their own research data are encouraged to enter the data in Table 1 into a spreadsheet and duplicate the results presented here as a check to make sure they are correctly using the analysis tools when analyzing data.

Summary

Yield monitors can be an effective tool for conducting on-farm research. Supplementing data collected with yield monitors with other information will increase confidence in results. Thought should be put into the design of the experiment and the analysis tools. Be meticulous about data collection and you have the opportunity to greatly expand your knowledge base on producing crops efficiently.