Searching for Efficiency: Least Cost Nonpoint Source Pollution Control with Multiple Pollutants, Practices, and Targets

Sergey S. Rabotyagov
Email: rabotyag@iastate.edu

Hongli Feng
Email: hfeng@iastate.edu

Catherine L. Kling
Email: ckling@iastate.edu

Manoj Jha
Email: manoj@card.iastate.edu

Department of Economics and
Center for Agricultural and Rural Development
578 Heady Hall
Iowa State University
Ames, Iowa, 50011-1070
USA

Selected Paper prepared for presentation at the American Agricultural Economics Association Annual Meeting, Portland, OR, July 29-August 1, 2007

I. Introduction

Nonpoint source (NPS) pollution from agriculture remains a major source of water degradation in the U.S. despite the devotion of substantial resources to its control over the past two decades. By definition, NPS pollution comes from many sources whose contributions to water pollution are hard to measure. Numerous economic studies have investigated the efficiency of different policy instruments in this context; Shortle and Horan (2001) provided an excellent survey of the literature.

However, some attributes of NPS pollution that can be critical in the design of policies have received little attention. The first is that there are multiple NPS pollutants which may interact with each other, including nitrogen, phosphorous, sediment, pathogens, and pesticides. Conservation practices that are effective at controlling one pollutant are not necessarily equally good at controlling other pollutants. A second understudied feature of NPS pollution control is the multiple spatial scales at which water quality standards may be set. For example, there may be distinct phosphorus standards at upstream and downstream subwatersheds based on the conditions of each subwatershed. Control measures taken to meet the upstream standard may be adequate to meet or exceed the downstream standards, or actions to meet those standards might actually exacerbate the downstream problems. Another characteristic of NPS pollution control is that multiple conservation practices can be implemented simultaneously in the same field and different fields within a watershed can have distinct practices. Some practices achieve more pollution reduction than others on a given field. Moreover, the effectiveness of a given conservation practice on a given field depends on the conservation practices and cropping systems in place elsewhere in the watershed. In other words, off-site impacts of land use on any given field in a watershed are endogenous to land use choices on other fields of the watershed.

In this study, we examine the policy implications of these three attributes of NPS pollution using a spatially explicit model of a large and critically important agricultural region: the Upper Mississippi River Basin in the central U.S. Specifically, we study (1) the tradeoffs between the costs of pollution control and the level of water quality; (2), the tradeoffs between meeting the water quality targets of different pollutants; and (3) the tradeoffs between meeting water quality targets at different spatial scales. While the fact that higher control costs are necessary to achieve larger water quality improvements is intuitive, the nature of the second and third tradeoffs is less obvious and will depend on the nature of the pollutants and the physical conditions of a watershed. In this paper, we will quantify these tradeoffs and explore the subsequent policy implications.

To empirically estimate these tradeoffs, we develop a modeling framework that (a) realistically incorporates the key attributes of NPS pollution and (b) is able to approximate the efficient solutions by optimally choosing the set of conservation practices for each field. Neither (a) nor (b) is an easy task, as manifested by the dearth of economic studies that reflect both features. Instead, economists have in general utilized simplified representations of the biophysical process of water pollution so that optimization could be performed with conventional approaches. For example, early studies used a simplified model with fixed, exogenous pollution delivery coefficients (e.g., Montgomery, 1972; Ribaudo, 1986 and 1989). Given such assumptions, it is straightforward to solve for cost-efficient allocations of pollution abatement using calculus-based constrained optimization techniques.

Development in the past two decades of realistic, physically-based, spatially distributed hydrologic simulation models highlighted the fact that field-level off-site impacts are endogenous and led several economists to incorporate some features of these models into their analyses via one of two approaches: full spatial optimization using a simplified version of the hydrologic process or incorporation of the hydrologic process but comparing the efficiency of a few select scenarios without explicit optimization.

One example of the first type is Braden et al. (1989), who separated a watershed into hydrologically independent flow paths and use a hydrologic model to estimate the impact of various management alternatives for the flow paths on the resulting sediment yield. As a result, a problem of finding cost-efficient sediment reduction solutions becomes a variant of the knapsack model in operations research. A study by Khanna et al. (2003) provides another good example of the ingenuity demonstrated by researchers to cope with the problem's complexity. The authors capture the interdependencies between upslope and downslope parcels by using coefficients derived from a hydrologic model. They restrict their attention to three parcels up from a stream, and to two alternatives on each parcel: crop production and land retirement.

A drawback to these approaches is that hydrologic models developed for the entire watershed are broken up; hence, one does not get the full benefit of a hydrologic simulation model. By contrast, studies that incorporate the complete hydrologic simulation models typically have not attempted optimization of land use choices. Instead, alternative land use change scenarios that achieve the pollution reduction goals are evaluated (e.g., Secchi, et al. (2005)).

Agricultural engineers have recently examined the cost of NPS control with integrated modeling systems that incorporate the full water quality models into optimization routines that are capable of finding the optimal or near optimal solutions to a problem otherwise intractable with conventional optimization methods. Arabi et al. (2006) and Srivastava et al. (2002) are two outstanding examples. However, these studies are done at a very small scale (smaller than $15km^2$ versus 492,000 km^2 of our study region). In addition, none of these studies examined explicitly the tradeoffs between different NPS pollutants and the tradeoffs between meeting targets at different spatial scales both of which are important policy issues in the control of NPS pollution.

In this paper, we develop a modeling framework that closely integrates an optimization methodology with a biophysical model. The full biophysical model, not simplified proxies, is employed. Second, the modeling framework is built at a regional scale to facilitate the investigation of relevant policy analyses related to the growing "dead zone" in the Gulf of Mexico and the tradeoff between regional and local pollution reduction targets. Third, we derive the conservation production possibility frontier that explicitly incorporates the tradeoffs between pollution control costs and water quality benefits, between different pollutants, or between different control targets. Although the empirical results of our paper may be specific to the region and pollutants considered in our study, the modeling framework and the issues raised in the paper have wide applications in the NPS control of any watershed or area.

The rest of the paper is organized as follows. In the next section, we first set up a conceptual framework for water pollution control in a watershed. Then we introduce an empirical modeling framework that integrated a water quality model and an optimization algorithm. After that, we describe the study region, the pollutants, and the implementation of the empirical algorithm. Results are presented in section 6 with regard to the three tradeoffs we discussed above. The final section provides concluding remarks.

II. Theoretical framework

Suppose there is a watershed with J subwatersheds. In each subwatershed j, there are K_j fields each of which has its own unique land characteristics and land management practices. A set of conservation actions, $\boldsymbol{x}_{jk}$, can be taken for field k of subwatershed j in order to improve the environmental conditions of the watershed. The vector $\boldsymbol{x}_{jk}$ has I elements, which indicate the adoption of a set of I distinct conservation practices. That is, if conservation practice, i, is taken, then $x_{jk}^{\boldsymbol{i}} = 1$; otherwise, $x_{jk}^{\boldsymbol{i}} = 0$ for all $i = 1, 2, ...I$. Note that more than one conservation practice can be adopted on a field so multiple elements of this vector can be non-zero. For ease of reference, we will denote the conservation actions in all fields of the whole watershed as X, i.e., $X = (\boldsymbol{x}_{11}, \boldsymbol{x}_{12}, ...\boldsymbol{x}_{1K_1}; \boldsymbol{x}_{21}, \boldsymbol{x}_{22}, ...\boldsymbol{x}_{2K_2}; ......\boldsymbol{x}_{J1}, \boldsymbol{x}_{J2}, ...\boldsymbol{x}_{JK_J})$. In other words, X is a collection of conservation actions planned for the watershed. The impacts of any $\boldsymbol{x}_{jk}$ are likely to be affected by conservation actions on other fields of the watershed. Thus, for convenience, we will refer to a set of conservation actions planned for the entire watershed, X, as a single plan.

The environmental impact of X is denoted as Y where Y is a vector with N elements, i.e., $Y = (y^1, y^2, ...y^N)$. Each element represents one environmental indicator; for example, y^n can be any pollutant (nitrogen, phosphorus, etc) loading at the watershed outlet; or some index of local water quality indicators. The relationship between Y and X is denoted as

(1) $$y^n = f^n(X; Z),$$

for all $n = 1, 2, ...N$, where Z is set of factors that affect y^n but are not part of the conservation plan such as soil and land characteristics, crop rotations and other crop management practices, climate, etc. Potentially, Z represents a collection of all the land and climate characteristics for each field in each watershed.

One important note on (1) is that an environmental indicator , y^n, can be affected by conservation actions on fields within its own watershed as well as any watershed that drains into the watershed.   One contribution of our paper is that our modeling framework realistically represents such interactions.

Denote T as the conservation possibility set that gives all feasible combinations of conservation plans and environmental outcomes. In other words, T is the set of all (X, Y) combinations that are technically feasible given the existing state of technology. The cost of a conservation plan is represented by a cost function $c(X)$. In general, different conservation plans will result in different costs and different environmental outcomes. Thus, one obvious goal of watershed management is to achieve a desirable tradeoff of costs, $c(X)$, and benefits, Y. In addition, watershed stakeholders may value different environmental indicators differently. For example, a local watershed may have a goal of reducing its phosphorous loading in its lakes but does not feel the need to have its nitrate loading reduced, which may be a concern regionally. Thus, there can be tradeoffs among the different elements of Y.

These tradeoffs can be identified through the following multi-objective optimization problem:

$$\begin{aligned} &\min \quad [c(X),\ y^1, y^2, \ldots\ldots y^N] \\ &s.t., \qquad (X, Y) \in T. \end{aligned} \tag{2}$$

The set of solutions to (2) consists of all conservation plans that are Pareto-optimal. A conservation plan X is Pareto-optimal if there is no $(X', Y') \in T$ such that $f^n(X') \leq f^n(X)$ and $c(X') \leq c(X)$, for all $n \in \{1, 2, \ldots, N\}$, and $m \in \{1, 2, \ldots, N\}$, such that $f^m(X') < f^m(X)$ or $c(X') < c(X)$. In other words, the solutions to (2) together make up the efficiency possibility frontier given T and $c(\bullet)$. Since this frontier is conceptually very close to the standard production possibility frontier (PPF) in production economics, we will simply refer to it as the conservation PPF.

To illustrate the tradeoffs, suppose N=2. We can consider y^1 and y^2 as the nitrogen and phosphorus runoff at the watershed outlet, respectively. Alternatively, these indicators might be the local and regional targets for phosphorus. Then the conservation PPF based on problem (2) has three dimensions: $c(X)$, y^1, y^2. One way to interpret this PPF is that it represents the least cost necessary to achieve the two environmental outcomes y^1 and y^2. In our analysis, we empirically identify this conservation PPF for our study region. To gain a clear picture of the tradeoffs, we can also break down the three dimensional PPF into different two dimensional PPFs that are most often used in economic analysis. For example, we can derive a PPF for $c(X)$ and y^2, which is essentially a cost curve for y^2, while holding y^1 at a prefixed value. In practice, this can mean that the watershed planner intends to identify the tradeoffs between conservation costs and phosphorous loading while insuring the attainment of a nitrogen target, set to meet some ecological needs such as the mitigation of the Hypoxic zone in the Gulf of Mexico. Similarly, we can derive a two dimensional PPF for y^1 and y^2 while holding $c(X)$ at a prefixed value. This is essentially an output transformation curve for a given cost. This curve depicts the combinations of nitrogen and phosphorous loadings that are possible under a given budget.

The shape of the output transformation curve at a given cost, and thus the shape of the cost curve of y^2 for a given y^1, will critically depend on T, the conservation possibility set. In other words, they will depend on the functions $f^1(X)$ and $f^2(X)$. At one unlikely extreme, there may be a one-to-one correspondence between y^1 and y^2, e.g., $y^2 = f^2(X) = \alpha y^1 = f^1(X)$. In this case, the output transformation curve for y^1 and y^2 at a given cost degenerates to a single point, as shown in Figure 1(a). This happens if the impacts of conservation practices are proportional with respect to y^1 and y^2, and such impacts are the same across different fields. At another unlikely extreme, all conservation practices that affect y^1 have no impact on y^2 and those that affect y^2 have no impact on y^1. A single practice that exhibits this property is the reduced application of nitrogen or phosphorous fertilizers. In that case, the output transformation curve for y^1 and y^2 at a given cost appears as in Figure 1(b). In particular, both ends of the curve touch the axes. Essentially, the funding is split between the control of y^1 and y^2. The share of each pollutant control can range from zero to one hundred percent. More generally, most conservation practices have some impact on both pollutants, which implies that a conservation plan will reduce both pollutants even though the degree of reduction may vary across the pollutants. Thus, the transformation curve often looks as illustrated in Figure 1(c), with some distances between both ends of the curve and the axes.

Basically, the different forms of the output transformation curves imply the different flexibility that watershed managers have in setting targets for different environmental outcomes. Figure 1(a) implies there is no flexibility at all—setting a target for y^1 is equivalent to setting a target for y^2. On the other hand, in a Figure 1(b), the targets are not linked at all which allows complete freedom in setting the targets. In the more realistic case, as shown in Figure 1(c), setting a target for one pollutant will have some implications for the other pollutant but there would be some flexibility for watershed planners. In practice, how much flexibility there is will depend on the characteristics of the watershed, the nature of the pollutants under consideration, and the practices that are included in the portfolio of pollution control.

Figure 1. An illustration of the output transformation curve at a given cost

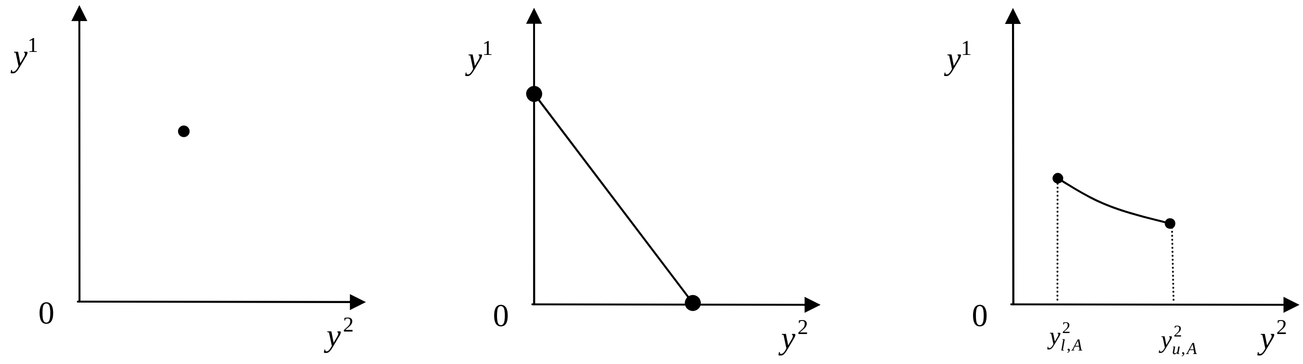

(a) One-to-one correspondence between y^2 and y^1

(b) Practices having separable impacts on y^2 and y^1

(c) The general case

III. An integrated empirical modeling framework

To solve problem **(2)**, we need to know the relationship between water quality outcomes and conservation practices, as represented by equation **(1)**. In economic analyses, biophysical relationships models are often expressed as a functional form which gives the impression that we can apply standard mathematical optimization procedures. However, the relationship between water quality and conservation practices is very complex, and often determined by multiple processes. For example, to model the nitrogen loading at a watershed outlet, the whole life cycle of nitrogen, where mineralization, decomposition, and immobilization are important parts, has to be modeled in the watershed. In our empirical analysis, we use the Soil and Water Assessment Tool (SWAT) model. In SWAT, the nitrogen cycle is simulated using two inorganic forms and three organic forms. Figure 2 describes the process involved.

Figure 2. Nitrogen cycle as simulated in SWAT (adapted from SWAT Theoretical Document).

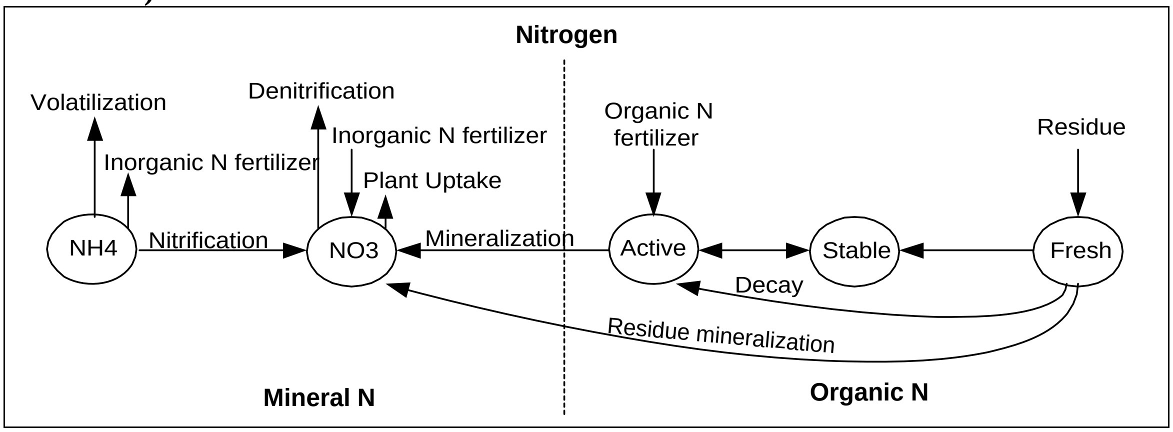

To most accurately represent the biophysical processes in a watershed, it is necessary to employ the full biophysical model rather than a set of more easily characterizable equations that proxy the model.

III.1. A water quality model—the Soil and Water Assessment Tool

The SWAT model (Arnold et al., 1998; Arnold and Forher, 2005) is a conceptual, physically based, long-term, continuous watershed scale simulation model that also operates on a daily time step. In SWAT, a watershed is divided into multiple subwatersheds, which are further subdivided into Hydrologic Response Units (HRU) that consist of homogeneous land use, management, and soil characteristics. Streamflow generation, sediment yield, and non-point-source loadings from each HRU are summed and the resulting loads are routed through channels, ponds, and/or reservoirs to the watershed outlet. Key components of SWAT include hydrology, plant growth, erosion, nutrient transport and transformation, pesticide transport, and management practices. Outputs provided by SWAT include streamflow and in-stream loading or concentration estimates of sediment, organic nitrogen, nitrate, organic phosphorous, soluble phosphorus, and pesticides. Previous applications of SWAT for streamflow and/or pollutant loadings have compared favorably with measured data for a variety of watershed scales (e.g., Arnold and Allen, 1996; Arnold et al., 1999, 2000; Santhi et al., 2001; Borah and Bera, 2004; Jayakrishnan et al., 2005; Gassman et al., 2007). Arnold et al. (2000) performed a stream flow validation study of the UMRB using input data based on the Hydrologic Unit Model for the United States modeling framework. The calibration and validation of SWAT for the UMRB region can be found in Gassman et al. (2006) and Jha et al. (2006).

III.2. Optimization method—the evolutionary algorithm

Using SWAT directly in lieu of the function, $f^n(X;Z)$, in problem (2) poses practical solution challenges. One approach is to run the model for all all possible conservation plans and evaluate the cost and the pollution outcome of each combination. The Pareto frontier would then be the set of least cost combinations associated with each combination of pollution reductions. However, for any realistic watershed problem, this brute force approach is infeasible. Specifically, given that there are I conservation practices possible for adoption on each field and there are $\sum_{j=1}^{J} K_j$ fields, this implies a total of $I^{(\sum_{j=1}^{J} K_j)}$ possible conservation plans to compare. In a watershed with hundreds of fields and several conservation practices, this comparison quickly becomes unwieldy. The combinatorial nature of the problem was recognized by Braden et al. (1989), and was one of the reasons for Khanna et al.'s (2003) decision to focus on a narrow band of land around streams.

Evolutionary algorithms provide a systematic way for searching through large search spaces. These algorithms mimic the process of biological evolution, which, in the words of Mitchell (1996), "in effect, is a method of searching for solutions among an enormous amount of possibilities." Researchers, beginning with Srivastava et al. (2002) and Veith et al. (2003) have used genetic algorithms (GA) in order to search for single cost-efficient watershed-level pollution reduction solutions. However, as discussed in the introduction, these papers focused on very small study areas and do not examine the important issues of NPS control considered in our study.

III.3. The language and logic of evolutionary algorithms

Beginning in 1950's and 1960's computer scientists came to a realization that the theory of biological evolution can be used as an optimization tool for engineering problems. Since the field of evolutionary computation owes its origins to observations of biological evolution, the terminology used has its analogs in biology, although, typically, the entities used to describe an optimization problem are much simpler than the real biological entities bearing the same name. A *genome* (or a *chromosome*) refers to a complete collection of *genes* and fully describes an *individual* (a candidate solution in an optimization problem). A set of possible values that any gene can take is referred to as an *allele set*, or *alphabet.* Often, a genome representing a candidate solution is a one-dimensional array, or vector. A gene then is an element of this array and encodes a particular element of a candidate solution. A value of a gene comes from its allele set, also a vector. Analogous to haploid organisms in real biology, *offspring* is created from two parent individuals. During sexual reproduction, *recombination* (*crossover*) occurs: the offspring's genome consists of portions of each of the two parents' genomes. As in biological evolution, offspring are subject to *mutation:* a random substitution of a gene's value with a value from its allele set.

In this study, the following correspondence between the terminology of evolutionary algorithms and entities related to nonpoint source pollution is made. Table 1 provides the necessary terms:

Table 1. Terminology of evolutionary algorithms in relation to watershed optimization.

Evolutionary computation term	**Its interpretation in a nonpoint source pollution setting**
Allele set	A set of mutually exclusive land use options and conservation practices
Individual (genome)	A distinct allocation of conservation practices and land use options in the watershed
Gene	Spatial unit of analysis (HRU)

In this application of evolutionary algorithms to spatial optimization, a genome is a vector of length *F*, where *F* is the number of spatial decision-making units. Each element of the vector (gene) is encoded with a value from the allele set *A*, and denotes a particular land use option.

As in biological evolution, individuals at every *generation* form *populations*, and are characterized by their *fitness*—a score which measures how well each individual is solving the optimization problem at hand (for example, a value of an objective function). Individuals possessing higher fitness scores are more likely to be selected for reproduction and therefore are more likely to pass along the characteristics associated with the candidate solutions they represent.

While there are many variations of evolutionary algorithms, most that can be called “genetic algorithms” have the following elements in common: populations of individual solutions, selection for reproduction according to fitness levels, crossover to produce new solutions (offspring), and random mutation of new offspring.

Given that in order to characterize the tradeoffs outlined above, a multiobjective optimization problem needs to be addressed, we turn to a class of evolutionary algorithms designed to solve multiobjective problems. Recent years have seen emergence of several multiobjective evolutionary algorithms. We use an algorithm called Strength Pareto Evolutionary Algorithm 2 (SPEA2), developed by Zitzler and Thiele (2002).

The search process starts with a population of candidate solutions from which a new population is created by the process of selection, crossover, and mutation. The fitness score of each individual in the population is a function of how many other individuals in the population it dominates (in the sense of Pareto) and by how many individuals it is dominated by. The algorithm provides an approximate solution to problem (2) by preserving Pareto-nondominated individuals, by eliminating Pareto-dominated solutions, and by iteratively creating new candidate solutions and assessing how well they perform on the multiple objectives outlined in (2). Furthermore, the algorithm takes into account the degree of "crowding" around an individual in order to preserve the diversity in the population and to explore a greater region of the objective space. Details of the fitness assignment in the algorithm are presented in the Appendix.

III.4. Integrating the optimization algorithm with the water quality model

In our application, three major components were integrated to arrive at the final modeling framework. The first component is the logic and the fitness assignment method of a multiobjective evolutionary optimization algorithm, SPEA2. The second component is a publicly available C++ library of genetic algorithms, GALib, originally developed by Wall (1996), with the current version available online. The third component is the hydrologic model, the SWAT model (2005 version), coupled with a Windows-based database control system, i_SWAT (CARD, 2007; Gassman et al., 2003). SPEA2 provides the fundamental multiobjective optimization logic, while GALib provides the tools that are needed to implement an evolutionary search algorithm. Finally, SWAT and i_SWAT provide a way to model the different conservation practices considered in this paper and model their watershed-level environmental impacts.

IV. The study region and the pollutants

The Upper Mississippi River Basin extends from the source of the Mississippi river at Lake Itasca in Minnesota to a point just north of Cairo, Illinois. The total drainage area is nearly 492,000 km^2, which lies primarily in parts of Minnesota, Wisconsin, Iowa, Illinois, and Missouri. Figure 3 contains a map of the Upper Mississippi River Basin and its position in the central U.S. Cropland and pasture are the dominant land uses in the UMRB, which together are estimated to account for nearly 67% of the total area (NAS 2000). Nutrient inputs (nitrogen and phosphorus) to fertilize the land are the primary sources of nonpoint source pollution in the UMRB stream system. These nutrients are also apparently the cause of a major oxygen-depleted hypoxic or "dead" zone in the Gulf of Mexico which has exceeded 20,000 km^2 (Rabalais et al., 2002).While the task force charged with assessing the causes of Gulf hypoxia in 2000 identified nitrate contributions, particularly from the UMRB, as the primary source of the nutrient loading causing the problem, more recent evidence suggests that both nitrate and phosphorous loads from the region (and elsewhere) are to blame.

Figure 3. UMRB and the watershed outlet at Grafton (from Kling et al. (2006)).

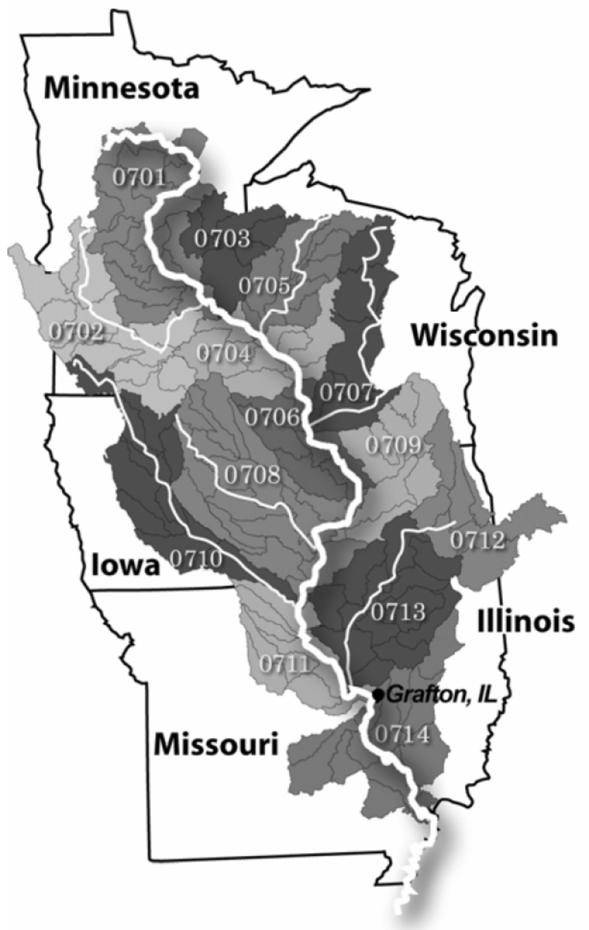

While nitrogen and phosphorous loads are believed to be the primary limiting nutrient in the dead zone in the Gulf, they are also the culprits of substantial local water quality problems within many areas of the UMRB. While phosphorous is more often a target in Total Maximum Daily Load programs in the UMRB, there are also many water bodies listed as impaired due to high nitrogen concentrations.

In short, the water quality problems in the UMRB are substantial and multi-faceted. On the one hand, nutrients from the region negatively affect water quality in lakes and streams locally throughout the basin, negatively affecting recreation opportunities, wildlife viewing, and ecosystem functioning. On the other hand, these nutrients travel out of the watershed and flow in to the Gulf of Mexico where they directly contribute to the large region devoid completely of life. No single regulatory authority has identified a standard or set of water quality standards for the many impacted lakes and streams in the region, but numerous Total Maximum Daily Load regulations, nutrient "criteria," and "targets" for nutrient reduction are in place or being developed by various state and federal agencies. Thus, the model described here representing a multitude of water quality targets at different spatial scales accurately describes the policy environment.

V. Algorithm implementation and the allele set

There are a number of abatement activities that individual farmers can undertake to reduce nitrogen and phosphorous transport from their fields. Various "in-field" conservation practices include conservation tillage (where residue from the previous year's crop is left on the ground to help reduce erosion), buffer strips, grassed waterways, as well as complete retirement of land from crop production in favor of other uses.[1] In addition, nitrogen and phosphorous loadings can be directly controlled by reducing the amount of application of nitrogen and phosphorous fertilizer to the crop. In this study, we consider several in-field conservation practices, a reduction in the quantity of nitrogen fertilizer applied, and retirement of land from crop production. With the exception of land retirement, all other practices are modeled in conjunction with the cropping system currently in place.[2]

[1] The Conservation Reserve Program (CRP) is a very large, federally funded program that makes direct payments to farmers to remove their land from active production and instead plant trees or other perennial ground cover.
[2] Since some conservation practices are currently in place, we assume that the existing conservation practices will either remain under new scenarios or be replaced by more costly and effective practices.

The following table presents the (unconstrained) allele set used in this study. As discussed above, for the HRUs which were observed to have the relevant conservation practice in the baseline, the allele set was constrained.

Table 2. Conservation options (allele set).

Option number	**Option description**
1	Conventional Till (CT)
2	Ridge Till (RT)
3	Mulch Till (MT)
4	No Till (NT)
5	CT+Contour
6	RT+Contour
7	MT+Contour
8	NT+Contour
9	CT+Grassed Waterway
10	RT+Grassed Waterway
11	MT+Grassed Waterway
12	NT+Grassed Waterway
13	CT+Terraced
14	RT+ Terraced
15	MT+Terraced
16	NT+Terraced
17	CT+RF
18	RT+RF
19	MT+RF
20	NT+RF
21	CT+Contour+RF
22	RT+Contour+RF
23	MT+Contour+RF
24	NT+Contour+RF
25	CT+Grassed Waterway+RF
26	RT+Grassed Waterway+RF
27	MT+Grassed Waterway+RF
28	NT+Grassed Waterway+RF
29	CT+Terraced+RF
30	RT+Terraced+RF
31	MT+Terraced+RF
32	NT+Terraced+RF
33	Land retirement

Reduced fertilizer (RF) in the table above refers to a 20 percent reduction in nitrogen fertilizer application. The allele set is constructed to take into account the fact that many of the practices we consider are not mutually exclusive and can be implemented jointly on any given field.

The practices considered are simulated using the SWAT model. In particular, land retirement is modeled by assigning a permanent grass cover to the HRU, fertilizer reductions are modeled by reducing nitrogen fertilizer applications (USDA-ERS) by 20 percent for all crop rotations where nitrogen fertilizer is used, and the in-field practices (tillage, grassed waterways, contour farming, and terraces) are modeled by adjusting the SWAT model parameters in a manner suggested by Arabi et al. (2007).

To meaningfully capture a tradeoff between water quality objectives and total control costs, detailed information on the costs of all the options in the allele set is needed and came from multiple sources. State-level costs of terraces, no till, and contouring were gathered from the Natural Resource Conservation Service website[3]. The costs of grassed waterways were obtained from the CRP program office, and converted to a per acre protected, annualized basis using a 5 percent discount rate and a 10 year useful life term.

The costs of land retirement are proxied by the cash rental rate and the costs of nitrogen fertilizer reductions were developed using the yield curves inferred from Iowa State University Extension's N-Rate Calculator information for geographic zones and corn-soybean crop sequences for Iowa, Minnesota, Illinois, and Wisconsin. State-level data for fertilizer application allowed us to compute the implied reduction in corn yields. Predicted yield reduction, multiplied by the price of corn, served as an approximation to the cost of reducing nitrogen fertilizer application. Details on the computed cost of nitrogen fertilizer reduction are provided in the Appendix.

[3] The cost of establishing a terrace had to be converted to an annualized, per acre, basis. To that end, a cost per foot reported by NRCS was multiplied by 166.7, as this many feet of a terrace can protect one acre of land (E. Palas, IDALS, personal communication). The resulting cost was annualized using a 5 percent rate of discount and a 25 year term representing the useful life of a terrace.

The algorithm was initialized with a population (set of conservation practices on each field in the watershed) of 40 individuals. In order to efficiently exploit our prior domain-specific knowledge, the initial population was seeded with an individual representing the baseline allocation of conservation practices, and an individual representing a scenario of all cropland in the UMRB being retired from production and placed under permanent grass cover. These individuals represent the boundary points on the conservation PPF: the baseline individual results in the lowest cost, and highest nutrient loadings, while the ‘all cropland retired’ individual results in the highest cost and lowest nutrient loadings. We further assist the algorithm in exploring the search space by seeding it with two more individuals, anticipated to be quite distinct in the objective space: an individual assigning all corn HRUs a 20 percent N fertilizer reduction (allele #17, “CT+RF”), and an individual assigning terracing, no-till, and fertilizer reductions (allele #32, “NT+Terraced+RF”). The rest of the initial population was generated by randomly assigning the cropland HRUs with one of the 33 alleles above (subject to the baseline constraint discussed above). We expect good initial coverage of the objective space, thus assisting the evolutionary algorithm in exploring a wider range of the search space.

VI. Empirical analyses and results

We apply the evolutionary algorithm to develop a conservation PPF which (approximately) solves the multiobjective optimization problem (2), using two distinct sets of objectives. The first set of results develops a PPF that relate to the regional water pollution problem of hypoxia in the Gulf of Mexico. Specifically, the three objectives to be minimized are: 1) the cost of nonpoint source pollution control; 2) the mean annual nitrate loadings at the overall UMRB watershed outlet (Grafton, Illinois), and 3) the mean annual total phosphorus loadings at the UMRB outlet.

The second set of results is developed for objectives that relate both to local water quality and Gulf hypoxia.. While the cost and nitrate loading objectives remain unchanged, the phosphorus objective assumes a different form. In particular, we wish to explore the set of tradeoffs between nitrate loadings and control costs, while constraining mean annual phosphorus loadings to be reduced by at least 30 percent in *each* of the subwatersheds in the UMRB (as represented by the 8-digit HUC watersheds).

The resulting frontiers for the two sets of objectives allows us to provide empirical answers to important policy questions . In particular, what is the nature of a tradeoff between the NPS control costs and NPS reductions? What are the costs of reducing nutrient loadings at the outlet for each of the nutrients separately, and jointly? Given a particular cost of control, what are the tradeoffs between nutrients? What practices should be used to control nitrates separately, phosphorus separately, and nitrates and phosphorus jointly? How do the answers change with a change in a spatial scale of nutrient reduction targets (i.e., when a subwatershed-level targets for phosphorus are employed)?

VI.1. Tradeoffs of NPS control costs and water quality benefits

The solution to a multiobjective optimization problem (2) is a three-dimensional conservation PPF. A set of Pareto-nondominated points surviving after 500 generations (iterations of the evolutionary algorithm) provides an approximation to the true frontier. Figures 4 and 5 provide two-dimensional projections and a three-dimensional visualization of the empirical frontier.

Figure 4. 2-dimensional projections of the empirical conservation PPF.

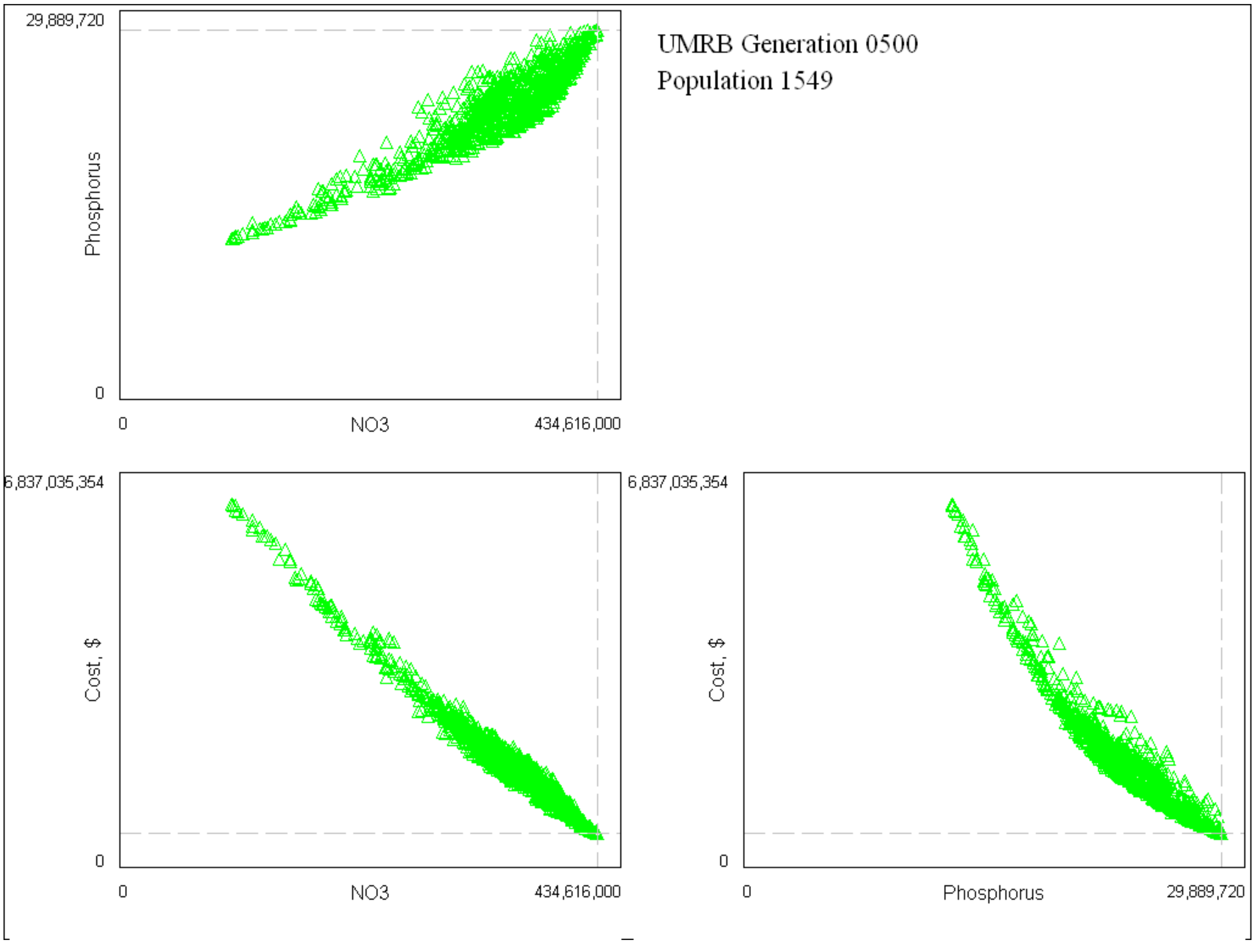

Figure 5. 3-dimensional vizualization of the empirical conservation PPF.

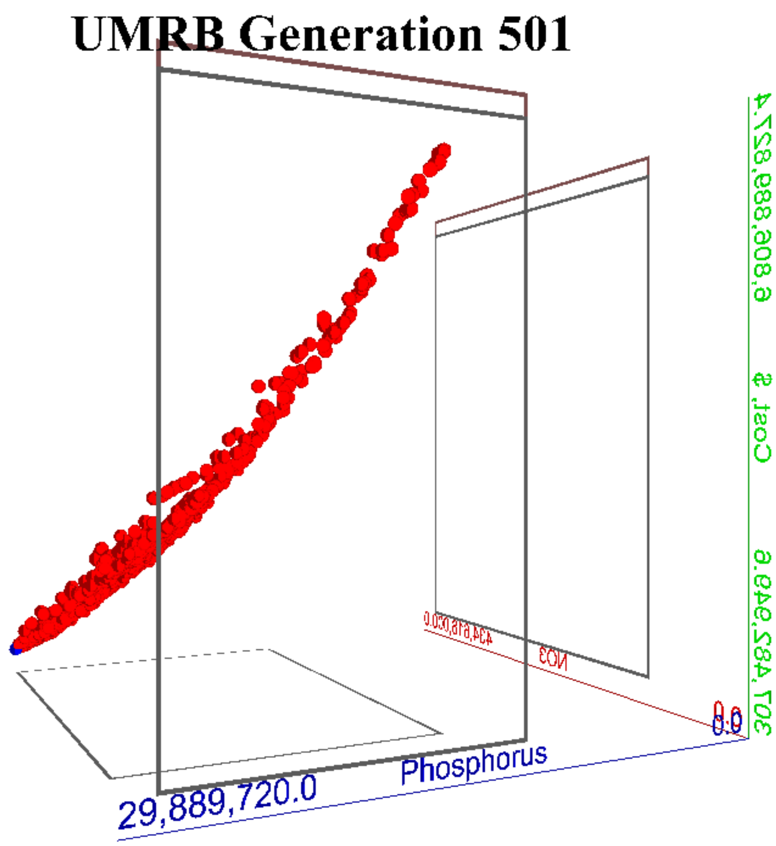

Figure 6 provides a mathematical approximation to the empirical results by fitting a second-degree polynomial through the points in cost-nitrate-phosphorus space.

Figure 6. Fitted conservation PPF.

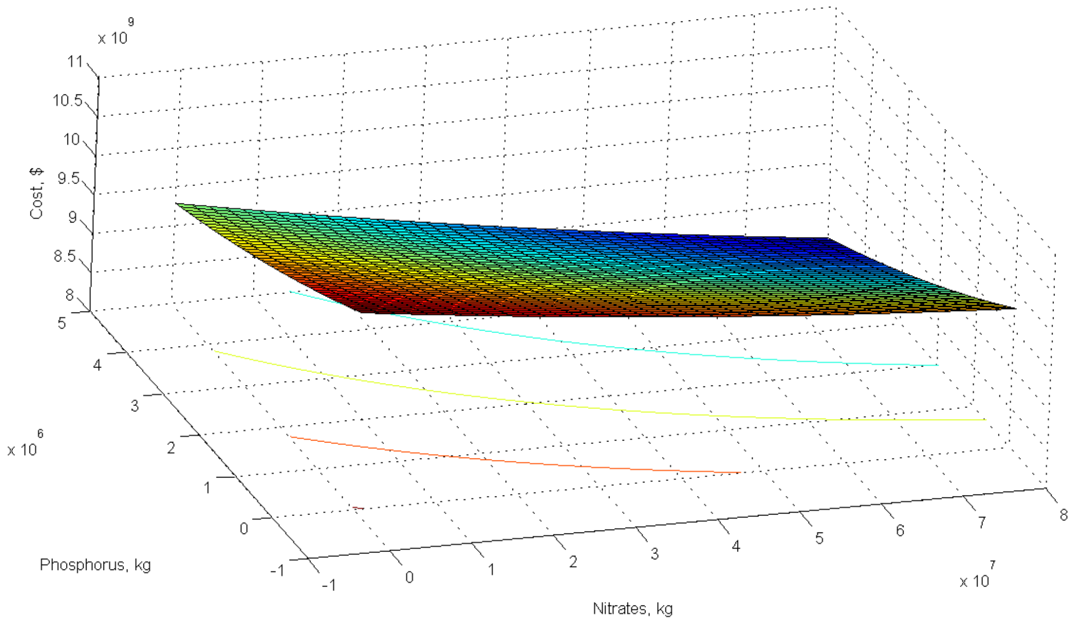

Both the empirical and the fitted frontiers indicate that tradeoffs between both nitrate and cost and phosphorus and cost can be represented by convex total cost curves[4].

Figure 7 contains cost curves for nitrate reductions under two different scenarios. Under one scenario, the cost curve is developed from the PPF in the absence of any constraint on phosphorus levels (as a lower envelope of the PPF in nitrate-cost space). Under an alternative scenario, a 30 percent concomitant reduction in phosphorus loadings is imposed as a constraint. As theory suggests, the constrained cost curve can be no lower than the unconstrained one, and that is indeed the case.

[4] The fitted polynomial relating cost to nitrate and phosphorus loadings at the watershed outlet was estimated to be: $Cost=1.05E+10-13.1814*N-316.718*P-1.2E-06*N*P+4.80E-08*N^2+1.35E-05*P^2$. $R^2=0.999$.

Figure 7. Cost-pollution tradeoff for Nitrate loadings at the outlet.

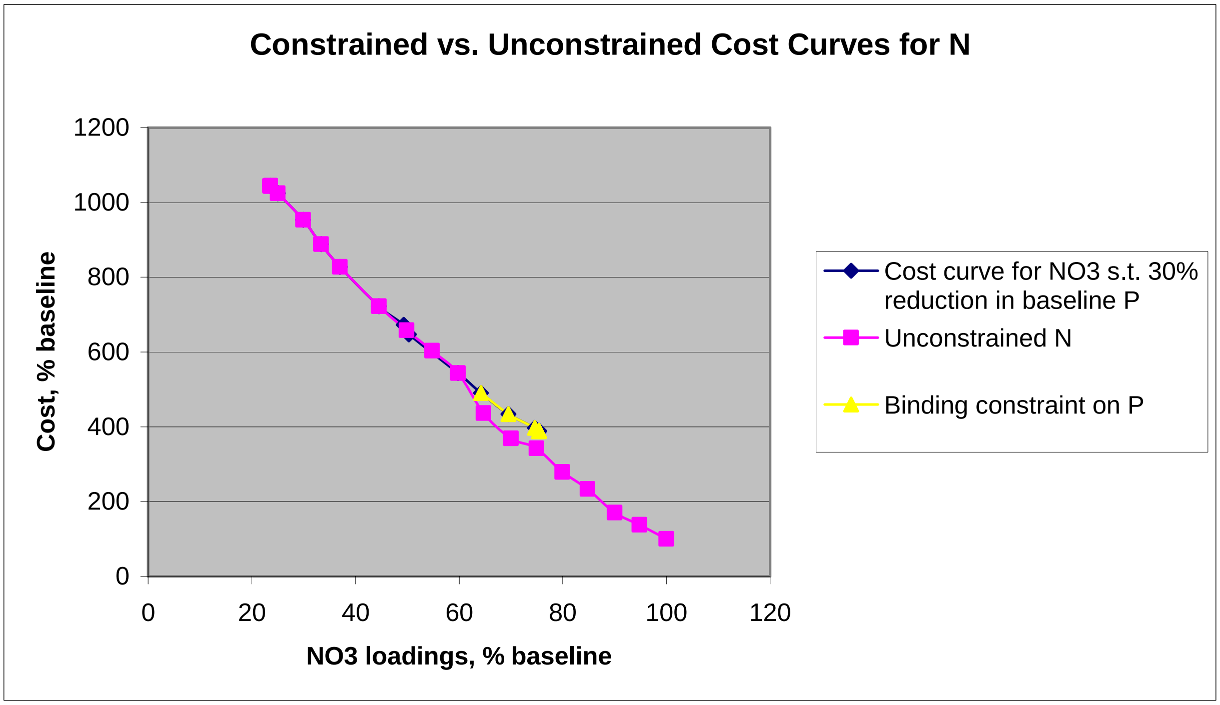

Figure 7 provides interesting insight on the interactions between conservation practices considered and the two nutrients. Note that while the unconstrained cost curve starts out at the baseline level of nitrate loadings, imposing a phosphorus constraint forces the curve to start at a level of nitrate loadings which is almost 25 percent lower than the baseline. In other words, given the set of practices considered, once phosphorus loadings are reduced by 30 percent, an automatic reduction of about 25 percent in nitrate loadings follows. Further evidence of such interactions is revealed by the fact that the phosphorus constraint is only binding up to about a 35 percent reduction in nitrates. Greater reductions in nitrates lead to simultaneous reductions in phosphorus, suggesting a complementarity in the set of practices used to achieve greater nitrate reductions.

Similar observations can be made for Figure 8, which depicts an unconstrained phosphorus cost curve and a constrained phosphorus cost curve, subject to the 30 percent constraint on nitrate loadings. In this case, imposing a nitrate constraint automatically reduces phosphorus loadings by about 17 percent, and a nitrate constraint is binding up to a 30 percent reduction in phosphorus, and is not binding thereafter.

Figure 8. Cost-pollution tradeoff for Phosphorus loadings at the outlet.

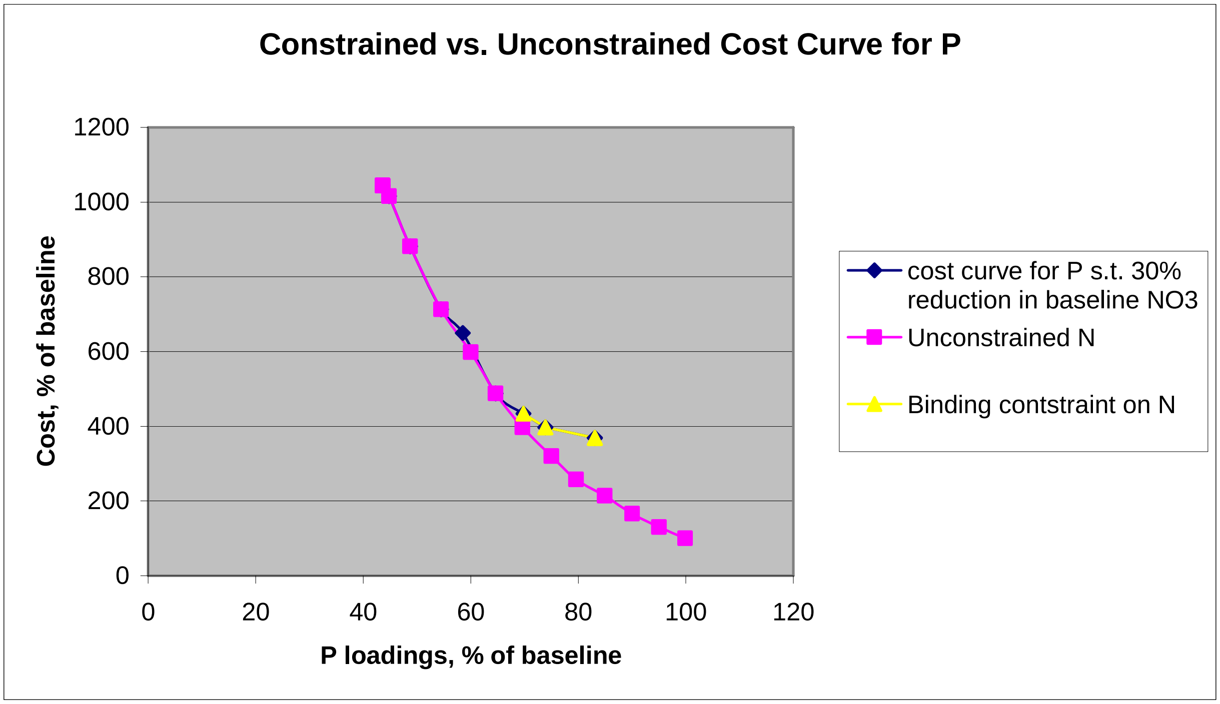

These findings can be explained as follows. First, the truncation imposed on a cost curve for nitrates (phosphorus) by a constraint on phosphorus (nitrates) is most likely due to the complementarities in controlling both nutrients imbedded in the land retirement option. As discussion below will show, significant additional investments in land retirement are needed to achieve 30 percent reductions in either nitrates or phosphorus. This immediately imposes a truncation on the constrained cost curves. Second, when greater reliance on land retirement becomes necessary to further control nitrates (phosphorus), leading to simultaneous reductions in the other nutrient, the constraint on phosphorus (nitrates) becomes non-binding.

VI.2. Tradeoffs between different pollutants (Nitrates and P)

As highlighted in the theoretical discussion above, the tradeoff between different pollutants for a particular level of control costs can range from a curve spanning the entire range of possible nutrient values to just a single point, or anything in between. Thus, theory alone provides fairly limited guidance as to what we should expect from a particular 'isocost' curve.

The task of looking at the tradeoffs between nutrients is further complicated by a fact that, in order to properly explore such a tradeoff curve, all the points on the curve have to have identical control costs. This works well in theory, but, in the empirical application, we have a finite number of points defining the PPF. Thus, potentially, only a few points may be located in a narrow band in a cost dimension to approximate such a tradeoff curve, and it is likely that we cannot identify distinct points along nitrate and phosphorus dimensions carrying identical control costs. These considerations make the empirical analysis of 'isocost' curves somewhat limited.

Nonetheless, given these caveats, a set of tradeoffs depicted in Figure 9, drawn for successively higher levels of control costs, tell an interesting story.

Figure 9. Tradeoffs between Nitrates and Phosphorus control.

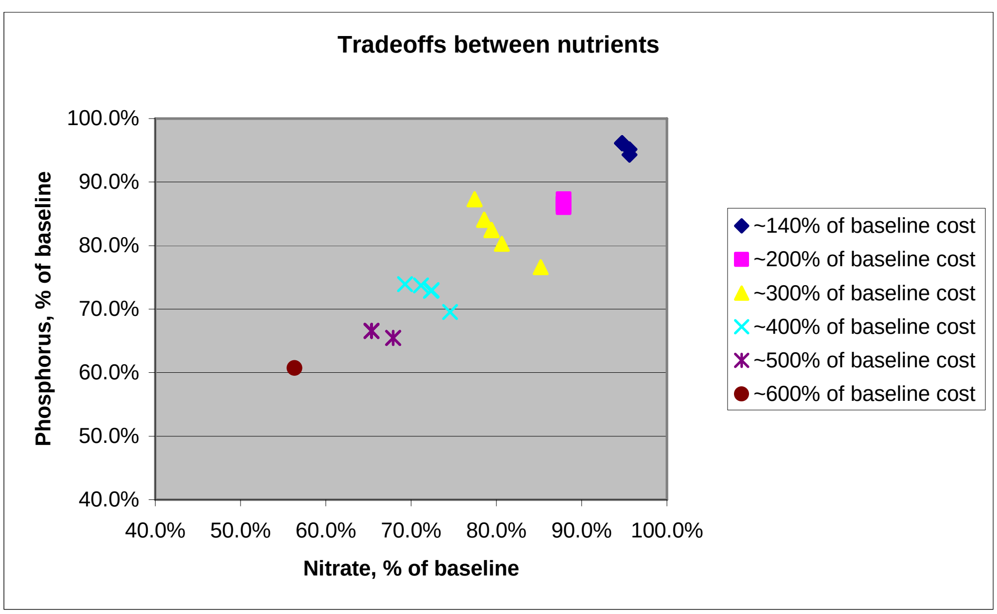

For lower levels of control costs (~140%, 200%), the set of tradeoffs is quite limited in its range. This may suggest that the set of practices that can be used to generate different watershed configurations subject to these budget constraints may not span a wide range in pollution space. However, as the cost levels rise to 300% of baseline cost, a much more pronounced frontier emerges, only to be shrunk again as the cost levels rise. This is consistent with the pattern observed in developing control cost curves for nutrients: as costs rise, use of land retirement becomes more widespread, which leads, eventually, to a collapse of the empirical isocost to a single point.

It should also be pointed out, that, given the diverse set of conservation practices being considered, we expect a set of tradeoffs to exist even for the situation where our algorithm only finds a single point, as the tremendous number of possible reallocations of conservation practices implies that there is also a large number of possible allocations of conservation practices even for a single level of cost. A full development of a (restricted) frontier in nitrate-phosphorus space could be undertaken using the methods employed in this study, and could serve as an interesting extension of this research. The current set of results can only be used to demonstrate that such tradeoffs indeed exist, and that their extent and shape varies with the cost level and the practices used.

VI.3. Effects of targeting nutrients separately or jointly.

The empirical frontier presented above consists of a fairly large number of individuals, each representing a distinct way of placing conservation practices in the watershed. While the frontier itself summarizes the tradeoffs for a range of control costs and nutrient reductions, each individual on the empirical frontier contains information on the ‘look’ of the watershed, that is, it is essentially a prescription for the application of conservation practices in the watershed. Potentially, a regulator makes use of the tradeoff information embedded in the frontier, and selects a set of appropriate nutrient reduction targets. A particular individual meeting these targets is then selected from the frontier, and it then specifies the subbasin-level distribution of conservation practices in the watershed.

Of course, whether nitrates and phosphorus at the outlet are targeted separately or jointly may have dramatic implications for which set of conservation practices should be used and where they should be located within the watershed. Further, this highlights the importance of developing and implementing a plan meant to reduce both pollutants, should these reductions be needed. If a plan meant to control only nitrates (or phosphorus) is quite different from a plan controlling both pollutants, then implementing water quality policy in a piecemeal fashion (e.g., control nitrates first, then focus on phosphorus) may be socially costly and inefficient.

Our empirical analysis confirms that such considerations are important, at least for a range of nutrient reduction targets. Next we demonstrate implications for the allocation of conservation practices for a 30 percent reduction goal in nitrates and phosphorus, depending on whether each nutrient is targeted separately or jointly.

First, Figure 10 below demonstrates how one identifies distinct individuals on the frontier, depending on the targeting strategy. Suppose a regulator wishes to reduce nitrate loadings by 30 percent. Such an individual (highlighted in blue) is located at the intersection of the lower envelope of the frontier in nitrate-cost space and the line specifying the loadings target (dashed line). This individual lies on the unconstrained cost curve for nitrates identified above. As a result, phosphorus loadings fall short of the 30 percent reduction goal (bottom right panel). Similarly, focusing on phosphorus reductions alone (red point) results in nitrate loadings greater than the 30 percent reduction goal. A point highlighted in purple, however, represents a watershed configuration which achieves both nutrient targets simultaneously. This point, for example, can be found on the constrained cost curves above.

Note that, while the three distinct individuals are located fairly ‘closely’ on the PPF, one cannot be sure that the actual allocation of conservation practices is similar. This can be seen by looking at the distribution of the watershed cropland among the 33 conservation practice options.

Figure 10. Graphical representation of targeting nutrients separately and jointly.

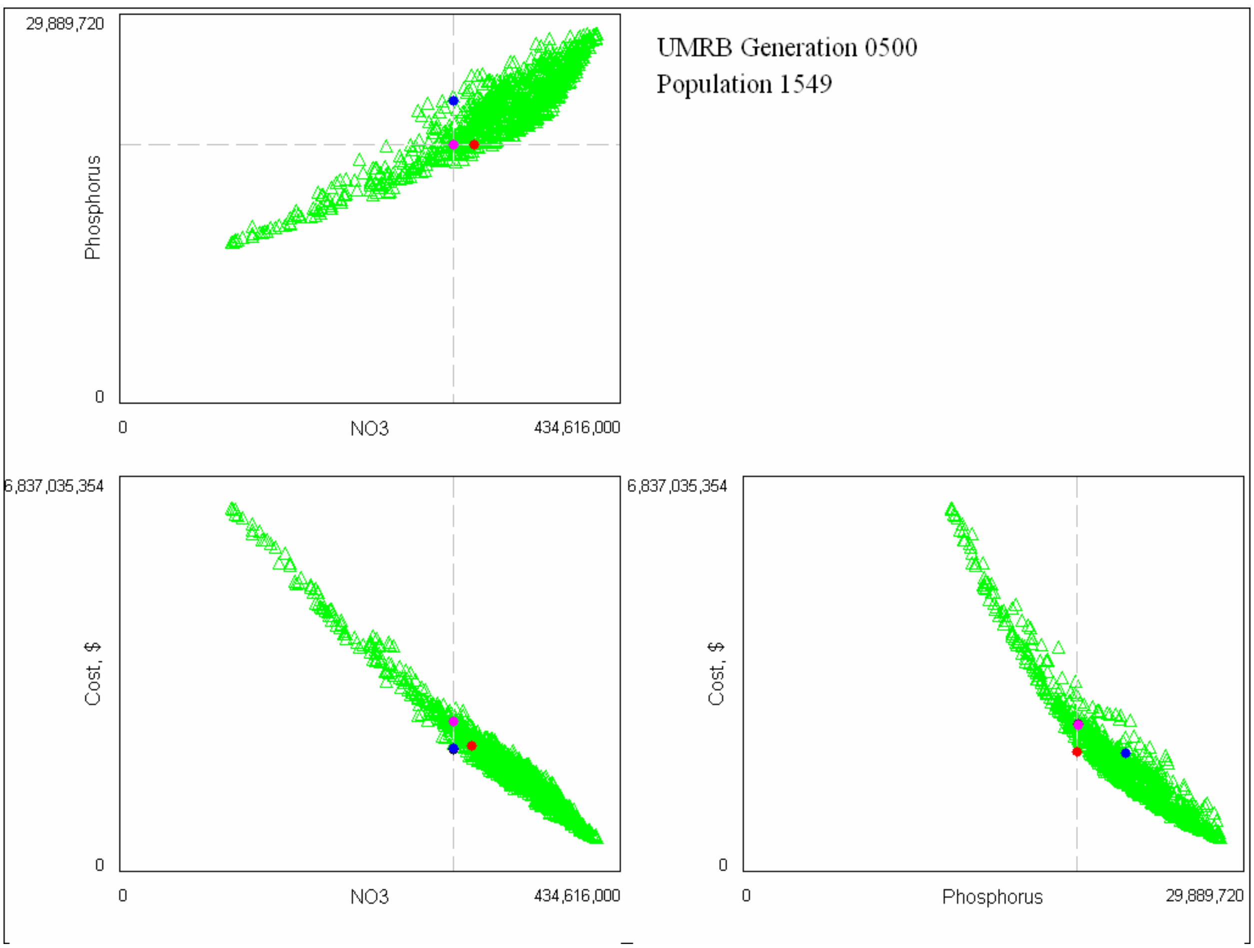

We thus identify 3 distinct individuals in the empirical PPF. Table 3 below identifies the cost and pollution consequences for the 3 scenarios:

Table 3. Consequences of targeting nutrients for a 30 percent reduction.

Ind #	NO_3 loadings	Control cost, $/year	P loadings	% of Baseline NO3	% of Baseline cost	% of Baseline P
2479 (blue)	289,740,000	2,225,360,000	23,688,400	70	369	83
4052 (red)	308,780,000	2,394,970,000	19,824,600	75	397	70
3829 (purple)	287,620,000	2,615,470,000	19,893,200	70	434	70

Each of these individuals prescribes a distinct placement of conservation practices. *A priori*, we expect to see greater use of options containing nitrogen fertilizer reduction for individual 2479. Also, since this individual reduces nitrates without an explicit target for phosphorus, we expect to see very few erosion control practices being selected (given that phosphorus is often bound to soil). That is, we expect to see very few options containing terraces, grassed waterways, and contouring having been selected in individual 2479. On the contrary, for the individual focusing on phosphorus (individual 4052), we expect to see few options containing nitrogen fertilizer reductions, and a greater area of the watershed devoted to practices typically considered helpful in controlling erosion (and thus soil-bound phosphorus): terraces, contouring, grassed waterways.

For individual 3829, we expect that alleles representing combinations of practices which could be beneficial for both nitrates and phosphorus to be selected to a greater extent than in both individual 2479 and 3829.

Land retirement is beneficial for both nitrate and phosphorus loadings, so no a priori ranking in its use between the three individuals is obvious. However, if ‘nutrient-specific’ options are not sufficient to reach a 30 percent reduction, then some use of land retirement is expected. This would also explain some complementarities observed (e.g., individual 2479 is predicted to reduce phosphorus loadings by 17 percent, while targeting nitrates alone).

Table 4. Distribution of cropland between alleles for the 3 individuals.

	Ind # 2479: 30% nitrates reduction		Ind # 4052: 30% P reduction		Ind # 3829: 30% NO3, P reduction	
Option description	**area, km^2**	**percent of total area**	**area, km^2**	**percent of total area**	**area, km^2**	**percent of total area**
Conventional Till (CT)	39485	16%	25411	10%	12500	5%
Ridge Till (RT)	12129	5%	14491	6%	9148	4%
Mulch Till (MT)	24841	10%	26854	11%	24460	10%
No Till (NT)	11351	5%	13039	5%	17626	7%
CT+Contour	1074	0%	1517	1%	1721	1%
RT+Contour	0	0%	1609	1%	3599	1%
MT+Contour	1323	1%	5606	2%	5014	2%
NT+Contour	402	0%	4647	2%	5173	2%
CT+Grassed Waterway	1622	1%	6174	2%	2311	1%
RT+Grassed Waterway	208	0%	2035	1%	2738	1%
MT+Grassed Waterway	896	0%	6439	3%	9858	4%
NT+Grassed Waterway	1483	1%	7744	3%	6980	3%
CT+Terraced	75	0%	2779	1%	1432	1%
RT+ Terraced	0	0%	2607	1%	2273	1%
MT+Terraced	812	0%	7352	3%	10217	4%
NT+Terraced	87	0%	7170	3%	6063	2%
CT+RF	26573	11%	814	0%	2203	1%
RT+RF	18024	7%	4657	2%	8494	3%
MT+RF	38439	15%	4679	2%	6158	2%
NT+RF	13297	5%	4494	2%	5392	2%
CT+Contour+RF	0	0%	512	0%	1121	0%
RT+Contour+RF	52	0%	3915	2%	5192	2%
MT+Contour+RF	866	0%	5628	2%	5900	2%
NT+Contour+RF	174	0%	5965	2%	9054	4%
CT+Grassed Waterway+RF	533	0%	2335	1%	3265	1%
RT+Grassed Waterway+RF	203	0%	2394	1%	6076	2%
MT+Grassed Waterway+RF	1613	1%	6822	3%	8167	3%
NT+Grassed Waterway+RF	709	0%	5780	2%	5590	2%
CT+Terraced+RF	0	0%	1467	1%	1636	1%
RT+Terraced+RF	0	0%	3192	1%	4352	2%
MT+Terraced+RF	2000	1%	5500	2%	8384	3%
NT+Terraced+RF	87	0%	6112	2%	5316	2%
Land retirement	50968	20%	49587	20%	41911	17%

Empirical results mostly confirm our expectations: individual 2479 allocates significant share of cropland to nitrogen fertilizer reduction options, while essentially ignoring such options as contouring, terracing, and grassed waterways. In turn, terracing, contouring, and grassed waterways are utilized to a much greater extent by individual 4052, while at the same time selecting nitrogen fertilizer reductions to a much smaller extent than individual 2479. Also, individual 3829 allocates relatively more land to options consisting of combinations of practices expected to be beneficial for both nitrates and phosphorus (e.g., combinations of terraces with nitrogen fertilizer reductions). All three individuals place a significant portion of the land area to the land retirement option.

Given the fact that land retirement is the costliest and also the most widely selected allele option, it is instructive to look at the subbasin-level maps of prescribed land retirement generated by the three individuals[5].

The following maps represent the distribution of land retirement as prescribed by the three targeting strategies, as well as a map of land retirement costs used in implementing the algorithm.

Several observations can be made about the maps of distribution of land retirement in the watershed. First, the algorithm clearly does not allocate land to be retired from production based on cost considerations alone, but rather on the impact that the spatial placement of land retirement will have on nutrient loadings at the outlet. In fact, all three individuals allocate fairly expensive areas in Illinois to be retired from production, while individuals that reach the 30 percent phosphorus reduction target (4052 and 3829) also retire expensive land in Iowa.

Second, the maps once again highlight the fact that if a regulator eventually may wish to control both nutrients, it would not, in general, be desirable to proceed in a sequential fashion.

---

[5] Similar maps can be produced for all 33 allele options, but that is unlikely to add to much to the discussion.

For example, if a regulator were to target land retirement so as to reduce phosphorus outlet loadings without simultaneously considering nitrates and subsequently decides on a need to reach the nitrogen target as well, the reallocation of land retirement would likely be quite costly.[6]

Figure 11. Subbasin distribution of land retirement.

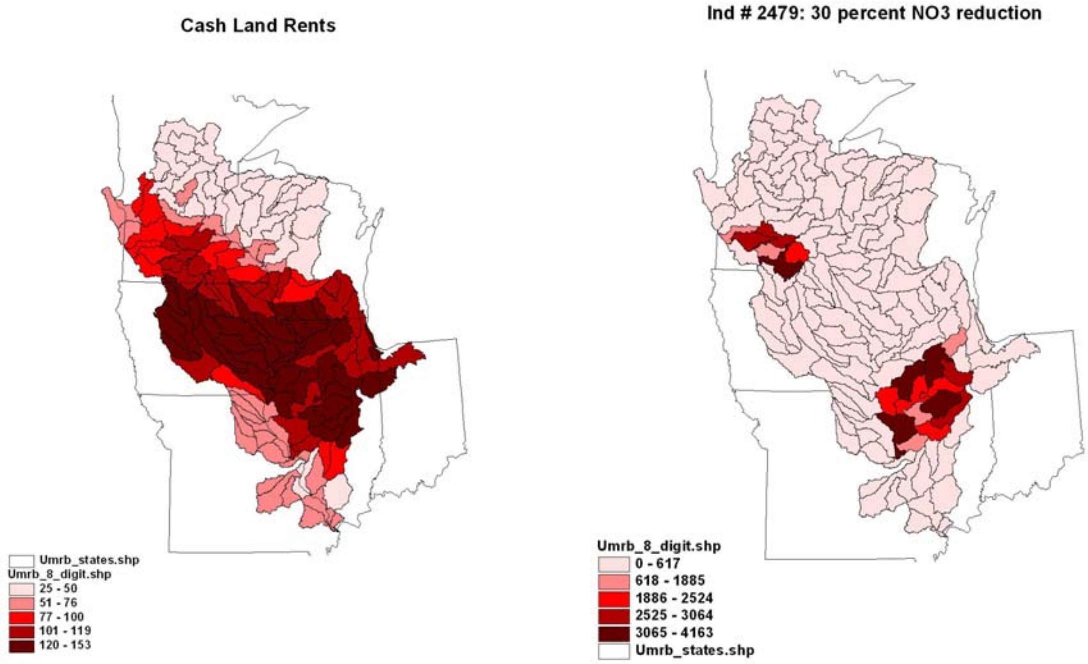

[6] Of course, given a particular investment in land retirement, making the watershed look like the map of Ind #3829 would also be inefficient. The main message here is that thinking about one nutrient as opposed to a combination of nutrients has significant implications for targeting of conservation practices.

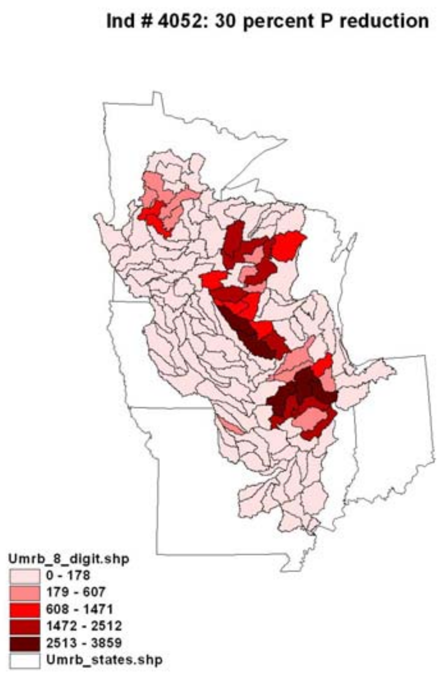

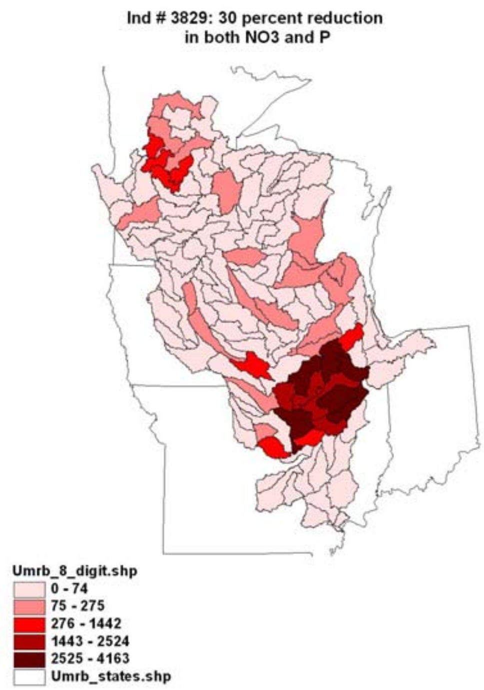

VI.4. Tradeoffs of alternative water quality targets

The analysis above was conducted under an objective of simultaneously reducing nitrate and phosphorus loadings at the outlet of the UMRB. Thus, the evolutionary algorithm only rewards those solutions which reduce nutrient loadings at the outlet subbasin, and does not reward any reductions occurring in other subbasins in the watershed. To illustrate, we again turn to individuals analyzed in the section above, and consider the spatial distribution of loadings of nitrates and phosphorus.

The first set of maps depicts the subbasin-level loadings of nitrates for inidividuals 2479 and 3829, expressed in terms of baseline loadings. As one can clearly see, setting a nutrient reduction goal in terms of reductions at the outlet of the watershed has profound implications for local water quality. When the goal is nitrate loading reductions at the outlet, the maps indicate that the algorithm allocates reductions quite unequally, with some subbasins where reductions are dramatic (over 80 percent), while many of the subbasins are not selected for nitrate reductions at all. This is of course what we would expect the algorithm to do, given the location of the watershed outlet.

Figure 12. Subbasin Nitrate load distribution, 30% Nitrate target vs. 30% P and Nitrates target.

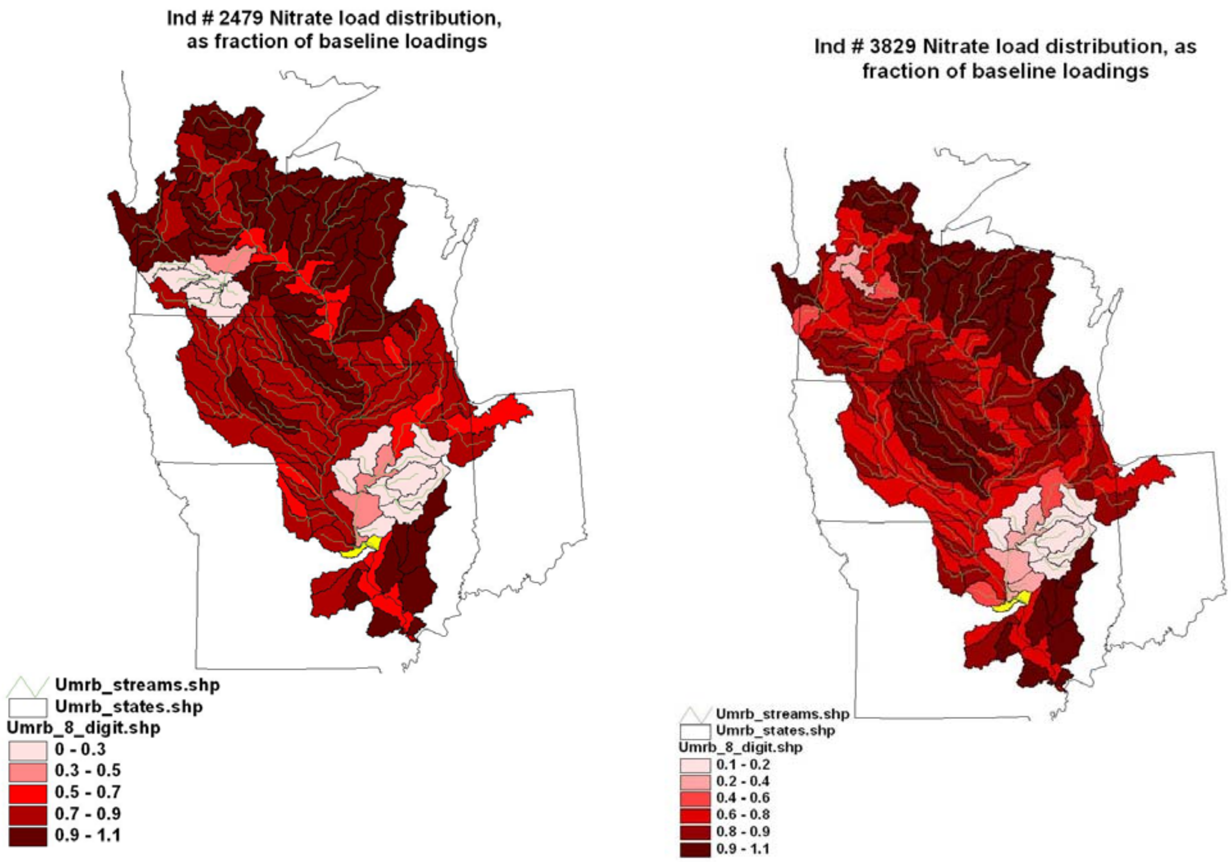

It is also interesting to point out that some subbasins selected for nitrate loading reductions follow the path of the Mississippi River (it is especially evident in Minnesota). Again, this is what we would expect the algorithm to do.

A similar pattern is observed when we look at the subbasin-level reductions in phosphorus loadings: while a phosphorus reduction of 30 percent is achieved, the distribution of loading reductions is such that almost a half of 131 subbasins do not experience a 30 percent reduction for both individuals 4052 and 3829. This may not be satisfactory if phosphorus pollution is a local water quality concern.

Figure 13. Subbasin P load distribution, 30% P target vs. 30% P and Nitrates target.

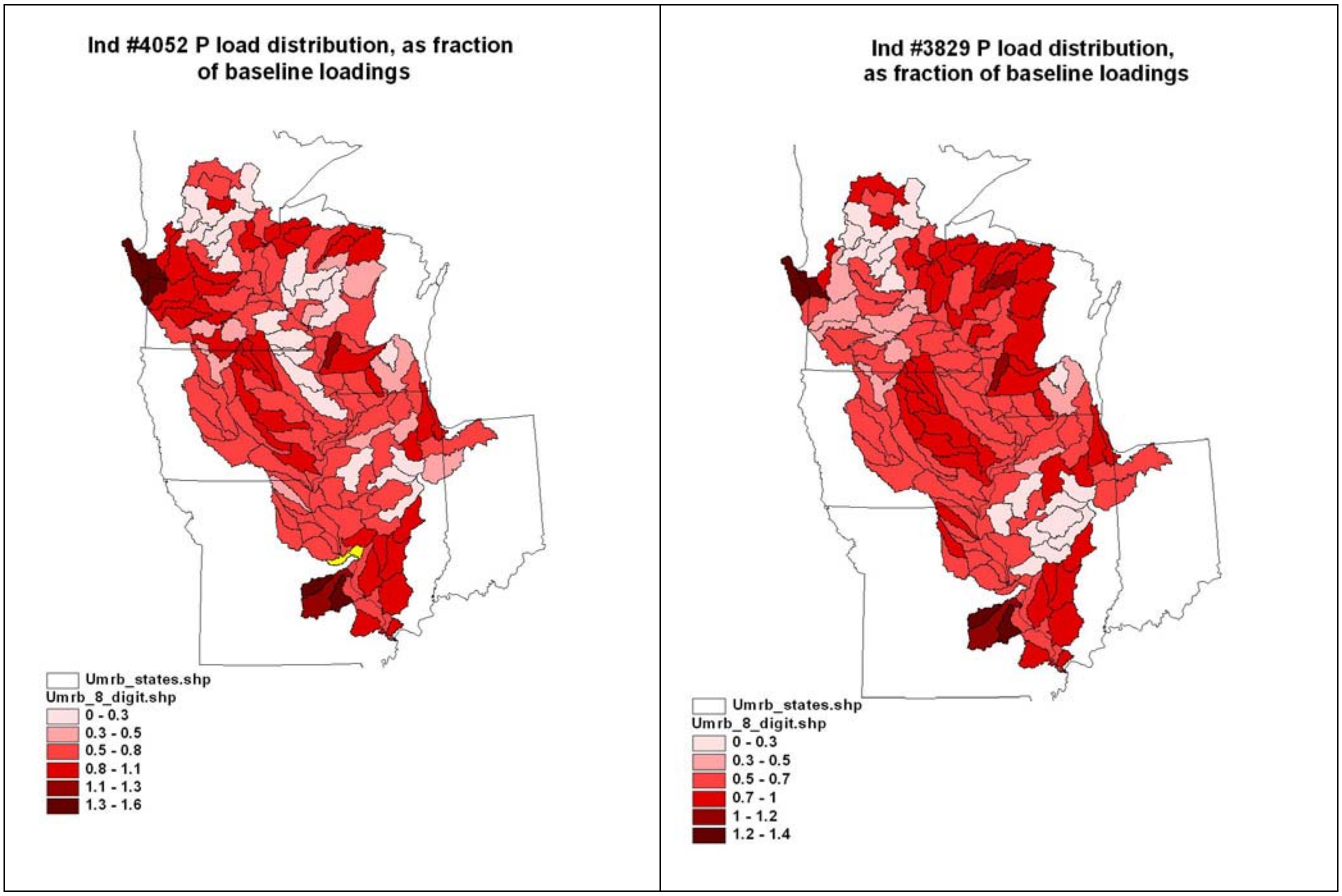

Next, we alter the objective function for the multiobjective evolutionary algorithm to identify a set of solutions which achieve a particular subbasin-level reductions in phosphorus loadings. Policy relevance of exploring such solutions is evidenced by a large number of TMDL plans calling for phosphorus reductions on a local, watershed, scale.

In this application, we impose a 30 percent reduction goal for every cropland 8-digit HUC subbasin in the UMRB. To this end, we consider the following three objectives for multiobjective optimization: nitrate loading at UMRB outlet, control costs, and the sum of excess phosphorus loadings from the subbasins. Formally, problem (2) becomes:

(3) $$\begin{aligned} \min \quad & [c(X),\ y^1, y^2] \\ s.t., \quad & (X,Y) \in T. \end{aligned}$$

where $y_1 = N_{outlet}$, and $y_2 = \sum_{i=1}^{131} \max\left(0, P_i - 0.7P_i^{baseline}\right)$. N_{outlet}, P_i, and $P_i^{baseline}$ represent nitrate loadings at the UMRB outlet, phosphorus loadings at subbasin i, and baseline phosphorus loadings at subbasin i.

The set of solutions where y_2 is zero defines a set of tradeoffs between nitrate reductions and control costs. Empirically, we select individuals for which the value of y_2 is 1 percent of baseline value or less, to represent the new set of tradeoffs. Figure 14 demonstrates the empirical tradeoff curve. Every point on this curve approximately achieves a subbasin-level phosphorus loading reduction of at least 30 percent.

Similar to the results obtained for the constrained cost curves, imposition of a local phosphorus reduction constraint automatically truncates the set of tradeoffs to a point where nitrate loadings at the outlet are reduced by approximately 35 percent.

Figure 14. Nitrates-cost tradeoff under a goal of reducing subbasin P loadings by 30 percent.

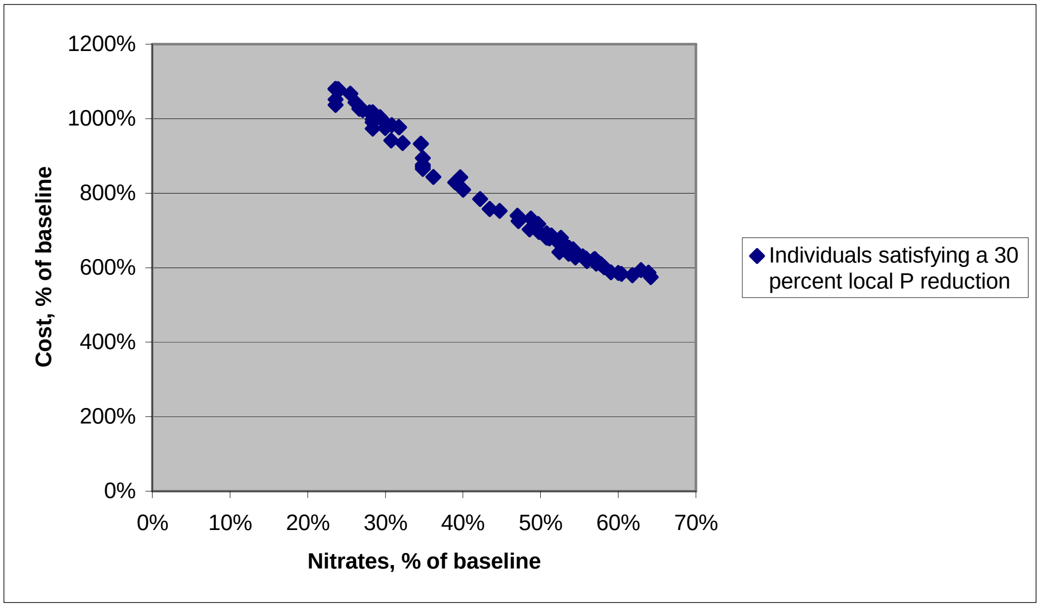

This boundary point is found to be individual number 3331. Incidentally, it achieves exactly a 30 percent reduction in phosphorus loadings at the outlet (much like individual 3829), but the distribution of conservation practices (and nutrient loadings) is dramatically changed.

Figure 15. Subbasin P load distribution, outlet P target vs. local P target.

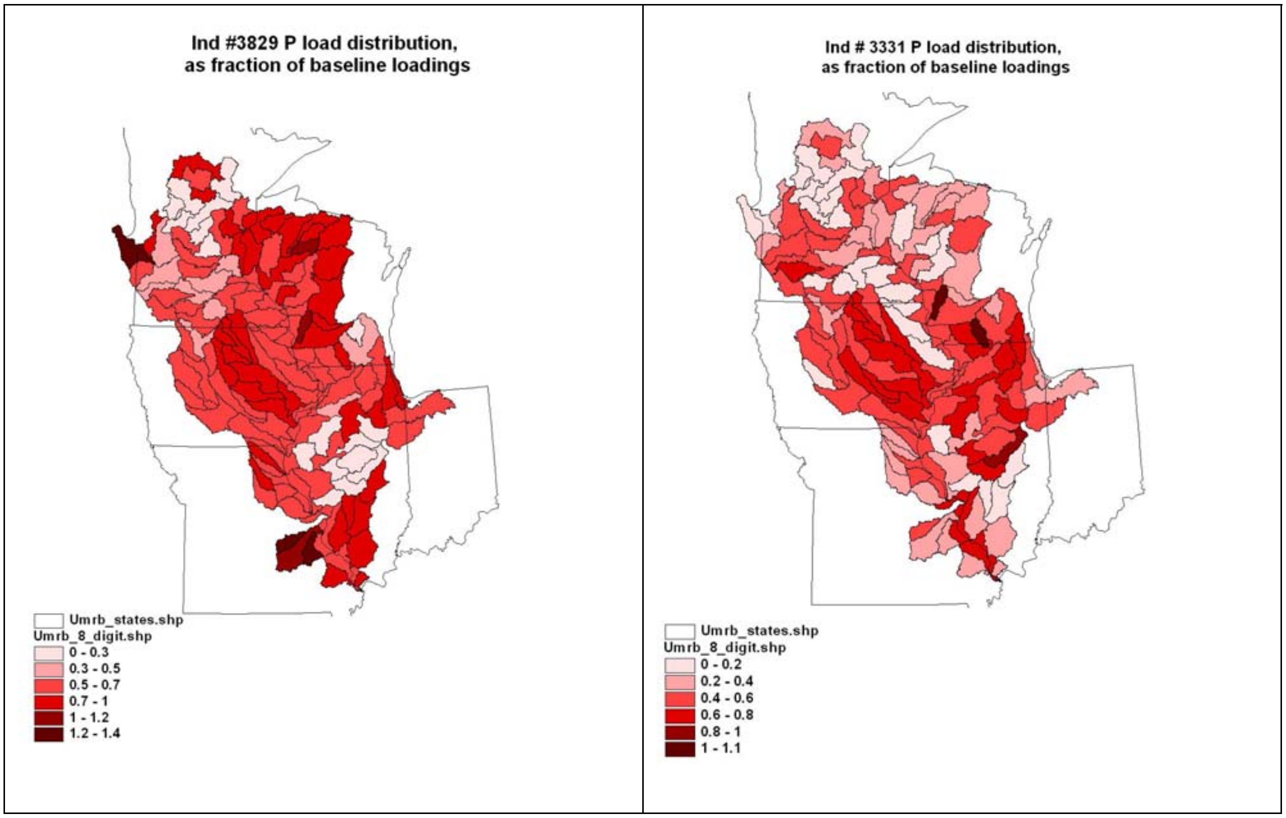

A side-by-side comparision of maps of subbasin-level phosphorus loadings tells the story, but a histogram of the number of subbasins achieving particular nutrient reductions is even more dramatic:

Figure 16. Number of subbasins with given subbasin-level phosphorus loadings, local P target vs. outlet P target.

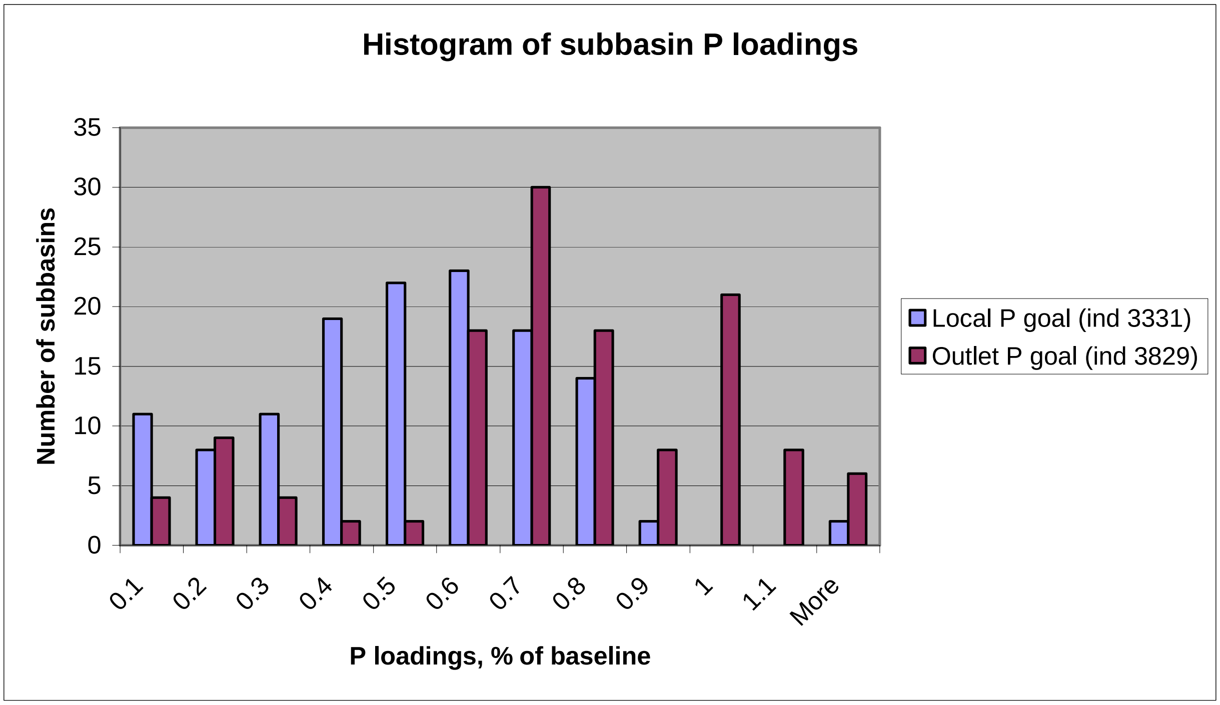

Under the outlet P goal (individual 3829), many subbasins do not experience any decreases in phosphorus loadings, while under a local P goal (individual 3331), an overwhelming majority of subbasins achieves at least a 30 percent reduction. The subbasins where a 30 percent reduction is not achieved contribute only a small fraction of excess loadings relative to the 30 percent goal.

Of course, extra phosphorus control comes at extra cost: the cost of controlling local P at or below the reduction goal is estimated to be almost 3.5 billion dollars per year (or 500 million dollars more than the cost of achieving a 35 percent nitrate reduction, but only achieving a 30 percent phosphorus reduction goal at the outlet).

VII. Concluding Remarks

In this study, we examined the policy implications of efficient control of NPS pollution using a spatially explicit model of a large and critically important agricultural region: the Upper Mississippi River Basin in the central U.S. We derived the conservation production possibility frontier that explicitly incorporates the tradeoffs between pollution control costs and water quality benefits, between different pollutants, or between different control targets. To empirically estimate these tradeoffs, we develop a modeling framework that (a) realistically incorporates the key attributes of NPS pollution and (b) is able to approximate the efficient solutions by optimally choosing the set of conservation practices for each spatial unit in the Basin. The regional scale of our modeling framework facilitates the investigation of relevant policy analyses related to the growing “dead zone” in the Gulf of Mexico and the tradeoff between regional and local pollution reduction targets.

Several caveats should be mentioned. First, the enormity of the search spaces precludes us from characterizing the solutions we obtain as truly efficient. However, we believe that the approximations we find are quite relevant to policy analysis. Second, our results are inexorably tied to the set of conservation practices and cost estimates. Although we made an effort to evaluate a wide variety of conservation practices discussed in water quality literature, inclusion of other possibly relevant practices (e.g., wetlands) may alter the results. This, however, is not so much a challenge to our modeling approach as an opportunity for further productive research.

Economists have long been able to point out that tradeoffs are ever-present in all of environmental policy, and in particular in nonpoint source pollution control. Tradeoffs are only meaningful when conservation policy options are efficient. Making such options explicit and thereby identifying numerous tradeoffs inherent in nonpoint source pollution control is the main contribution of this paper.

VIII. References.

Arabi, M., R. S. Govindaraju, and M. M. Hantush. 2006. Cost-effective allocation of watershed management practices using a genetic algorithm. *Water Resources Research* (42), W10429.

Arabi, M., J. Frankenberger, B. Engel, and J. Arnold. 2007. Representation of agricultural management practices with SWAT. *Hydrological Processes* (submitted).

Arnold, J.G., and P.M. Allen. 1996. Estimating hydrologic budgets for three Illinois watersheds. *Journal of Hydrology* 176: 57-77.

Arnold, J.G., and Fohrer, N. 2005. Current capabilities and research opportunities in applied watershed modeling. *Hydrological Processes* 19: 563-572.

Arnold, J.G., R. Srinivasan, R.S. Muttiah, and J.R. Williams. 1998. Large area hydrologic modeling and assessment part I: Model development. *Journal of American Water Resources Association* 34(1): 73-89.

Arnold, J.G., R.S. Muttiah, R. Srinivasan, and P.M. Allen. 2000. Regional estimation of base flow and groundwater recharge in the Upper Mississippi Basin. *Journal of Hydrology* 227: 21-40.

Arnold, J.G., R. Srinivasan, R.S. Muttiah, P.M. Allen, and C. Walker. 1999a. Continental scale simulation of the hydrologic balance. *Journal of American Water Resources Association* 35(5): 1037-1052.

Borah, D.K., and M. Bera. 2004. Watershed-scale hydrologic and nonpoint-source pollution models: Review of applications. *Transactions of ASAE* 47(3): 789-803.

Braden, J.B., G.V. Johnson, A. Bouzaher, and D. Miltz. 1989. Optimal spatial management of agricultural pollution. *American Journal of Agricultural Economics* 61: 404–13.

Gassman, P.W., S. Secchi, M. Jha, and L. Kurkalova. 2006. Upper Mississippi river basin modeling system part 1: SWAT input data requirements and issues. In: Coastal Hydrology and Processes (ed. V.P. Singh and V.J. Xu). Water Resources Publications, Highlands Ranch, Colorado, pp. 103-116.

Gassman, P.W., M. Reyes, C.H. Green, and J.G. Arnold. 2007. The Soil and Water Assessment Tool: Historical development, applications, and future directions. *Transactions of ASABE*. (forthcoming).

Gassman, P.W., Campbell, T., Secchi, S., Jha, M., and Arnold, J.G., 2003. The i_SWAT software package: a tool for supporting SWAT watershed applications. In: SWAT2003:

The 2nd International SWAT Conference, 1-4 July, Bari, Italy. Instituto di Ricerca sulle Acque, IRSACNR,Bari, Italy. pp. 66-69.

Jayakrishnan, R., R. Srinivasan, C. Santhi, and J.G. Arnold. 2005. Advances in the application of the SWAT model for water resources management. *Hydrologic Processes* 19(3): 749-762.

Jha, M., P.W. Gassman, S. Secchi, and J. Arnold. 2006. Upper Mississippi river basin modeling system part 2: baseline simulation results. In: Coastal Hydrology and Processes (ed. V.P. Singh and V.J. Xu). Water Resources Publications, Highlands Ranch, Colorado, pp. 117-126.

Jha, M., P.W. Gassman, S. Secchi, and J. Arnold. 2006. Upper Mississippi river basin modeling system part 2: baseline simulation results. In: Coastal Hydrology and Processes (ed. V.P. Singh and V.J. Xu). Water Resources Publications, Highlands Ranch, Colorado, pp. 117-126.

Khanna, M.,W. Yang, R. Farnsworth, and H. Onal. 2003. Cost effective targeting of CREP to improve water quality with endogenous sediment deposition coefficients. *American Journal of Agricultural Economics* 85: 538–53.

Kling, C. L., S. Secchi, M. Jha, H. Feng, and P.W. Gassman. 2006. Upper Mississippi river basin modeling system part 3: conservation practice scenario results. In: Coastal Hydrology and Processes (ed. V.P. Singh and V.J. Xu). Water Resources Publications, Highlands Ranch, Colorado, pp. 127-134.

Kurkalova, L.A., Burkart, C., and Secchi, S., 2004. Cropland cash rental rates in the Upper Mississippi River Basin. 2004. Technical Report 04–TR 47. Center for Agricultural and RuralDevelopment, Iowa State University, Ames, Iowa. Available at: http://www.card.iastate.edu/publications/DBS/PDFFiles/04tr47.pdf.

Mitchell, M. 1996. An introduction to genetic algorithms. Cambridge, Massachusetts: the MIT Press.

Montgomery, W.D. Markets in licenses and efficient pollution control programs. *Journal of Economic Theory* 1972; 5; 395-418.

NAS. 2000. The changing face of the UMR Basin; agriculture: Selected profiles of farming and farm practices. National Audubon Society, Upper Mississippi River Campaign, St. Paul, Minnesota. http://www.umbsn.org/news/documents/chg_face.pdf.

Rabalais, N.N., Turner, R.E., and Scavia, D. 2002. Beyond science into policy: Gulf of Mexico hypoxia and the Mississippi River. *BioScience* 52(2): 129-141.

Ribaudo, M.O. 1986. Consideration of off-site impacts in targeting soil conservation programs. *Land Economics* 62: 402–11.

Ribaudo, M.O. 1989. Targeting the conservation reserve program to maximize water quality benefits. *Land Economics* 65: 320–32.

Santhi, C., J.G. Arnold, J.R. Williams, W.A. Dugas, R. Srinivasan, and L.M. Hauck. 2001a. Validation of the SWAT model on a large river basin with point and nonpoint sources. *Journal of American Water Resources Association* 37(5): 1169-1188.

Secchi, S., M. Jha, L. A. Kurkalova, H. Feng, P.W. Gassman, C.L. Kling. 2005. The designation of co-benefits and its implication for policy: water quality versus carbon sequestration in agricultural soils. CARD Working Paper 05-WP 389, Iowa State University, Ames, Iowa.

Shortle, J. S. and R. D. Horan. 2001. The Economics of Nonpoint Pollution Control. *Journal of Economic Surveys* 15(3): 255-289.

Srivastava, P., J.M. Hamlett, P.D. Robillard, and R.L. Day. 2002. Watershed optimization of best management practices using AnnAGNPS and a genetic algorithm. *Water Resources Research* 38(3): 1-14.

United States Department of Agriculture, Economic Research Service. Fertilizer input data available at http://www.ers.usda.gov/Data/FertilizerUse.

United States Department of Agriculture, National Resource Conservation Service. Program cost data available at http://www.nrcs.usda.gov/PROGRAMS/EQIP/.

Veith, T.L., M.L. Wolfe, and C.D. Heatwole. 2003. Development of optimization procedure for cost-effective BMP placement. *Journal of the American Water Resources Association* 39(6): 1331-1343

Wall, M. 1996. GAlib: A C++ Library of Genetic Algorithm Components. Version 2.4.6. Available at: http://lancet.mit.edu/ga/.

Zitzler E., Laumanns M., Thiele L. 2002. SPEA2: Improving the Strength Pareto Evolutionary Algorithm for Multiobjective Optimization. Evolutionary Methods for Design, Optimisation, and Control, CIMNE, Barcelona, Spain, pp. 95-100.

Appendix

Table 1A. Summary of cost estimates, by state.

State name	Annualized cost of GW, per protected acre, $	Mean cash rental rate, $/acre	Cost of No-Till, $/acre	Annualized cost per terrace-protected acre, $
Illinois	7.4	110.1	22.2	22.0
Iowa	5.3	123.2	9.6	51.6
Minnesota	5.3	70.5	10.8	40.2
Missouri	3.9	65.4	12.9	13.7
Wisconsin	13.1	66.2	51.9	24.0

Table 2A. Estimates of cost of 20 percent nitrogen fertilizer application reduction.

yield zone	State	N application, lb/acre	20% reduced	Corn-Corn Yield drag, bu	Cost, C-C, $/year	Corn-SB Yield drag, bu	Cost, C-S, $/year
1	Illinois (North)	157.1	125.7	6.0	13.3	2.3	5.1
2	Illinois (Central)	157.1	125.7	4.4	9.7	6.0	13.2
2	Missouri(North)	153.4	122.8	4.8	10.5	6.1	13.5
3	Illinois(South)	157.1	125.7	5.4	12.0	4.2	9.3
3	Missouri (Central)	153.4	122.8	5.6	12.4	4.4	9.7
4	Iowa	125.3	100.2	8.0	17.7	3.0	6.6
5	Minnesota	114.1	91.3	6.3	13.9	5.1	11.3
6	Wisconsin	87.8	70.2	6.2	13.7	4.6	10.1

Figure 1A. Yield zones in the UMRB.

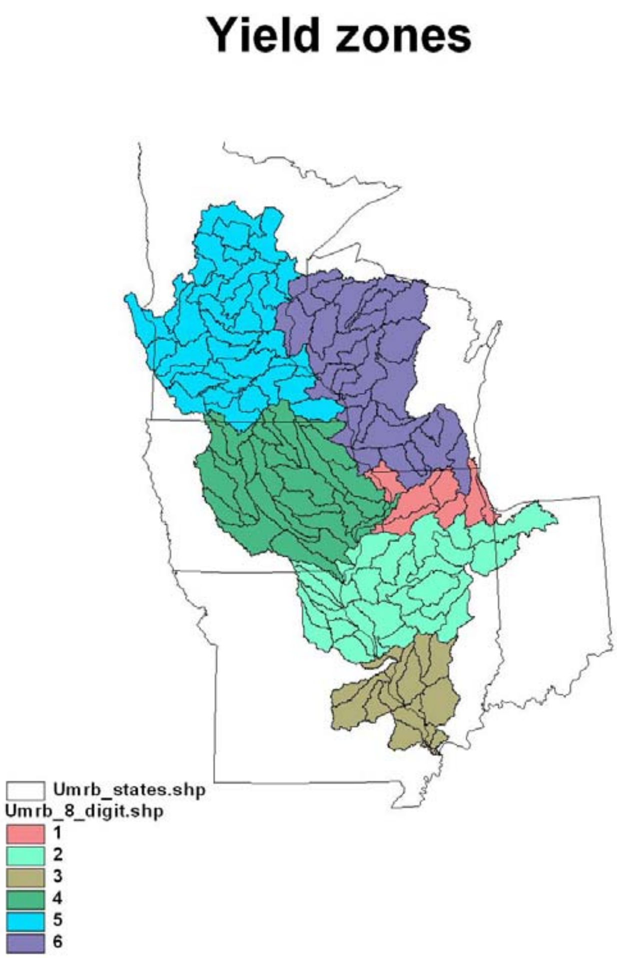

SPEA2 Fitness assignment.

An individual $\mathbf{i}$ is assigned a strength value $S(\mathbf{i})$ which equals to the number of solutions it dominates:

(0.4) $$S(\mathbf{i}) = \left|\{\mathbf{j} \mid \mathbf{j} \in \mathbf{P}_t \cup \overline{\mathbf{P}}_t \wedge \mathbf{i} \succ \mathbf{j}\}\right|,$$

where $\overline{\mathbf{P}}_t$ is the original population at generation t, $\mathbf{P}_t$ is the temporary population created, $|\cdot|$ denotes the cardinality of a set, and $\succ$ corresponds to the Pareto dominance relation. On the basis of this definition of strength values, the raw fitness for individual $\mathbf{i}$ is calculated:

(0.5) $$R(\mathbf{i}) = \sum_{\mathbf{j} \in \mathbf{P_t} \cup \bar{\mathbf{P}}_t, \mathbf{j} \succ \mathbf{i}} S(\mathbf{j}).$$

Thus, the raw fitness of an individual is determined by the strength of the dominators (individuals that dominate $\mathbf{i}$). Then, the raw fitness value of $R(\mathbf{i}) = 0$ corresponds to a nondominated individual, while a high raw fitness value corresponds to an individual that is dominated by many other individuals (which in turn dominate other individuals). In light of this interpretation, fitness minimization used in the formulation of the algorithm makes intuitive sense. Figure 2 demonstrates the fitness assignment process and highlights the fact that individuals that are located in the “crowded” areas of the objective space get a higher raw fitness value, and therefore are less likely to be selected into a future generation. For instance, point F dominates points B, C, and A, and therefore gets a strength value of 3. Since point F is nondominated, its raw fitness is zero. Point D, on the other hand, dominates only A, and thus gets the strength value of one, but is dominated by point G, which itself dominates 3 points. Thus, point D gets the raw fitness value of 3. Point A is the ‘worst’ point in the objective space, as it is associated with the highest cost and pollution levels. It itself does not dominate any other points, but is dominated by points F, G (with a strength value of 3), H (with a strength value of 2), D (with a strength value of 1), and E (with a strength value of 1). Therefore, the raw fitness value for point A is 3+3+2+1+1=10. Recalling that in this algorithm, individuals with the lower fitness scores are considered ‘more fit’, it is clear that individual A is far less likely to survive into the next generation than, for example, point F.

Such assignment of raw fitness scores also takes into account the relative ‘isolatedness’ of candidate solutions in the objective space. Conceptually, we would like the resulting Pareto-optimal frontier to span a large portion of the objective space. Therefore, candidate solutions on the interior of the frontier are somewhat less preferred than those close to the edges. In the figure, for example, while both points B and C are dominated, point C is dominated by both points F and G by virtue of its 'interior' location in the objective space; whereas point B is dominated only by point F and not by point G: its pollution level is lower than that of G. As a result, point B has a raw fitness score of 3 as opposed to the score of 6 for C, and its 'genetic makeup' is therefore less likely to be eliminated in the subsequent generations.

Figure 2A. Raw fitness assignment in SPEA2.

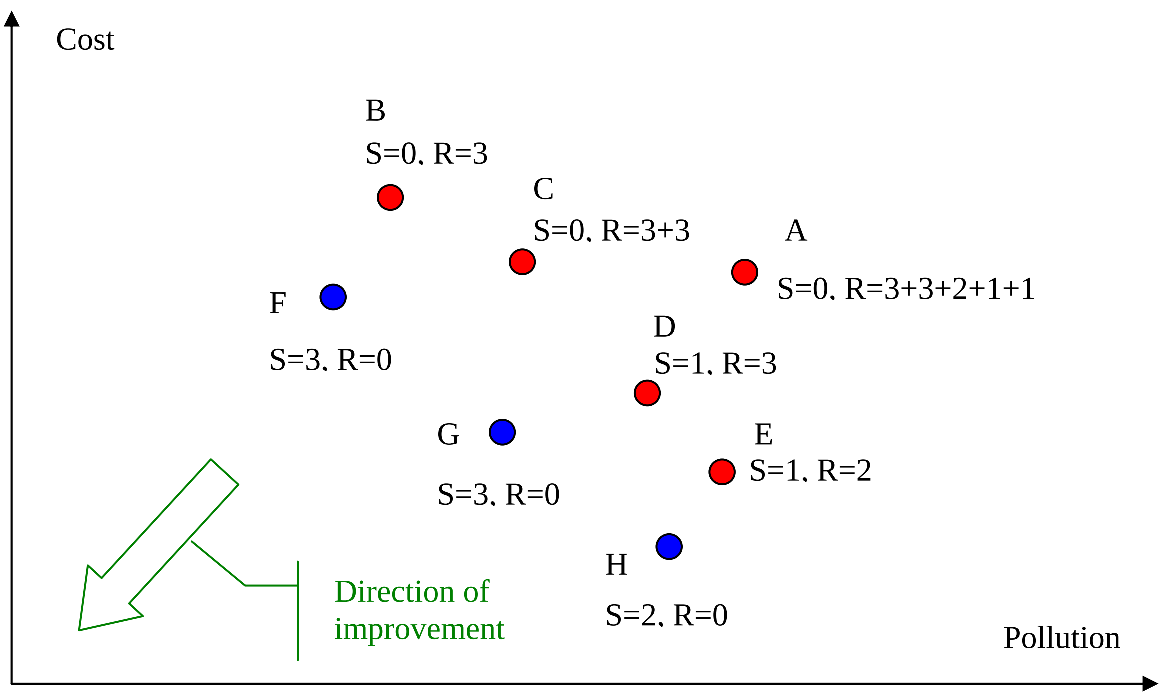

Finally, while the raw fitness score assignment outlined above incorporates some information on the location of the solutions in the solution space, additional density information is also incorporated into the calculation of a fitness score. Density estimation technique is used to further differentiate between individuals that are located in the "crowded" areas of the objective space (less preferred) from those located in the relatively sparse areas of the objective space (more preferred). The density estimation technique used in SPEA2 is an adaptation of the k-th nearest neighbor method, where the density at any point is a decreasing function of the distance to the k-th nearest data point. For each individual $\mathbf{i}$, we calculate the distances (in objective space) to all the individuals in the population and the temporary population, and store them in a list. After sorting the list in an increasing order, the k-th element yields the distance, denoted as $\sigma_{\mathbf{i}}^{k}$. k is chosen to equal to the square root of the sum of the population size and the size of the temporary population ($\sqrt{40+12} \approx 7$). The density is computed as:

(0.6) $$D(i) = \frac{1}{\sigma_{\mathbf{i}}^{k} + 2},$$

where 2 is added to the denominator to ensure that the value of the density is greater than zero and less than one.

Given the raw fitness score and the estimated density, the fitness of an individual $\mathbf{i}$ is calculated as:

(0.7) $$F(\mathbf{i}) = R(\mathbf{i}) + D(\mathbf{i}).$$

This is the fitness score used for selecting individuals in the algorithm implemented.