From Data to Decisions:
a Case Study in Variable Rate Fertilizer

6th Annual Kansas Precision Agriculture Conference, Great Bend, Kansas, January 15-16, 2003
NeATA Conference, Grand Island, Nebraska, February 3-4,2003
Terry L. Kastens, agricultural economist, Kansas State University
Dietrich L. Kastens, farmer and geographer
Draft January 2003

BACKGROUND

Slow adoption of VRA fertilizer

In the early 1990's, which were also the early days of precision agriculture (PA), the technology's pioneers pictured a world where farmers would quickly adopt geo-referenced on-the-go combine yield monitors. From there, farmers were expected to quickly adopt variable rate application (VRA) of fertilizer and possibly other crop inputs. In fact, university variable rate fertilizer trials were already in place in the 1980's and farm magazine articles of the mid 1990's were already touting the expected benefits of VRA.

From a yield monitor standpoint, it does appear that the early PA thinkers were more-or-less right. That is, figure 1 shows that, while the percent of custom harvesting combines having yield monitors has risen to around 60% by the 2001 harvest year, the portion with GPS has gotten larger as well (go to www.agecon.ksu.edu/kdhuyvetter to view the harvest reports from United States Custom Harvesters Inc.'s (USCHI) Custom Harvester Analysis and Management Program (CHAMP)). From a VRA standpoint, it appears they also were more-or-less right. That is, based on CropLife Magazine's 2002 dealership survey, figure 2 shows that nearly 60% of responding crop input dealerships were offering either 1-product or muli-product VRA services of fertilizers and chemicals by 2002.

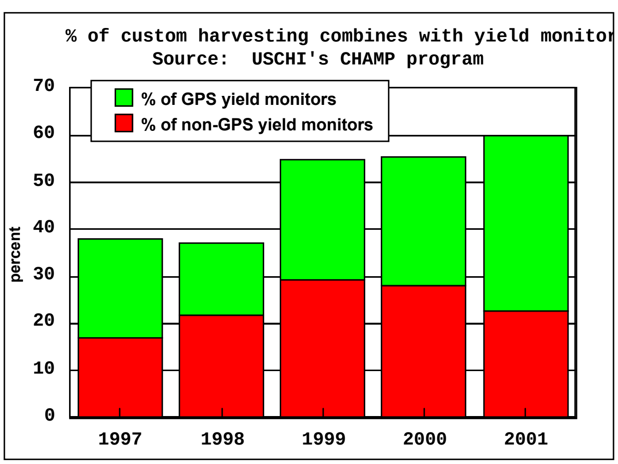

Figure 1

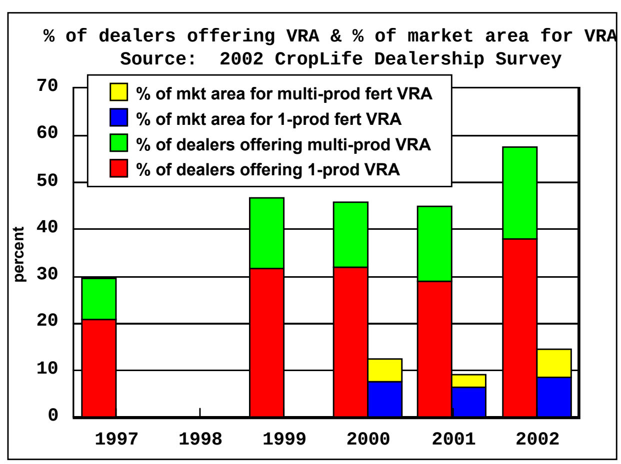

Figure 2

Though yield monitor and VRA offerings appear to have increased substantially over the last several years, it should be remembered that service offerings do not imply service use. That is, custom harvesters regularly note that, despite having access to yield monitors, many customers do not actually request the resultant yield maps. Further, as seen by the shorter bars in figure 2, the percent of the custom fertilizer application market area (i.e., percent of acres) actually penetrated with 1- or multi-product VRA, at under 15%, is still quite small. Although not shown in the figures, the CropLife survey suggested that 17% of customers' acres were being harvested with combines with yield monitors (unclear whether that meant with GPS-capable yield monitors). Further, although not shown here, the CropLife surveys of the last several years reveal that dealerships continually have been overly optimistic about their customers' future use of PA services. That is, each year reveals lower adoption rates than those anticipated by the dealerships in earlier years.

But, some PA technologies are adopted more quickly . . .

Despite seemingly slower-than-expected adoption of fertilizer VRA, some PA technologies seem to have been adopted rather rapidly. For example, the 2002 CropLife dealership survey revealed that 44% of the dealers now use lightbar guidance for their fertilizer applications (up from 24% in 2000). Further, it seems that anecdotal evidence from farmers also suggests rapid adoption of lightbar guidance.

It has long been known that early adopters of new agricultural technologies are the ones who profit the most. That is, technology-based positive economic profits are quickly bid into land values and rents. Higher rents mean higher costs, which mean that many farmers find themselves either out of business or that they have to adopt the new technologies just to break even. More succinctly, the non-adopting farmer finds that he is not "holding his own" economically, but rather that he is steadily going broke in the face of what he views as rents that are too high.

In the economic technology treadmill of modern agriculture over the last couple of centuries, why is it that some technologies are adopted more quickly than others? The answer lies in the fact that different technologies vary greatly in terms of a) the magnitude of expected profitability associated with them and b) the degree of confidence an adopter assigns to that magnitude of profitability. Technologies with small expected gains, but where those gains are very clear and distinct, such as Roundup-Ready soybeans, are adopted rapidly (according to Darr and Chern's Ohio State University study, U.S. Roundup-Ready soybean acres went from only 7.4% in 1996, the year of introduction, to 44.2% as early as 1998 – less formal sources suggest current numbers are much higher, perhaps 60% or greater). That is, Roundup-Ready soybeans could easily be assessed by comparing well-known higher seed costs against well-known herbicide cost reductions. More colloquially, these are belly-button or "duh" technologies as everyone can easily calculate their advantages. To some extent, lightbars are belly-button technologies. Their potential advantages in either cost or accuracy are easy to assess against their competition of foam or mechanical markers. It is likely that GPS-guided automatic steering devices will also be adopted relatively rapidly. That is, who can argue with savings associated with reduced overlap or even with the unknown-in-magnitude-but-clearly-positive reduction in operator fatigue associated with automatic steering devices? Any real (easier to get rented land?) or perceived (personal satisfaction) gain from straight rows is an added benefit.

Just as the gains for some technologies are obvious, the gains associated with others are illusive. Though it might be hard to imagine from today's perspective, tractors and fertilizer are two such examples. Tractors had to evolve for many years before they were clearly superior to draft animals. In the early years of fertilizer, many fields were still adequately fertile for the existing crop potential at the time. That meant farmers would regularly argue for years about the gains to commercial fertilization. A more recent example of such technologies is no-till farming. Many regions of the U.S. have been slowly migrating in the no-till direction for decades, as it becomes increasingly more obvious just where that technology might pay and as herbicides and related machinery become ever better along the way.

For many years production agriculture has been associated with substantial economies of size. That is, many of a farm's cost categories are associated with both fixed and variable costs. The per-unit (e.g., per acre or per bushel) cost of the fixed component always falls with increased farm size, providing an intrinsic economic advantage to larger farms over smaller farms. Over the decades, many agricultural technologies have been labor reducing (e.g., larger tractors). That means one operator has been able to farm more and more acres over the years, taking advantage of any other, non-labor, economies of size without increasing labor (e.g., volume discounts for input items). But, it also means that farmers have not generally gained experience in incorporating employees into their operations. Such discreet barriers (e.g., the reluctance to hire that first employee) lead to a vicious circle of giving preference to labor-saving technologies. Thus, for many farms, it is the labor-saving aspect of no-till farming that makes it appealing, rather than the higher yields or lower crop input costs. That is, farms move towards less tillage so that they can spread their "fixed" labor costs over more acres (by expanding their farms), thus reducing the labor cost per acre, and thereby allowing them to capture other non-labor economies along the way. Consequently, no-till adoption would be much slower yet were it not for the intrinsic labor savings. In short, technologies with associated labor savings tend to be preferred to those which might be labor increasing.

In the technology scheme, where does VRA of fertilizer or other crop inputs fit?

Increasingly, it seems that fertilizer VRA will be a slow-to-be-adopted technology. First of all, despite uniform application of fertilizer being a now-well-accepted technology, the gains to VRA are not well known. That is, the subtleties of yield response to fertilizer are not well known by agronomists or farmers. Though it might be generally believed that "approximately 100 lb/acre of N should be applied to this field given its level of soil test N," the expected yield difference associated with applying 10 lb more or 10 lb less than the recommendation is not at all well known. Further, because yield response curves are generally quite flat near the point of optimal fertilizer rates (e.g., small changes in yield for large changes in fertilizer), to be generally profitable, site-specific fertilizer application rates must ultimately be driven by site-specific non-fertility characteristics (shifting the yield response curve up or down). Yet, mathematical algorithms to translate such non-fertility characteristics into fertilizer rate recommendations are far from adequately developed. Further, the cost associated with gathering needed data is regularly viewed as too high relative to the expected payoffs. For example, against a total fertilizer cost requirement of perhaps $15/acre, a western Kansas wheat farmer might view the cost of 2.5 acre grid soil sampling ($7.30/acre according to CropLife's 2002 survey) to be "too high," at least in an annual framework. Moreover, interpolating between 2.5

acre centroids, often required for making site-specific fertilizer recommendations, can be highly inaccurate (see www.agecon.ksu.edu/kdhuyvetter for a copy of the Kastens and Staggenborg paper, Spatial Interpolation Accuracy, presented at 2002's K-State Conference on Precision Agriculture). Finally, to date, fertilizer VRA has not been viewed as a labor reducing technology. In fact, farmers who choose to engage in the required activities themselves likely would view the technology as a labor increasing one.

With all of the negatives just noted around fertilizer VRA, why might some PA adopters choose to move in that direction? Precisely because of those negatives. That is, because fertilizer VRA clearly is not a belly-button technology, its gains will not be quickly capitalized into land values and rents, rather they will accrue to the adopter for years and perhaps decades.

Given the foregoing discussion, it appears PA adoption can be divided into two classes. The first class involves those PA technologies whose economic gains will be fairly easy to perceive. Again, use of a lightbar is a good example. For such technologies, a quick "test drive," a quick calculation of expected gains, or a relative cost comparison with the nearest substitute is often sufficient to entice the farmer to participate. Certainly, years of data gathering are probably not required for making the adoption decision for such technologies. Indeed, farmers who deliberate too long on such technologies will find they have missed the profitability period. The second class involves those PA technologies whose economic gains are hard to delineate, and which probably require farm-level database building to help make decisions for the long run. Again, VRA of crop inputs is a good example. Because such PA technologies are intrinsically long-run, they should always be thought of in the context of other long-run agricultural technologies, for example those related to long-run soil fertility goals, tillage practices, and crop rotation strategies. Doing that might mean the costs of PA adoption can be spread over the gains from multiple technological changes, enhancing the profitability of adoption. Regardless, adoption of long-run PA technologies requires a substantial investment of time and money, meaning that a clear vision or purpose probably ought to attend.

THE CASE FARM

What follows is a case study of PA adoption. It is not meant to be definitive, but rather exemplary. Other farms might arrive at substantially different conclusions. The hope is that the reader might benefit from being exposed to the goals, the underlying reasons, and the actual activities undertaken by a particular farm as it has gradually adopted PA technologies.

General information about the case farm

The case farm is the Kastens farm in northwest Kansas (Rawlins County). The farm has a normal (30-year average) annual rainfall of around 21 inches. It has around 5,700 acres of dryland crop acres and about 3,700 acres of pasture. The pasture is used in a cow-calf program of around 300-350 cows. The main crops are wheat and corn, with a little alfalfa and forage sorghum, and occasionally a few acres of sunflowers or soybeans. Currently, approximately 1/5 to 1/4 of the cropland is fallow each year. The farm is rapidly moving towards 100% no-till and hopes to decrease the portion of fallow land over time. Nearly 100% of the farm's cropland is terraced with contour terraces, an historical must given conventional tillage and the silt loam soils of the area. However, the reduction in tillage has diminished the importance of terraces, leading to more and more back-and-forth farming.

Though the farm has historically benefitted from the symbiotic relationship between crop production and a cow-calf operation (e.g., crop residue grazing), it is beginning to question whether it might be better off reducing the impact of the cow-calf program on the crop production program – in terms of both labor competition between enterprises and in terms of following-crop yield reductions due to residue grazing.

People wise, the Kastens farm is comprised of a hands-on day-to-day manager guy (Gary Kastens), a business, economics, and analysis guy (Terry Kastens), a computer, geography, and electronic technology guy (Dietrich Kastens) and a can-do operational guy (Lester Yoos).

General goals of the case farm's PA program

Some PA technologies have been and will be adopted on the basis of little economic analysis. For example, a used (261 separator hours) JD 9600 combine with a GPS-based GreenStar yield monitor was purchased in December of 1997. The yield monitor investment was chiefly based on faith that it would be economically profitable. Similarly, a Starfire LB5 lightbar was purchased in 2001, with essentially no explicit economic analysis. However, it was believed that the lightbar experience would be useful for whenever the farm might need to consider guidance methods for future machines such as a fertilizer rig, a no-till drill, or a crop sprayer. In general, the farm will continue to make a number of PA adoption decisions based on only small amounts of analytical time. For example, the farm likely will invest in GPS-guided steering for at least one of its tractors in the next 12 months. The decision will be based on a quick analysis of swath overlap economics (savings of tractor and implement time, labor, and crop inputs), with some allowance for reduced operator fatigue.

In concert with the discussion in the first few pages of this paper, the farm's main PA thrust has focused on and will continue to focus on VRA of crop inputs – because such a focus is expected to result in per-acre profits that will be higher than those of neighboring farms for many years to come. Because VRA of crop inputs is intrinsically a long-term investment, it has been very slow and deliberate in its development on the case farm. Slowly moving machinery decisions have paralleled slowly moving agronomic, geographic, economic, and computational decisions. To date, the focus has been principally on moving towards VRA of fertilizer (N and P). However, it is expected that investments there will ultimately be used to aid VRA of other crop inputs (e.g., seeding rates and herbicide rates).

Dietrich and Terry are the two PA people for the case farm. As a geographer, Dietrich works mainly with the spatial aspects of the farm's PA program, typically working in the GIS software MapInfo. Dietrich works both at the front end (data collection) and at the back end (using software to guide the PA machinery, e.g., the VRA application). As an agricultural economist, Terry works mainly with the analytical aspects of the program. Using principally Matlab and spreadsheet software, he is positioned between Dietrich's front end and back end activities. Together, Dietrich and Terry are faced with having to convince both themselves, as well as Gary and Lester, that the farm's PA investment is worthwhile.

PA investments part of investments meeting broader goals

Since the farm anticipates increased use of no-till technologies, the PA machinery investments have been and will be considered in the broader context of changes in the machinery line to accommodate less tillage. For example, making the investment in a VRA fertilizer rig is more feasible when moving towards no-till means different fertilizer application methods anyway. Then, VRA machinery investment can be viewed as add-on investment rather than new-machine investment. Further, increased back-and-forth farming due to less tillage means it is easier to justify GPS-based guidance or steering. Finally, it should be easier to justify a VRA rig when the farm's goal is to build up soil test P and it knows that fields and sites vary substantially in regard to that measure (that goal was displayed especially by the Kastens, Schmidt, and Dhuyvetter paper <u>Combining Farm and University/Industry Information for Variable Rate Fertilizer Decisions</u>, available at www/agecon.ksu.edu/kdhuyvetter, and presented at the 2002 Kansas Precision Agriculture Conference).

A two-year informational search took place around the selection of a fertilizer rig that could simultaneously accomplish 2-product VRA and also work well in no-till cropping conditions. At first, the goal was to apply N as anhydrous ammonia because of the cost advantage over other types of N, along with liquid or dry P. Despite a large bias in the direction of anhydrous, over time it became apparent that finding or building a reliable no-till anhydrous rig was going to be too difficult. The decision was to mount Flexicoil openers at 15 inch intervals on a 9x6 (54 foot) undercutter frame and a used undercutter was purchased. Design problems with the undercutter, along with the fact that the openers would not allow application of fertilizer to growing wheat, led to an examination of university research comparing the N-efficiency for deep placement with that of coulter injection and with that of surface application. A small N-efficiency advantage of deep placement over coulter injection across the expected cropping program was weighed against the potential advantages of a coulter injection system over a deep placement system such as that planned with the undercutter. The expected advantages associated with the coulter injection system were that it would allow for faster travel speeds and less draft, and that it would have lower repairs. Finally, the farm settled on a 40 foot DMI coulter injection (15 inch spacing) machine for liquid fertilizer. A further key advantage to this machine is that it was factory- rather than shop-built, and that it has widespread use in the Midwest. That means that if would be easy to re-sell on the used market in the event it does not meet expectations – reducing the risk of investment adoption.

DATA AND DECISIONS

This section focuses on the data-collection and number-crunching aspects of the case farm. Likely, for a number of reasons, it is this part of any farm's PA program that causes the most consternation. First, it is labor intensive and, as noted earlier, farmers prefer adopting technologies that are labor reducing. Second, because of the endless possibilities, the result is frustration due to mental overload and repeatedly un-attained goals. Third, managers become frustrated with their own inabilities associated with data analysis – due to lack of the necessary formal education, training, or experience. Or, they become frustrated that they cannot locate trustworthy individuals who do have the abilities to analyze data. Fourth, they find themselves frustrated with software that repeatedly fails to meet their expectations.

What follows is a discussion about how the case farm has reduced the consternation and frustration associated with data collection and analyses. Unfortunately, because "everything builds on everything else" with PA, the most appropriate ordering of this discussion is not immediately clear. Consequently, we somewhat arbitrarily chose to develop the discussion around a number of principles that seem to have emerged on the case farm – principles that have helped the case farm continue moving forward in reaching its long run PA goals. Some of the principles are general in that they could apply to many or most farms adopting long run PA strategies. Others are more specific to the case farm. Regardless, they are all intertwined and they all overlap with each other. Finally, some of the principles are axiomatic commitments of the case farm and others have only been learned the hard way.

Principle 1: Adhere to scientific (causal and statistical) principles

We believe that ascribing and adhering to causal and statistical principles is itself an important principle of data collection and analyses. There are at least two good reasons for this. The first has to do with early adoption of technology. With little doubt, those managers who are making substantial progress in reaching their long-term PA goals are early adopters, which immediately implies that they have the *potential* to be more profitable than their neighbors for many years to come.

Typically, farm managers rarely are comfortable with making massive changes to their management – at least until such suggested changes are well documented scientifically. This is as it should be. After all, good farm management does not arise out of thin air, but evolves over time from substantial intuitive feedback. Essentially, most managers rely on substantial empirical evidence before making changes to their management. Besides from published research results, that evidence often comes in the form of personal experiences (observed treatments and associated yields, etc.). Yet, such evidence is scant for the early adopter. The profitable early adopter is the one who correctly assimilates important agronomic and economic principles or causal forces. As a case-farm-related example, consider the following. When the case farm adopted no-till corn production around 1989, it did not do that on the basis of observing other farms in the area. Such data did not exist! Rather, it did it on the correct assumption that no-till would save moisture that should ultimately result in increased grain production over conventional tillage. Long-term PA programs (e.g., VRA of fertilizer) are no different. If appropriate principles are adhered to, the associated changes will lead to profit.

The first reason for adhering to principles was more-or-less axiomatic, or faith-based. The second reason is that we actually have substantial empirical evidence that adhering to causal principles results in more profitable outcomes. For example, a mathematical model of crop yield that is constrained to be consistent with agronomic and economic principles (i.e., mathematically allowed to make only believable predictions from the standpoint of trained agronomists and economists) will be a more accurate model. Increased yield model accuracy implies more accurate crop input decisions and hence more profit. As evidence of this, the reader might consider Combining Farm and University/Industry Information for Variable Rate Fertilizer Decisions (available at www/agecon.ksu.edu/kdhuyvetter).

Principle 2: Bound the problem

Much frustration arises from the gradual realization that a data collection or analysis "problem" might be too large to ever get "solved" in one's lifetime or with one's current or future skills. Bounds come in many forms, one which has already been noted in principle 1. That is, agreeing to adhere to principles keeps a problem bounded in that no consideration is given to components of a problem that do not adhere to principles. For example, when yield causal principles are ascribed to, "turn the crank" unsupervised software-based database searches for poorly-understood, yet systematic, data relationships are sharply curtailed. This idea alone can easily save the manager hundreds of hours of frustrating computer work. More colloquially, there will be no searches for the Holy Grail.

Another way to bound the problem is to stay focused on smaller parts of a larger problem. For the case farm, this means focusing it's long-run PA efforts first on VRA of crop inputs, and more narrowly, on only VRA of fertilizer. Though a framework for consideration of VRA of non-fertilizer inputs is allowed to emerge over time, nearly 100% of the case farm's current analytical efforts are aimed at VRA of N and P. Staying goal-focused has another advantage in that it will keep the manager from constantly reviewing new software on the basis of the number of "cool" things it might be able to accomplish. Rather, more pointed questions will be asked of the software, such as, "Can it accomplish task A, which is needed to further my goal of fertilizer VRA?"

A third important way to bound the problem is to simply declare some parts of an analytical process as fixed, at least for a substantial length of time (as in several years). This is especially true for program features that lack clear scientific consensus. An example might be the size of a fertilizer management unit. Should it be a zone based on soil type? Should it be 2.5 acres? On this, because there has been no scientific consensus, a manager might simply declare, "My management unit shall be 1 acre. And, it shall only change in the presence of substantial evidence to the contrary. And, I have no intention of seeking that evidence for the time being." PA researchers and managers constantly wonder if they are doing the right thing ("maybe the management unit should be 3 acres . . . "). That is the nature of a new technology. But, too much wondering will lead to inaction and never getting the program implemented and hence never capturing the gains associated with early adoption.

Principle 3: Recognize that there probably are no comprehensive point-and-click softwares

Holy Grail-type software shoppers have been discussing their complaints across PA message boards and chat rooms for years. Consequently, this is not a principle that had to be learned the hard way for the case farm. Rather, the preponderance of empirical evidence has been sufficient enough. This issue likely is endemic within PA because PA adopters repeatedly find a lack of commonality among their software goals and expectations. Hence, it should not be surprising to find that software writers also lack the necessary commonality. This contrasts, for example, with more clearly purposed software such as accounting software. There, a manager is much more likely to find a comprehensive software meeting his needs. Once this principle is recognized, the manager can focus efforts towards finding specific softwares for accomplishing specific tasks.

This principle is not to say that comprehensive softwares do not exist, only that they will not be point-and-click. Certainly, those managers who are comfortable with code (e.g., macros or scripts) writing can easily find general softwares that have substantial flexibility in accomplishing desired tasks. For example, MapBasic is an add-in to MapInfo that allows for automating virtually any features possible in the GIS package MapInfo. Further, softwares like Matlab allow even more freedom yet to construct mathematical algorithms for desired purposes.

Principle 4: Long-run PA activities are labor-intensive, especially early on

This issue was discussed earlier in the paper. Failure to recognize it can cause self-deprecation for managers who repeatedly realize that "I haven't accomplished nearly as much as I expected by this point." Business partners begin to wonder what the PA guy has been doing. Likely, the worst scenario arises when a PA thrust is ended following some fixed amount of labor that had been allocated to the project. In short, PA labor should be considered a capital investment, much like an investment in land or in a tractor. This is not to say that PA must always remain a labor-intensive activity. As years pass, many tasks probably will become more automated. Most importantly, that an activity happens to be labor-intensive should be irrelevant. If it is expected to pay, the investment should be made.

Principle 5: Recognize that you may not be working with a comprehensive database

First off, we should make it clear that we use the word "database" to represent a digitally stored collection of data, not one that had been generated expressly using a traditional database program such as Microsoft's Access. In fact, the case farm has never seriously considered using traditional database software in database development, likely because a database software's unique aspects have not been considered sufficient to merit consideration. This is not to say that such an approach would necessarily be wrong, only that it has never been pursued either Dietrich or Terry.

This is one principle that the case farm had to learn the hard way. Initially it was the goal to combine many pieces of information into a common ascii text file, raw and transformed data alike, as well as those pieces of information needed for each of Dietrich and Terry. During that attempt, the number of data variables began to grow immensely, even when each year was allowed to "have" its own database. Further, computationally combining many diverse pieces of information into a common file with a common scale proved to be tedious and hard to debug. This is not to say such a goal is impossible to attain, only that it might be frustratingly slow getting there. Further, a large database can hamper electronic transmission over rural phone lines, something that is required given that Dietrich and Terry are 300 miles apart. Consequently, as database development continued for the case farm over a period of weeks and months, Dietrich and Terry repeatedly found themselves reminding each other to "quit thinking so big."

Part of the problem of database construction was associated with attempts to meet the needs of each of Dietrich and Terry – within the same database. Because Dietrich and Terry have PA sub-goals that are unique to their activities, and because they generally work with different softwares, problems arose in terms of data ordering and data type (textual or numerical variables). Eventually, Dietrich and Terry independently came to the same solution, as exemplified by the following statement from one to another: "Just give me the data you have in whatever format it is, and I will extract and transform the information I need. Then, let me know what information you need back, and I will be sure my output to you contains that somewhere." Finally, pulling this off required a number of data identifier columns so that data could always be re-ordered as needed.

The whole point to this principle is that, even after you have somehow transformed data to a common scale (e.g., point yield monitor data and 1-acre cell information to guide VRA), you still may want to have separate databases for separate tasks.

Principle 6: Consider different classes of data, by purpose

By now, crop consultants, crop input dealerships, and PA farmers alike each have scores of disks and CD's full of data, along with all of the guilt associated with "not having done many meaningful things" with those myriad data. The case farm is no exception. However, it does seem that it might be useful to think of different data classes, by purpose. For example, we might classify data as LaterData, IdeaData, or WorkingData. Essentially, this data classification scheme has emerged from problems encountered during the case farm database building efforts of Dietrich and Terry.

LaterData are those data which are inexpensive to collect buy which might have profit potential at a later date. One example might be simple tractor running time data such as time-of-day, ground speed, location, and elevation. Such data might very well be used in developing reasonable elevation maps at a later date that ultimately can be used to improve mathematical models of yield. They might also be used at a later date to calculate field operational efficiency measures, or simply as a documentation of historical tasks. LaterData should simply be collected and stored in their raw form, with no numerical transformations and no guilt attached to that fact. The goal with LaterData should be to accumulate the information with as little expenditure of time as possible.

Occasionally, some LaterData are transformed to **IdeaData**, typically in the form of maps. IdeaData are important to foster the creative ideas which might ultimately point a manager to decisions that provide him a comparative advantage over his competitors. But, it should be recognized that many LaterData will never reach IdeaData status, but rather disappear into oblivion. That should not be viewed as a large problem as long as the cost of LaterData collection was kept low.

WorkingData are those data with the highest degree of immediacy. They are destined to impact management decisions very soon, and hence the manager should not be hesitant to expend the labor required of collection. WorkingData typically can be stored and handled without viewing them in a map format. On the other hand, the end product of WorkingData might very well be a map, e.g., one for VRA fertilizer application.

Principle 7: Document everything

Because both Dietrich and Terry have worked in academic disciplines for a number of years, they have long recognized the importance of procedural documentation. Without procedural documentation it will be impossible for others to duplicate your work. And, more importantly, it will be most difficult for you to duplicate your own work at a later date. Data collection procedures, data transformation procedures, and procedures for making decisions from data typically arise or evolve from substantial efforts in creative thinking, deductive thinking, and consensus building among multiple individuals. Without documentation of those procedures and the thought processes behind them, the PA manager will constantly find himself second-guessing his prior decisions, trying to recreate his own and others' thought processes that "must have" led up to the current situation. Such re-creations are expensive and compete for one's creative energies. Beyond that, careful documentation makes it easier to decide which points might most benefit from an ongoing scientific debate among a farm's PA program participants. Finally, careful documentation is especially important if activities and results are to be meaningfully conveyed from an educator to a student or from a consultant to a client.

Principle 8: Be wary of excessive database automation in the early stages of databasing

In collecting, transforming, and aggregating data into meaningful databases, the temptation is to computationally automate all procedures. That is, as already noted several times, the tendency is to focus on labor reduction (or, in this case, less of a labor increase). But, computational automation can have a large initial fixed cost in the form of writing the necessary computer coding, or perhaps in making the necessary spreadsheet cell links, and especially in developing the necessary debugging techniques for ensuring quality results. Normally, this would be a good investment of time. However, in the early stages of PA program development, a number of dead ends often will be encountered, which are especially frustrating when they are encountered *after* substantial investment in computational coding.

In short, this principle suggests that PA managers should not loathe hand data entry for some data in the early stages of program development. Similarly, it could be better to break a computational exercise into smaller pieces, where each piece might be automated but where the whole is not. Smaller pieces of computation coding are much easier to debug.

Principle 9: Rely on others to provide constructive criticism

If one thing is clear with PA, it is the fact that nothing is clear; nothing is black and white; everything is gray. In that setting, defending one's PA program against constructive criticism is a most important task for ensuring its long run success in terms of ensuring that it will stay economically competitive. In the academic world, degree of "truth" is judged by acceptance of other researchers. Absent the critical eye of others in the discipline, research can quickly become dominated by preconception and bias, resulting in unexpected results and hence reduced profitability for those relying on that research. Farm-level PA programs are no exception.

Despite the importance of presenting one's thoughts to critical thinkers, PA managers might be reluctant to "give away their trade secrets." Although there may be some validity to that line of thinking, it probably should not be carried too far. First of all, a neighbor whose main program is a copycat one likely will always be behind you on the technology and profitability curve for three primary reasons: a) informational transfer takes time, b) some of the features of your PA program, since they were developed specifically from your farm's data, may not be directly transferable to your neighbor's farm, and c) informational transfer will be incomplete – your neighbor will miss some of your points and some of the changes you make over time. That means you will always have the benefit of "knowing your competition," and that you get the added benefit that someone else is "testing your system" and presumably providing constructive criticism along the way. Secondly, the negative aspects of neighbor-to-neighbor competition are mainly in the form of competition for owned and rented land, meaning they are not important for farms that are farming far enough apart that they don't directly compete for land. That leaves a considerable geographical area (perhaps up to the rest of the world) from which you can search out a critical eye.

A DESCRIPTIVE NARRATIVE ABOUT THE CASE FARM'S VRA PROGRAM

This section of the paper is a loosely tied together narrative of how the Kastens farm has approached its fertilizer VRA program in general, and data issues in particular. It is not comprehensive. Rather, certain points are delineated for the reader – points that might help those wishing to initiate or further develop the long run PA programs for their own farms. The format is phrasal bullets interspersed with text to support the decisions made.

a PA program with low annual costs but potentially large initial investment cost

Because of the generally low $/acre cost of fertilizer in a dryland cropping program (e.g., $15-$20/acre for wheat) there will not be much room for large annual expenditures associated with specifically the PA component of an overall fertilizer program. That likely would rule out such activities as annual 2.5-acre grid soil sampling. On the other hand, it would not rule out a substantial fixed cost component of the program which can be spread over many years and many acres, including those acres yet to be acquired in a farm expansion. Consequently, our program is heavy on the side of data and its numerical analysis. Further, such an investment is far from a belly-button technology, meaning that it should allow for positive economic profits for years to come.

large dependency on the investment aspects P fertilizer

Over the last several years, Terry and others (especially John Schmidt, K-State agronomist and Kevin Dhuyvetter, K-State agricultural economist) have conducted substantial research in the area of P fertilizer rates. Their work, which regularly has been presented at Kansas Precision Agriculture Conferences, predicts a substantial economic benefit from considering soil test P (STP) as a long term investment. In short, profits arise from "building up" STP in certain areas, with the build-up schedule determined optimally from the yield model, soil test information, and economic factors such as crop and fertilizer prices, interest rates, and expected length of land tenancy. It is this building up of STP that is expected to be the largest driver of profitability associated with the VRA program.

a production function (yield model) framework

Based on substantial research over the last five years by Kastens, Schmidt, and Dhuyvetter, we consider crop yield as a mathematical function of causal variables such as fertilizer and soil test N and P, soil organic matter percent, soil pH, soil texture variables, and soluble salts. The yield model will guide profit-maximizing fertilizer inputs. That means that discovering the mathematical yield model will be an important task of our data.

a 1-acre management unit

Because of the general lack of evidence in the science community supporting a zone management scheme, we have elected to work with a fixed-size management unit, and have declared it to be 1 acre. That means fertilizer rates will be constant for a particular acre, but can vary from acre to acre. Other than being a convenient legal unit, there is little special about one acre over say, 2 acres. However, with the current size of our machinery (30' drill, 40' planter, 40' fertilizer rig), it seems the expected overlap between acres will not be too excessive (an acre has 209 feet on a side). Moving towards herbicide VRA might mean 1-acre management units will be too small given sprayer boom width. On the other hand, individual nozzles might be varied. Further, herbicide VRA is probably a few years down the road for the case farm. And, the case farm does not intend to apply fertilizer with a spray rig. Because each 1-acre fertilizer rate will be determined by the yield model, that means an estimate of each of the causal variables for each acre for each year a crop is grown must be obtained. Hence, that is another important task of our data.

yield data from MPGM program

Several years ago the MPGM (multi-purpose grid mapper) program was developed by Dietrich, Randy Taylor, K-State agricultural engineer, and Terry. This program was designed to a) ensure that a field's site-specific yields sum to the total yield (production) supported by elevator scale tickets, b) correct certain yield monitor measurement errors, c) correct for yield errors associated with swath width measurement errors generated in fields with numerous point rows, and d) to assign yields to every location in a field, even when one of the combines is harvesting without a yield monitor. The MPGM assigns a yield to each 60 foot square cell in the field. Then, as needed, these yields can be aggregated to larger areas (as when yields are required for each acre on the farm) or assigned to specific points such as soil sampling points (e.g., by using a nearest neighbor routine).

soil sampling based on 3 samples per field

At the 1999 Kansas Precision Agriculture Conference, Terry presented a paper entitled <u>Site-specific Fertilizer Management with Less-than-perfect Information</u>. That paper briefly introduced the idea that a mathematical analysis technique referred to as maximum entropy could be used to extract valuable information from only a few measured points per field – 3 to be exact. In the year before, the case farm had begun an annual soil sampling routine that depended on 3 soil samples per field, a practice which has continued to date.

In the 3-samples-per-field sampling scheme, one sample is a field composite with geo-referenced collection sites (about 25-30 cores from 8-10 sites in a field) and the other two are site-specific samples (localized composites from around the sampling vehicle) taken from opposite sides of the field. The composite soil test is assigned to each of the 8-10 collection sites when used for yield model estimation (providing 8-10 data points), and it is also used to represent the whole field when needed. The two site-specific samples are used for annual fertilizer trials and also for yield model estimation. Collectively, assuming 8 collection sites for the composite sample, one each for the two site-specific samples, and that around 15 fields per year per crop are soil sampled, this sampling scheme has the potential to add about 150 data points to the yield model estimation routine each year, for each of wheat and corn.

two unique sets of data and two unique goals

That data collection, transformation, and analysis should have a purpose or goal is implicit within the core principles discussed earlier. For the case farm, this means two distinct data frameworks with two distinct goals.

The first data framework is the one just described – building a core data set comprised of soil-test data, crop yields, and other variables as they become available. Each yield in this data set is the average MPGM yield from the area immediately surrounding a site-specific soil sampling point or one of the composite sample collection sites. This data framework has one purpose, to be used for estimating a wheat yield model and a corn model for the farm. With each year come new data points for use in improving the yield model. Geographically speaking, since measured soil-test information is only available for a fraction of the farm's cropland, this data set covers only a fraction of that total cropland. But, inferences from the resultant yield models are used everywhere on the farm. Currently, the wheat and corn databases each contain around 500 observations.

The second data framework comprises a separate data set for each harvest year, with a unique observation for each 1-acre management unit (cell) of the farm. Each of these data sets have around 7,100 observations, which is more than the total acres of cropland because field boundary cells typically are less than 1 acre. Geographical coordinates for cell centroids, along with various numerical identifiers, tie each data set to those of previous years or future years. In this data framework, each cell with a harvested crop is assigned a crop yield from average MPGM yields in that cell. Each 1-acre cell will contain an estimate of each variable value needed to determine fertilizer rates. Cells that happen to contain measured soil-test data will be assigned those values rather than ones which have been estimated. This data framework ultimately is used to develop the fertilizer prescription maps needed to guide the VRA application of N and P.

little reliance on spatial interpolation

Based on the research we have reviewed (for one example see the Kastens and Staggenborg paper Spatial Interpolation Accuracy), we believe that spatial interpolation techniques are significantly less accurate than potential users are led to believe by the nicely-smoothed colored maps so readily provided by the PA software industry, especially when interpolation is across 2.5-acre grid cell centroids. Consequently, our PA program is designed around another process, referred to as the KSD approach, which stands for Kastens, Schmidt, and Dhuyvetter, who have been working on the technique (though no formal paper is yet available).

The KSD approach starts with a causal yield model. However, the causal variable values for each 1-acre cell, which are required to make VRA prescriptions in this framework, are not measured values but rather must be estimated. Fortunately, there is some information available for each 1-acre management unit, for example crop yield and fertilizer rates. The yield model itself is combined with this limited amount of spatially dense information (yields and fertilizer rates) using maximizing entropy – to provide reasonable estimates for each causal variable. Then, the entropy-based predictions are averaged with crop-removal based estimates of soil test variables to further improve the accuracy associated with predicting the causal variable values. Finally, the composite soil tests are used to keep predicted 1-acre soil factor values from straying too far from that which is believable.

use of electrical conductivity (EC, i.e., Veris) data in the future

The KSD approach to estimating the causal variable values for each 1-acre management unit can especially benefit from acquisition of more spatially dense data. Consequently, the Kastens farm recently has purchased a Veris 3100 cart to collect spatially dense information on electrical conductivity. These data are expected to improve the accuracy of the causal variable value estimates, as well as the accuracy of the yield model itself. Other such spatially dense, yet inexpensive, information that probably will be included in the future are elevation data and satellite-based remote sensing imagery.

software used

Dietrich primarily works in the GIS software MapInfo and Terry primarily works in the matrix language program Matlab. MapInfo is a software well-adapted for map making and works well with textual variables as well as numerical ones, but it is a less-than-adequate program for numerical analysis. Matlab is especially adapted to number crunching (numerical analysis) but cannot handle textual variables and is not a GIS. Both Dietrich and Terry are able macro (computer code) writers, and rely heavily on those techniques to guide their softwares. Finally, FarmWorks software is used to collect field operation information and too create fertilizer prescription maps for guiding the DMI machine's VRA application of fertilizer, and Red Hen software is used for geo-referenced soil sampling.

In addition to working with MapInfo and Matlab, Dietrich and Terry each depend heavily on simple computer spreadsheets (e.g., Excel or Lotus). Spreadsheets are useful for error-checking, finding data outliers, and especially for determining exactly what it is you want to do with the data. That is, spreadsheets expediently allow for all of the required "what if" situations you might encounter in your analysis. Then, the spreadsheet procedures are duplicated in the other softwares of interest (writing macros), and the spreadsheet serves as an error check on the process. For example, it expedites questions such as "Does my Matlab run get the same answers as my spreadsheet?"

Spreadsheets also are probably the best-understood numerical software analysis packages available. That means spreadsheets are a good way to explain procedures or actual data (e.g., recommended fertilizer rates) to others. A main disadvantage with spreadsheets is that, when spreadsheets contain numerous and complicated mathematical formulas, they can easily get quite large even by today's computational standards. Because spreadsheets are such an intuitive way to store and analyze data, after getting the necessary information from Dietrich (generated in MapInfo), Terry attempted to do all of the analyses needed for calculating VRA fertilizer rates within a spreadsheet – even the maximum entropy parts.

Although the process of generating all numerical analyses in a spreadsheet did work, just to turn out the farm's recommended fertilizer rates for 2003 corn required one spreadsheet of 75 meg, one of 130 meg, and about four that were much smaller (5-10 meg each). From a computational standpoint and from a data transmission standpoint, these file sizes are not feasible. That is, the spreadsheet unpredictably closes itself out with no input from the user and the computer regularly crashes under the computational load. Consequently, and despite our desires, most of the data analysis work was carried out in Matlab. However, the recommended fertilizer rates are provided by Terry to Dietrich in a spreadsheet form. Then, Dietrich converts the fertilizer recommendations (for each management cell) to a prescription map to guide the VRA fertilizer rig.

a program where accuracy and profitability should improve over time

Because our PA program is heavy on numerical analysis and weak on measured soil-test data, it is crucial that our KSD-based procedures will at least be accurate enough to be profitable. Actually, our desired criteria are more stringent than that. In particular, we desire the following. First, the KSD approach should result in improved accuracy of measurement over time, at least for those causal factors that are not expected to be impacted by management. Second, it is hoped that the KSD approach will result in prediction accuracy that is similar to what it might be with annual 2.5-acre grid soil sampling. Third, we expect profits associated with the program to increase over time. Fourth, we expect the KSD approach to result in substantially more profit over time than an annual 2.5-acre grid sampling approach.

a simulation test of the program

In the overall development of the KSD-based fertilizer VRA program for the case farm, implementation of the recommended site-specific N and P fertilizer rates is just now beginning, with the first crops to be impacted being those that will be harvested in 2003. Consequently, even if the farm chose to test its decisions by operating part of its acres with a KSD-based program and part with some competing approach, say 2.5-acre grid soil sampling, it likely still would be many years before a reliable accuracy and profitability comparison of the competing approaches might be undertaken. Thus, to provide a more immediate evaluation of the expected economic impact of the farm's fertilizer VRA program, an extensive numerical simulation was recently completed. A full discussion of the statistical simulation process noted in the previous paragraph is beyond the scope of this paper, so only a brief description follows.

Several competing fertilizer programs were tracked in a wheat production simulation of 10 hypothetical fields, with 500 data points each, over 19 years (actually 19 future crops, with each cropping intervals assumed to be 1.5 years to mimic a wheat-corn-fallow cropping sequence). To ensure numerical stability of simulation results, the simulation was repeated three times, where each of the three runs used approximately 170 hours of computer run time. Unlike some PA research simulations, which might assign ("assign" means randomly draw from a statistical distribution) a single "true" soil test value to a whole finite land area such as 2.5 acres (potentially biasing economic results in favor of VRA), this simulation assigned a different "true" soil test value to each data point in the field. What this means is that no posited data collection scheme is assumed to be perfectly accurate. For example, although 2.5-acre grid soil sampling will be more accurate than using a single composite soil sample for the whole field, it is still an inaccurate measure of the true soil test values.

The simulation exercise was based on a wheat model where wheat yield was assumed to be a function of fertilizer N (fertN), fertilizer P (fertP), soil test N (STN), soil test P (STP), soil organic matter percent (OM), sand percentage (SAND), clay percentage (CLAY), soil pH (PH), and soluble salts (SALT). The wheat model was assumed to have substantial error about it, mimicking reality, where weather events largely drive crop yields. "True" values for OM, SAND, CLAY, PH, and SALT were assigned once and then assumed constant across time. On the other hand, "true" assigned values for STN and STP were dependent on previous crop removal of that nutrient and previous crop fertilizer rates, rainfall (random) from fertilizing date to estimate date for STN (to approximate N loss), and random error. Relative to the other factors, this made predicting STN and STP a moving target.

In the simulation exercise, 10 short-run and 10 long-run fertilizing schemes (N and P) were considered, where each depended on different informational assumptions. For example, to get at potential profitability, one scheme (albeit an impossible one) assumed that fertilizer rates were prescribed with perfect soil test knowledge everywhere. Another assumed uniform fertilizer rates prescribed by a 25-core field composite soil sample for each field each year. Another assumed that 2.5-acre fertilizer rates were prescribed by annual soil samples taken from the centroid of each 2.5-acre cell. Another assumed that 1-acre fertilizer rates were prescribed by the KSD approach, where maximum entropy was used along with an annual 25-core field composite soil sample for each field. And so on. Long-run schemes were duplicates of short-run schemes, except that the farm manager was assumed to consider fertilizer P in a multi-year profit maximization context rather than in a single-year one. A 10-year rolling time horizon was assumed for long-run schemes. That is, each year, the manager assumes he will have control of the land for exactly 10 more years, and then applies fertilizer P accordingly. The next year he again assumes he will have the land for 10 more years, and so on.

It should be noted that, in the simulation of 2.5-acre annual grid sampling, we assumed accuracy commensurate with the best interpolative method uncovered in the Kastens and Staggenborg paper on interpolation accuracy noted earlier. Thus, we are assuming that the grid sampling manager has prior knowledge about which interpolation method to use (e.g., among methods such as kriging, inverse distance squared, etc.). This biases the results in favor of grid sampling. On the other hand, the Kastens and Staggenborg study of interpolative accuracy was based on predicting hypothetically discarded data points in actual grid sampling studies (most used 2.5-acre cells). Since a typical discarded data point was, by definition, a substantial distance away from the points used to make its estimate, this biases the results against grid sampling. That is because the typical data point to be estimated with a 2.5-acre grid sampling approach would be one-half the distance between 2.5-acre centroids, not the full distance. Of course, whether or not interpolative methods can predict close-up points substantially more accurately than more distant points is based on the nugget effect (how variable are soil cores taken only inches apart?), which can be quite large. All in all, we hope our two intrinsic offsetting biases more-or-less balance each other out, resulting in a reasonable characterization of a real time 2.5-acre grid sampling program's accuracy.

simulation results

What did the simulation reveal about accuracy of estimates of soil test values? As a measure of accuracy we use RMSE (root mean squared error), which is the square root of the average squared error of a series. Figure 3 depicts the prediction accuracy for one typical variable, OM, across three informational schemes. Clearly, both the grid sampling approach and the KSD approach are more accurate than using a single composite soil sample to represent a whole field. Also, the figure makes it clear that grid sampling is more accurate than KSD in early years. But, since accuracy improves over time with the KSD method, it eventually becomes more accurate than an annual grid soil sampling regime.

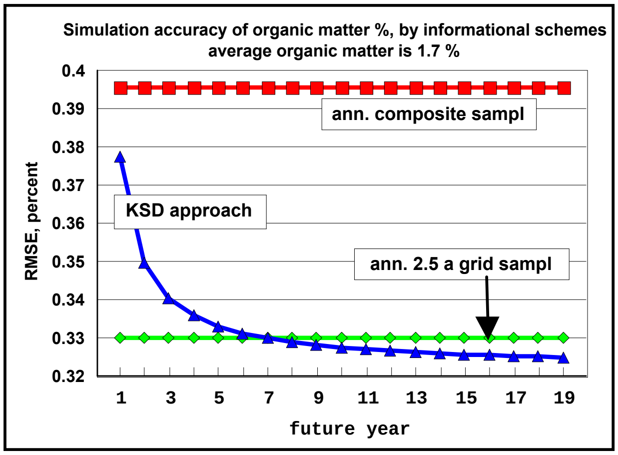

Figure 3

The simulation revealed that KSD accuracy improves over time for each of the fixed soil factors (OM, SAND, CLAY, PH, and SALT). That is not true for STN and STP, which are essentially moving targets. For STN, KSD prediction accuracy is approximately constant across time and the accuracy is slightly worse (higher RMSE) than it is with annual grid sampling. Of course, both methods are more accurate than using an annual field composite sample. For STP, KSD accuracy is also slightly worse than annual grid sampling accuracy. Again, both methods are more accurate than an annual field composite sample. For STP, both the KSD and grid sampling approaches get less accurate over time, at least when accuracy is measured in terms of RMSE. However, all informational schemes tested result in building up STP over time, as seen by figure 4. Consequently, it should not be especially surprising if accuracy of STP estimation, as measured by RMSE, gets worse over time (not shown). It should be noted however, that relative accuracy, as measured by RMSE divided by average STP each year, stays nearly constant across time.

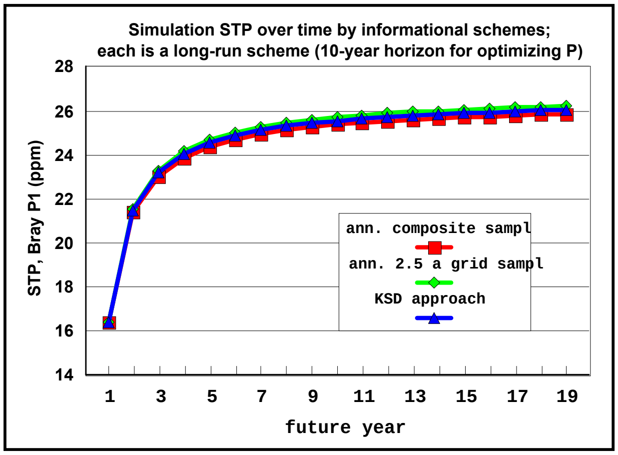

Figure 4

What about profitability. Figure 5 depicts that story, showing the expected returns over fertilizer costs associated with the grid sampling scheme and the KSD scheme, and comparing them both to a "perfect information" system, where all soil factor values are known with certainty (but not weather-related factors). The benchmark against which the three lines in figure 5 are to be compared is the short-run annual composite soil sampling scheme (optimal uniform VRA). By that benchmark, STP also builds over time, but not nearly as fast as using the three long-run methods depicted in figure 4. That is, STP in year 19 is around 19 ppm (not shown) rather than the 26 ppm shown in figure 4.

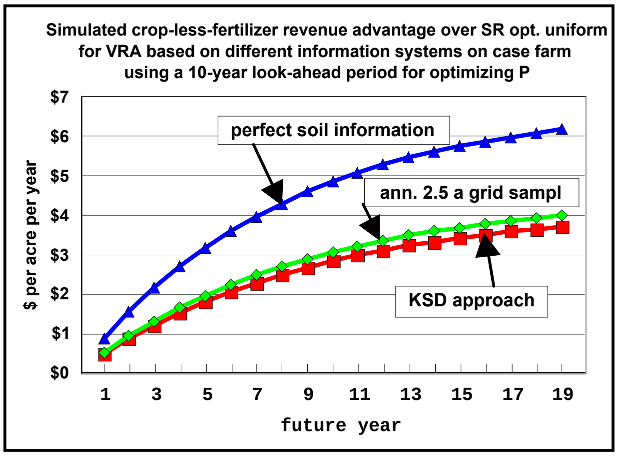

Figure 5

Clearly, figure 5 shows that the KSD approach is expected to have returns to fertilizer that are similar to the grid sampling approach, and that they are increasing over time as desired. What is not shown in the figure, however, is that the grid sampling approach would have annual costs associated with collecting the samples and obtaining the laboratory analyses. Those costs might easily come to $5-$10 per acre per year, which would quickly negate the PA benefits shown in figure 5. That might be the reason 2.5-acre annual grid sampling has not become widespread in dryland farming of northwest Kansas. On the other hand, the informational costs associated with the KSD approach are principally the costs of numerical analysis, which, on a per acre basis, should fall over time as procedures become more automated, and which also should fall as the farm expands because they are largely driven by a fixed cost component (education and computer power). Finally, if a "less-frequent-than-annual 2.5 acre grid sampling line" were added to figure 5 it would be lower than the corresponding annual line. Also, if a "less-frequent-than-annual 2.5 acre grid sampling in conjunction with the KSD approach" line could be added, it would be above the corresponding KSD line.

For the case farm, the "profits" associated with the KSD scheme in figure 5 have been deemed large enough to proceed with the related VRA of fertilizer program. But, there are a number of reasons that figure 5 may be understating the expected profits associated with the overall KSD-based fertilizer VRA program. First, recall that even with the optimal short-run uniform fertilizer program, STP increased from 16 to 19 ppm over the 19 years examined. That is because it was profitable to do so. The profit associated with that aspect of the PA program is not shown in figure 5 as it is already embedded in the benchmark scheme (the 0 line of figure 5). Without pursuing a PA research and development program on the farm over the last several years, would we have known to build up STP over time? Maybe, but maybe not. That means that the profit embedded in the benchmark (perhaps $1-$2 per acre – future work will test this) might also be *due to* the farm's PA program. Second, it is expected that additional inexpensive spatially dense data (e.g., EC, elevation, remote sensing) will improve the accuracy of soil factor estimation, pushing the red and green lines of figure 5 higher as a larger portion of the pie is captured. Third, these spatially dense data should improve the yield model's accuracy, making the overall pie larger (pushing the blue line up). Since the KSD approach to estimating soil factor values depends heavily on the yield model, anything that raises the blue line probably will raise the red line of figure 5 as well.

Will (can) the case farm's PA program work for other farms?

The short answer is yes. Anyone can choose to follow the research and procedures discussed in this paper to duplicate a similar program on his farm. A more realistic answer is no. Most farm managers will not choose to pursue such a numerically intensive PA program, which intrinsically depends on a large investment in human capital (education). But, they might very well choose to purchase such a PA program, assuming that market forces cause such programs to eventually become available through consultation or through software. Regardless, it is expected that farm managers developing their own PA program should benefit from adhering to a number of the principles set forth in this paper.

SUMMARY

Providers of precision agriculture (PA) services have repeatedly overestimated the adoption rate of a number of PA technologies. For example, early players in the PA movement expected adoption of variable rate application (VRA) of crop inputs to occur quickly. But, after nearly ten years have now passed, this does not seem to be the case. Yet, some PA technologies, such as lightbars, seem to be adopted much more quickly. This paper discusses some of the reasons that different PA technologies tend to be adopted at different rates. The short answer is that it is easier to calculate the expected economic gains associated with some technologies. Roundup-Ready soybeans is an example, where small but sure profits led to high adoption rates over night.

Early adopters of technologies garner positive economic profits. Yet, these gains are short-lived for belly-button (everyone has one) technologies such as Roundup-Ready beans – because they are quickly capitalized into higher land values and rents, driving up production costs accordingly. With PA, the farm manager wishing to have an economic edge over his neighbors that will last for years and decades must pursue those PA activities that are by nature hard to understand. VRA of crop inputs is an example of such an activity. Farmers, service providers, software writers, and researchers alike struggle to understand the complex relationships underlying crop yield; and, that understanding is sorely needed to acquire profits from crop input VRA.

By now, many potential PA practitioners have been gathering data for years. That has been the easy part of the PA movement. Now, these same practitioners find themselves drowning in a sea of data, not knowing which data are valuable, which data should be analyzed, how they should be analyzed, and what should be the resultant decisions. In hopes of providing increased understanding by example, this paper carefully delineates the thoughts, reasons, procedures, and mistakes for one actual case farm as it has tried to move forward over the last several years in developing a profitable PA program centered around VRA of fertilizer N and P. The case farm is a dryland wheat and corn farm in northwest Kansas. First, a number of general principles are discussed that should be broadly applicable to anyone struggling with data collection, analysis, and decision making issues. Next, a more detailed case-farm-specific descriptive narrative is provided. Finally, the paper closes by describing the results of a simulation exercise designed to help determine the expected profitability associated with the case farm's PA program.

The case farm is somewhat unique among PA practitioners in that it has intentionally (because of underlying scientific research) avoided dependence on spatial interpolation procedures. Instead it has focused on procedures that depend on a mathematical yield model estimated from farm data, where crop yield is characterized as a mathematical relationship depending on a number of soil factors such as fertilizer N and P, soil test N and P, soil organic matter, soil texture, soil pH, and soluble salts. Further, rather than relying on intensive and expensive site-specific soil sampling procedures (grid soil sampling) to estimate the values of yield-causing soil factors, the case farm has relied instead on complex numerical and statistical procedures to extract meaningful site-specific information from only a small number of laboratory-analyzed soil samples per year. Based on an extensive simulation exercise, this paper has shown that the case farm's approach to fertilizer VRA is expected to be profitable, and substantially more profitable than an approach depending on grid soil sampling. Consequently, the farm described in this paper is moving forward with the VRA procedures it has developed. It expects to garner positive profits related to its PA program for many years to come. And, over the next several years, it actually expects those profits to increase.