Handling Practicalities in Agricultural Policy Optimization for Water Quality Improvements
COIN Report Number 2017007

Brad Barnhart[a], Zhichao Lu[b], Moriah Bostian[c], Ankur Sinha[d], Kalyanmoy Deb[b], Luba Kurkalova[e], Manoj Jha[e], Gerald Whittaker[f]

[a]*U.S. EPA ORD NHEERL WED*
[b]*Michigan State University*
[c]*Lewis & Clark College*
[d]*Indian Institute of Management*
[e]*North Carolina A&T*
[f]*Oregon State University*

Abstract

Bilevel and multi-objective optimization methods are often useful to spatially target agri-environmental policy throughout a watershed. This type of problem is complex and is comprised of a number of practicalities: (i) a large number of decision variables, (ii) at least two inter-dependent levels of optimization between policy makers and policy followers, and (iii) uncertainty in decision variables and problem parameters. Given agricultural and economic data from the Raccoon watershed in central Iowa, we formulate a bilevel multi-objective optimization problem that accommodates objectives of both policy makers and farmers. The solution procedure then explicitly accounts for the nested nature of farm-level management decisions in response to agri-environmental policy incentives constructed by policy makers. We specifically examine the spatial targeting of a fertilizer-reduction incentive policy while seeking to maximize farm-level productivity while generating mandated water quality improvements using this framework. We test three different evolutionary optimization algorithms – m-BLEAQ, NSGA-II, and SPEA2 – and show that m-BLEAQ is well suited for handling the bilevel optimization problems and the considered practicalities.

1. Introduction

Nonpoint source pollution from agriculture remains a leading policy challenge for water quality improvement [23]. Key complicating factors include measurement problems for emissions, spatial heterogeneity of production over the agricultural landscape, and the hydrological dynamics within the watershed [22]. One promising area of research in agricultural and natural resource economics surrounds the use of integrated models that link management practices and production decisions to physical models of nutrient flow [25]. This integrated approach allows for optimal spatial targeting of management practices over the agricultural landscape and more explicit estimation of policy cost [29, 8].

In general, problems involving players, where the actions of one player (leader) are followed by the optimal response from another player (follower), are referred to as Stackelberg games. Such kind of games were first studied by Stackelberg [42]. Later in the 1970s, these problems were formulated as nested mathematical programs [9] and found applications in a wide variety of areas like defense [11, 44], transportation [26, 10], model production processes [27], chemical engineering [41, 12], and optimal tax policies [24, 34, 36], among others. A number of studies followed that included algorithms [3, 6, 4, 43, 45] These types of problems require nested optimization algorithms with one algorithm within the other, which leads to high computational requirements. Recently, meta-modeling approaches have gained importance for bilevel programming that exploit the properties of bilevel problems through mappings to solve the problem with lesser computational expense. A few studies in this direction are [33, 35, 1, 40, 38]. Multiobjective studies on bilevel optimization are few, both in the classical [16, 19] and evolutionary computation [47, 13, 37, 39] domain.

In this paper, we consider an agri-environmental policy targeting problem for the Raccoon watershed in Central Iowa, USA. The upper level decision maker (i.e., the policy maker in this case) faces multiple objectives of minimizing pollution and maximizing social benefit from agricultural production. However, the policy maker is incapable of solving her own optimization problem without taking the actions of the followers (i.e., the farmers) into account. For different policy decisions that the policy maker makes, the farmers react differently by maximizing their own objectives. This interaction makes the policy maker's problem a nested optimization task with the farmers' response embedded within the policy maker's problem. A deterministic version of this problem can be formulated as a bilevel optimization problem, which is a two level mathematical program with one program nested within the other. A natural extension of a deterministic problem is a problem that takes into account real-world uncertainties arising from the uncertain behavior of the players and other environmental parameters. We study both the deterministic as well as the uncertain version of the agri-environmental policy targeting problem in this paper.

We first provide a general formulation of a multiobjective bilevel problem, and one of the common techniques used to reduce it to a single-level problem. Thereafter, we discuss the agri-environmental policy targeting problem and provide its empirical models. Specifically, an empirical model for farm-level profits and yield is introduced along with another empirical ecohydrological model for the Raccoon watershed that we use to simulate the effects of fertilizer usage on water quality at the outlet of the watershed. Though the model considered in this paper has a bilevel structure, the regularity of the lower level allows us to write the problem as a multiobjective optimization problem. This allows us to solve the problem using a standard multiobjective optimizer when the optimal solution to the lower level can be written in closed form. We also solve the problem with a bilevel multiobjective evolutionary optimization algorithm to which the lower level optimal solution is not available directly. We compare the results for the two approaches and then incorporate additional real-world difficulties (risk-based uncertainties) in the model. Finally, we end the paper with some insights on future studies and conclusions.

2. Multiobjective Bilevel Model for Policy Problems

In this section, we provide a general formulation for policy problems where the upper level decision maker (policy maker) faces multiple objectives and the lower level decision maker faces a single objective. In other cases, where the lower level decision maker also has multiple objectives, the lower level problem can be replaced with a value function corresponding to the lower level decision maker, thereby reducing the lower level problem to a single objective optimization task. A bilevel problem has decision variables, constraints and objectives corresponding to each player. The constraints and objectives for each level may involve decision variables of the other player, thereby making the two levels interdependent. The lower level optimization is a parametric optimization task for which the upper level variables act as parameters, and the problem must be solved with respect to the lower level variables. Assuming a single leader and a single follower, the problem can be defined as follows:

Definition 1. For the upper-level objective function $F : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^p$ and lower-level objective function $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$, the bilevel optimization problem is given as:

$$
\begin{aligned}
&\underset{x,y}{\text{minimize}}\ F(x,y) = (F_1(x,y), \dots, F_p(x,y)) \\
&\text{subject to} \\
&\quad y \in \underset{y}{\text{argmin}}\{f(x,y) : g_j(x,y) \leq 0, j = 1, \dots, J\} \\
&\quad G_l(x,y) \leq 0, l = 1, \dots, L
\end{aligned}
$$

where $G_l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}, l = 1, \dots, L$ denotes the upper level constraints, and $g_j : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}, j = 1, \dots, J$ represents the lower level constraints.

The agri-envorinmental policy targeting problem is a different bilevel problem in which there are multiple followers, each acting independently and rationally, after the decisions of the leader. In such a case, the respective bilevel problem gets modified as:

Definition 2. For the upper-level objective function $F : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^p$ corresponding to a single leader and lower level objective functions $f_k : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ corresponding to followers $k \in \{1, \dots, K\}$, the bilevel optimization problem with single leader and multiple follower is given as:

$$
\begin{aligned}
&\underset{x,y}{\text{minimize}}\ (F_1(x, y^1, \dots, y^K), \dots, F_p(x, y^1, \dots, y^K)) \\
&\text{subject to} \\
&\quad y^1 \in \underset{y^1}{\text{argmin}}\{f^1(x,y^k) : g_j^1(x,y^1) \leq 0, j = 1, \dots, J\} \\
&\quad y^2 \in \underset{y^2}{\text{argmin}}\{f^2(x,y^k) : g_j^2(x,y^2) \leq 0, j = 1, \dots, J\} \\
&\qquad\qquad \vdots \\
&\quad y^k \in \underset{y^k}{\text{argmin}}\{f^k(x,y^k) : g_j^K(x,y^k) \leq 0, j = 1, \dots, J\} \\
&\quad G_l(x, y^1, \dots, y^K) \leq 0, l = 1, \dots, L
\end{aligned}
$$

where $G_l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}, l = 1, \dots, L$ denotes the upper level constraints, and $g_{k_j} : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}, j = 1, \dots, J$ represents the lower level constraints for followers $k \in \{1, \dots, K\}$, respectively.

Due to the presence of multiple followers each involving a different optimization problem, this problem is more computationally challenging to solve than the usual bilevel problems presented in Definition 1.

2.1. *Single Level Reduction*

In bilevel optimization, it is possible to replace the lower level problems in Definitions 1 and 2 with its Karush-Kuhn-Tucker (KKT) conditions when the lower level problems adhere to certain regularity conditions. Despite the reduction to single level, such problems are still hard to handle because of the constrained search space induced because of the KKT conditions. In certain situations, it might even be possible to write the optimal solution to the lower level problem in a closed form, i.e. y(x); however, this is rare. If a closed form solution can be written, then the bilevel problem effectively gets reduced to an ordinary optimization task. For multiple-follower problems, closed form optimal solutions or any approximations to the lower level problem are extremely beneficial.

3. Agri-environmental Policy Targeting Problem

Next, we consider our agri-environmental policy targeting problem for the Raccoon River Watershed, which covers roughly 9,400 km^2 in West-Central Iowa.

Agriculture accounts for the majority of land use in the study area, with 75.3% of land in cropland, 16.3% in grassland, 4.4% in forests and just 4% in urban use [20]. The watershed has been widely studied [20, 17, 21], and more recently in a bilevel framework [7], due to its high concentration of nitrate pollution from intensive fertilizer and livestock manure application. Nitrate concentrations routinely exceed Federal limits, with citations dating back to the late 1980s. The watershed also serves as the main source of drinking water for more than 500,000 people in Central Iowa. As a consequence, the Des Moines Water Works (DMWW) has for many years operated the world's largest nitrate removal facility. In 2015, a spike in nitrate levels forced the facility to run its nitrate removal equipment for a record 177 days. This culminated in the well-publicized federal lawsuit filed that year by the DMWW against three area counties for failing to control nitrate pollution from agricultural non-point sources.

3.1. *Problem formulation*

Given the above issues, one of the objectives the policy maker faces is to reduce the extent of pollution caused by agricultural activities by controlling the amount of fertilizer used. However, at the same time the policy maker does not intend to hamper agricultural activities to an extent of causing significant economic distress for the farmers.

Consider a producer (follower) $k \in \{1, \ldots, K\}$ trying to maximize her profits from agricultural production through N inputs $x^k = \{x_1^k, ..., x_N^k\}$ and M outputs $y^k = \{y_1^k, ..., y_M^k\}$. Out of the N inputs, x_N^k denotes the fertilizer input for each farm. The policy maker must choose the optimal spatial allocation of taxes, $\tau = \{\tau^1, \ldots, \tau^K\}$, for each farm corresponding to the Nitrogen fertilizer usage $x_N = \{x_N^k, \ldots, x_N^k\}$ so as to control the use of fertilizers. Tax vector $\tau = \{\tau^1, \ldots, \tau^K\}$ denotes the tax policy for K producers that is expressed as a multiplier on the total cost of fertilizers. Note that the taxes can be applied as a constant for the entire basin, or they can be spatially targeted at semi-aggregated levels or individually at the farm-level. For generality, in our model we have assumed a different tax policy for each producer. The objectives of the upper level are to to jointly maximize environmental benefits, $B(x_N)$—which consists of the total reduction of non-point source emissions of Nitrogen runoff from agricultural land—while also maximizing the total basin's profit $\Pi(\tau, x, y)$. The optimization problem that the policy maker needs to solve in order to identify a Pareto set of efficient policies is given as follows:

$$\max_{\tau,x,y} \; F(\tau, x) = \big(\Pi(\tau, x_N),\, B(x_N)\big) \tag{1}$$

$$\begin{aligned}
\text{s.t.} \quad & (x^k, y^k) \in \underset{x^k,\, y^k}{\operatorname{argmax}} \{\pi^k(\tau^k, x^k, y^k) : (\tau^k,\, x^k,\, y^k) \in \Omega^k\} \\
& \qquad\qquad \forall\, k \in \{1, \ldots, K\}, \\
& x_n^k \geq 0,\ \forall\, k \in \{1, \ldots, K\}, n \in \{1, \ldots, N\}, \\
& y_m^k \geq 0,\ \forall\, k \in \{1, \ldots, K\}, m \in \{1, \ldots, M\}, \\
& \tau^k \geq 1,\ \forall\, k \in \{1, \ldots, K\},
\end{aligned} \tag{2}$$

The fertilizer tax, τ, serves as a multiplier on the total cost of fertilizer, so that $\tau \geq 1$. The environmental benefit is the negative of pollution and therefore can be written as the negative of the total pollution caused by the producers, i.e. $B(x_N) = -\sum_{k=1}^{K} p(x_N^k)$. Similarly the total basin profit can be written as the sum total of individual producer's profit, i.e. $\Pi(\tau, x, y) = \sum_{k=1}^{K} \pi(\tau^k, x^k, y^k)$. The lower level optimization problem for each agricultural producer can be written as:

$$\begin{aligned}
\max_{x^k, y^k} \pi^k(p^k, w^k, \tau^k) = \; & p^k y^k - \sum_{n=1}^{N-1} w_n x_n^k - \tau_N^k w_N x_N^k \\
\text{s.t.} \quad & y^k \leq P^k(x^k),
\end{aligned} \tag{3}$$

where w and p are the costs and prices of the inputs x and outputs y, respectively. $P^k(x^k)$ denotes the production frontier for producer k. Heterogeneity across producers, due primarily to differences in soil type, may prevent the use of a common production function that would simplify the solution of (1). Likewise, the environmental benefits of reduced fertilizer use vary across producers, due to location and hydrologic processes within the basin. This also makes the solution of (1) more complex.

3.2. Empirical Model for Iowa Crop Production

The Iowa crop production model [32] is used to analytically model the economics of agricultural producers in the Raccoon watershed. The model operates on a 56 meter square grid based on USDA GIS remote-sensing crop cover maps. To account for producer heterogeneity across grid units, we use the ICCPI soil quality index, a measure of soil productivity that comes from GIS-based soil data. We model three crop rotations currently practiced in Iowa: continuous corn, corn-soybeans, and corn-corn-soybeans. The model captures the differentiated impacts of rotations on the yields. Crop production outputs, y_m, $m = 1$ (corn), 2 (soybeans), are expressed as functions of soil productivity, q, previous crop, tillage system, s, and nitrogen fertilizer application, x_N.

The corresponding corn and soybean current yield functions, in bushels per acre, are:

$$y_1^k = q^k[\alpha_0 \gamma_2^{s_2^k} \gamma_3^{s_3^k} \gamma_4^{(1-d_c^k)s_2} \gamma 5^{(1-d_c^k)s_3^k} + \alpha_1 d_c^k \\ + \alpha_2(x - \rho d_c^k) + \alpha_3(x_N^k - \rho d_c^k)^2] \tag{4}$$

and

$$y_2^k = q^k[\alpha_0 \gamma_2^{s_2^k} \gamma_3^{s_3^k} + \alpha_1 d_c^k + \alpha_2 d_{cc}^k], \tag{5}$$

where d_c indicates corn as the previous crop, d_{cc} indicates corn as the previous two crops, s_2 indicates mulch as the tillage choice, and s_3 indicates no-till as the tillage choice. The ICCPI productivity measure, q, serves as a multiplier, ranging from 0 to 1. There are three possible rotations: corn-corn (cc), corn-soy (cs), corn-corn-soy (ccs). Note that nitrogen fertilizer is only applied to corn, and that current yields depend on position within a given rotation.

Given the above functions, the optimal response of the followers can be written as a closed form solution in terms of the upper level decision variables. Profit maximization implies the optimal nitrogen fertilizer input demands as

$$x_N^{k^*} = \rho + \frac{\tau_N w_N}{p_1 2\alpha_3 q^k} - \frac{\alpha_2}{2\alpha_3}, \tag{6}$$

for corn after corn, and

$$x_N^{k*} = \frac{\tau_N w_N}{p_1 2\alpha_3 q^k} - \frac{\alpha_2}{2\alpha_3}, \tag{7}$$

for corn after soy.

3.3. The Biophysical Model

To model the water quality effects of prospective policy incentive schemes, we input the corresponding production model's fertilizer use into an empirical model for producing Nitrogen run-off. The model is $\text{runoff}^k = e^{-0.9819+1.668\log(1.12085 x_N^k)}$ [46]. The sum of all runoff from all farms is defined to be the runoff for the entire basin, which is minimized as an upper level objective. In our future work we intend to utilize a more accurate spatially explicit biophysical watershed model known as the Soil and Water Assessment Tool (SWAT) [2] to dynamically simulate Nitrogen runoff at the basin's outlet using local information related to land cover, soil type, and water management practices (e.g., tile drains). A number of studies apply SWAT to non-point policy for this reason [5, 31, 30].

4. Computational Algorithms Used in this Study

We use three different approaches to identify the Pareto-optimal set of policies for the upper level decision maker. In this section, we provide a brief description of the approaches that we have used to solve the above problem.

4.1. DEAP Implementation

Our first approach implements the Distributed Evolutionary Algorithms in Python (DEAP) [18], which is a flexible optimization framework. We used the traditional $(\mu + \lambda)$ evolutionary algorithm with SPEA2 selection [48] using DEAP. The algorithm creates a parent population of μ individuals and subsequently generates λ offspring within the population. After $(\mu + \lambda)$ individuals are tested, the results are sorted, and the μ best individuals from the $(\mu + \lambda)$ population are selected for the next generation. For this application, we set $\mu = 100$ and $\lambda = 100$; these values were chosen to produce a suitably large population without incurring substantial computational burden. Crossover and mutation probabilities were set to 0.90 and 1/(number of farms), respectively, following conventions [18]. DEAP is highly integrated with the SCOOP library in Python [28], which enables parallel implementation of the algorithm to reduce computational burden.

4.2. m-BLEAQ Implementation

The second approach utilizes m-BLEAQ [39]. The algorithm draws insights from the mathematical structure and properties of bilevel optimization problems and attempts to reduce the computational costs. The method creates localized quadratic approximations of the rational reaction set of the followers as a function of the leader's decision variables. This helps in reducing the number of lower level optimization calls that are required to solve the bilevel optimization problem. If the localized quadratic approximation of the rational reaction set created by the algorithm is accurate, then the optimal response of the follower in the vicinity of the local approximation can be directly obtained by using the approximation instead of solving the lower level optimization problem. Such local approximations of the rational reaction set are useful to guide the algorithm to the regions that might contain the bilevel optimal solutions without wasting computations in poorer regions. As the population members converge towards the Pareto-optimal solutions, the approximation continues to improve. In this study, while using m-BLEAQ we do not provide the algorithm the lower level optimal solutions; instead, we let the algorithm handle the entire bilevel problem.

4.3. NSGA-II Implementation

We also use NSGA-II [15] to solve the bilevel problem in this study. Since NSGA-II is capable of handling single level multi-objective problems, we therefore eliminate the lower level completely and directly provide the optimal lower level solution to the upper level using the empirical solutions in 6 and 7.

5. Results

In this section, we provide the simulation results comparing m-BLEAQ, NSGA-II, and SPEA2 on the empirical model of the Iowa Racoon watershed. Results on a small size, scaled size of the problem, robustness study will be illustrated one at a time in the following of the section.

5.1. Ten-Farm Problem

In this section, we provide the simulation results of the three methods using a small-sized problem, in which only ten farms are randomly chosen from the total 1,175 farms available. The Pareto-optimal frontiers are plotted in Figure 1. It is clear that all three methods achieve similar results. We calculated the hypervolume for all three methods as well (see Figure 2). The m-BLEAQ method is a steady-state algorithm, in which two members in the population at each generation are compared with two offspring created from three parents. Hence, the hypervolume comparison is plotted versus the function evaluations. The hypervolume is normalized with the maximum hypervolume achieved by each of the three methods.

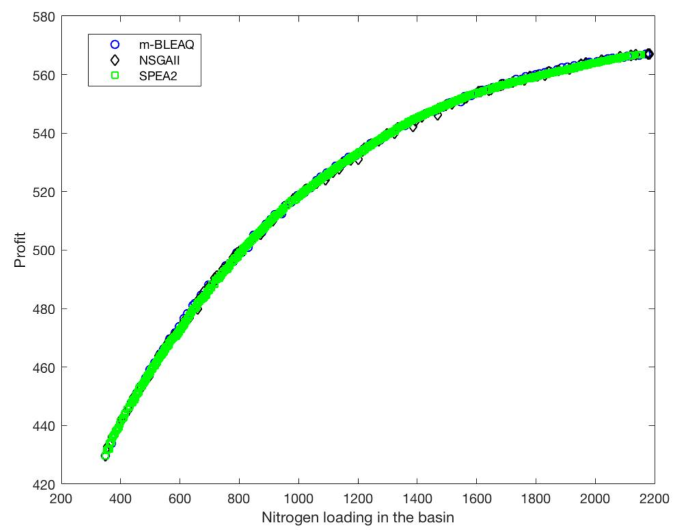

Figure 1: Pareto-optimal frontiers comparison on 10-farm case.

5.2. Scaling Up

In this section, we extend the dimensionality of the problem by increasing the considered farms in the lower level to evaluate each method's scalability.

We first present the results on 500 farms in the lower level. The Pareto-optimal frontiers shown in Figure 3 and the hypervolumes shown in Figure 4 suggest that the SPEA2 method implemented in DEAP encounters difficulty in handling the 500-farm problem, but the m-BLEAQ and NSGA-II methods do not. In addition, m-BLEAQ takes fewer function evaluations to approach the Pareto-optimal frontier as compared to the other two methods.

Finally, we include all 1,175 farms into our lower level consideration. Because of the high-dimensionality in the decision-variable space, we used the problem information to seed the initial population for the m-BLEAQ and

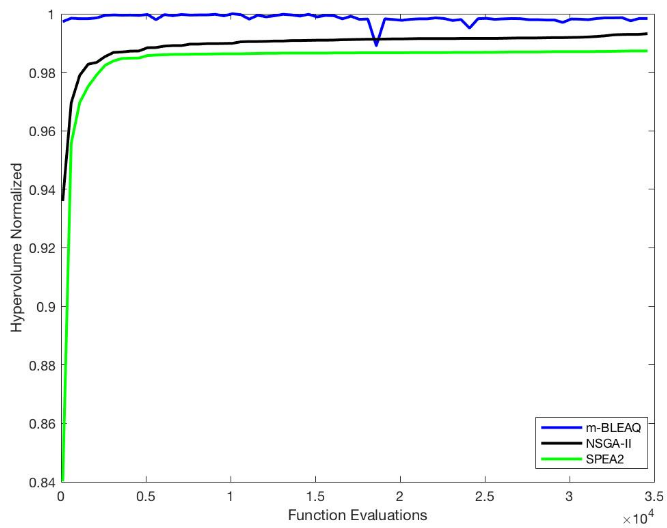

Figure 2: Hypervolume comparison on 10-farm case.

NSGA-II algorithms to improve the performance and convergence speed. We calculate the tax policy decision that will make the optimal nitrogen fertilizer input to be zero for each of the 1,175 farms by using the Equation 6 and 7 depending on the rotation. This tax policy should correspond to one of the extreme points on the Pareto-optimal frontier that minimize the nitrogen run-off in the watershed. We present the simulation results comparing the three algorithms in Figures 5 and 6.

According to the Pareto-optimal frontier comparison in Figure 5, m-BLEAQ performs the best among the three considered algorithms, and it is clear that the majority of the solutions on the frontiers of NSGA-II and SPEA2 get dominated by the solutions from m-BLEAQ, except for the extreme points. DEAP's SPEA2 implementation still struggles, and it achieves reduced profits at the lowest runoff value compared to the other two algorithms. This is likely due to the seeding mechanism used, since the DEAP SPEA2 algorithm was only seeded with large (i.e., 10x) tax values for each farm rather than setting the optimal nitrogen fertilizer input to be zero. This may have led to reduced production (and therefore profits) compared to the other two algorithms.

From the hypervolume comparison in Figure 6, one can clearly observe the difference in the final hypervolume achieved by the three algorithms,

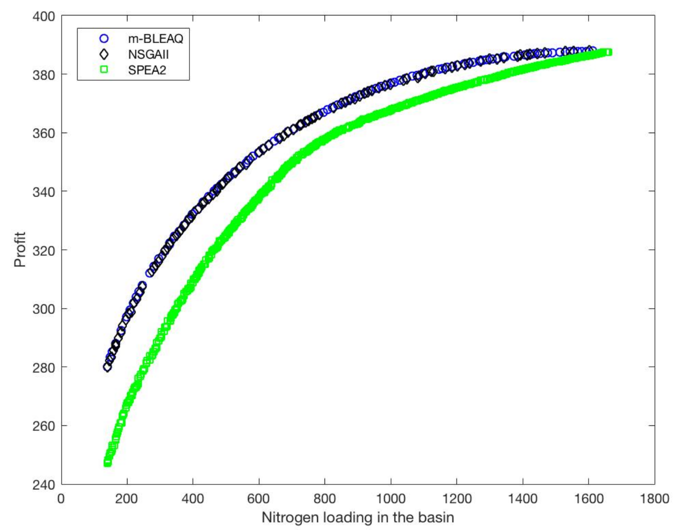

Figure 3: Pareto-optimal frontiers comparison on 500 farms case

and surprisingly m-BLEAQ could still approach the Pareto-optimal frontier with very few function evaluations.

5.3. *Effect of Recombination Operators in NSGA-II*

The dramatic difference between the results obtained from m-BLEAQ and NSGA-II is surprising to the investigators since the m-BLEAQ solution procedure is very similar to NSGA-II when the lower level optimal reaction function is readily available. Despite the steady-state nature, m-BLEAQ uses a vector-wise recombination operator, Parent Centric Crossover (PCX), that is different from the variable-wise recombination operator, Simulated Binary Crossover (SBX), used in NSGA-II. Also, studies have shown that PCX is superior to SBX in handling variable linkages and complex problems [14]. Therefore, we now compare NSGA-II, NSGA-II with PCX, and m-BLEAQ to further validate this finding.

It is clear from Figure 7 that the variable-wise recombination operator cannot properly handle the high-dimensionality in the problem, but a vector-wise recombination (PCX) using NSGA-II achieves similar results to that obtained by the bilevel m-BLEAQ approach.

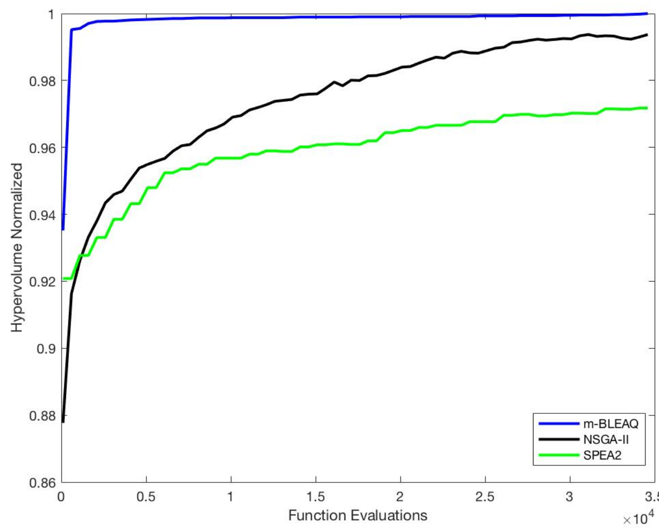

Figure 4: Hypervolume comparison on 500 farms case

5.4. Uncertainty Handling and Robust Solutions

In this section, we study the effect of uncertainties that may arise in corn production given a certain tax policy decision from the policymaker in the upper level. We assume the empirical model of the corn production presented in section 3.2 is an optimistic (over) estimation of the corn yield in actual implementation, and we also assume that the uncertainties that may arise in the corn production follows a Gaussian distribution. We then model the uncertain corn yield function $y_1'^k$ as below:

$$y_1'^k = y_1^k - \sigma \cdot y_1^k \cdot |N(0,1)|$$

where $N(0,1)$ is a normal random variable and σ (standard deviation) is the parameter that we used to control the severity of the uncertainties. The resulting uncertain corn yield function with different values of σ then follows the probability density function shown in Figure 8.

To account for the uncertainties, we sample a number of instantiations from the above uncertain corn yield function model. Then, the averaged function value is used for the subsequent corn production profit calculation. We present the simulation results comparing the final Pareto-optimal frontiers under different values of σ in Figure 9. The sample size used in the

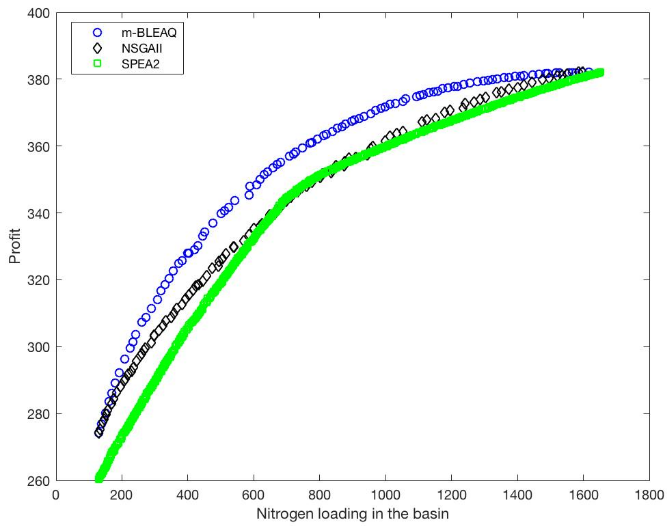

Figure 5: Pareto-optimal frontiers comparison on 1,175-farm case.

simulation is 31.

One can clearly notice the reduction in profit at every nitrogen loading compared to the deterministic case when the corn yield function is subject to uncertainties. Also, the amount of reduction increases as the standard deviation (σ) of the uncertain corn yield function increases. These results are as expected since the deterministic frontier corresponds to the optimistic estimation of the corn production and as the fluctuation in the corn production increase, more sacrifice in profit should be predicted.

However, in the meantime, more robustness should also be expected, since the standard deviation of the uncertainties that the problem is optimized against increase. We perform a follow-up study to quantify the robustness (or the sensitivity) achieved by comparing the deterministic and the robust frontier. We first collect the decision variables (tax policy) corresponding to the deterministic Pareto-optimal frontier. Then we re-evaluate these decision variables 31 times using the uncertain corn yield function with $\sigma = 0.1$ and record the minimum, median, and maximum objective function values. We repeat the same procedure for the decision variables from the robust Pareto-optimal frontier with $\sigma = 0.1$. The results are presented in Figure 10 and 11.

Figure 10 provides an overview of the profit fluctuation comparing the

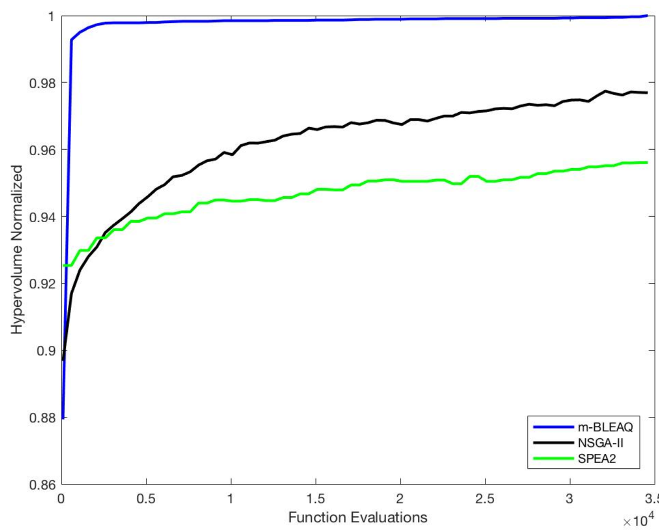

Figure 6: Hypervolume comparison on 1,175-farm case.

deterministic and robust solutions from the 31 trials performed in the above experiment. Since the deterministic frontier is obtained from the optimistic estimation of the corn production, the variation only fluctuates below the frontier. Examples at a few nitrogen loading levels show that noticeable reductions in the profit value fluctuation can be achieved from a slight sacrifice in median profit (as shown in Figure 11). These results indicate that a deterministic optimal solution is *brittle* when an uncertainty is expected in any problem parameter. However, when the same problem is solved for a robust solution, it is less vulnerable to the uncertainties in the problem parameters. Often, this comes at the sacrifice of the profit (objective of the study), and a robust optimization allows users to weigh the extent of sacrifice on the objective compared to increased robustness in the solution to make the whole approach practical.

6. Conclusions

Agricultural policy is a nested problem that includes interactions between policy makers and policy followers. Therefore, we presented a method for efficiently targeting agricultural policy using multiobjective bilevel optimization methods.

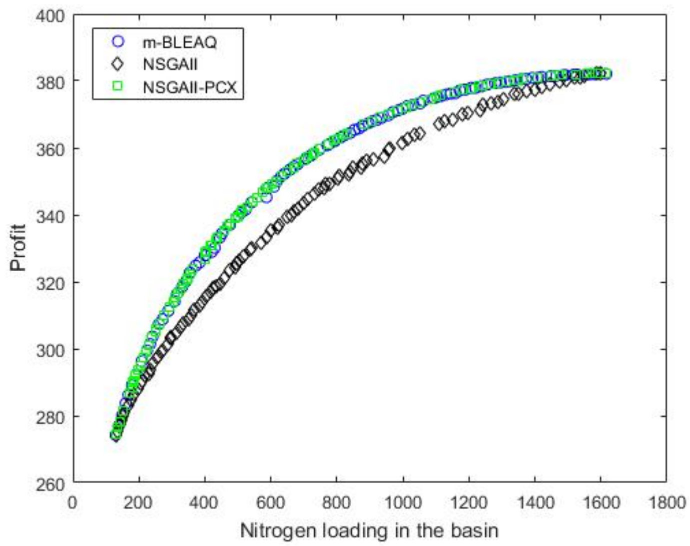

Figure 7: Comparison of results obtained from two recombination operators in NSGA-II on 1,175-farm problem.

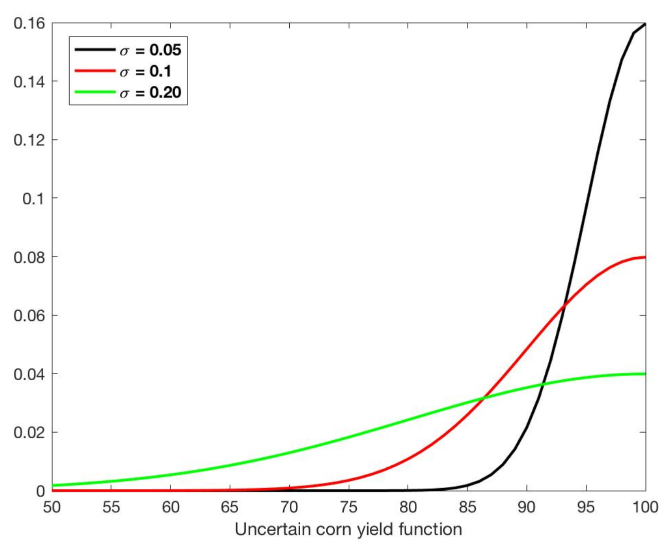

Figure 8: The probability density function of $y_1'^k$ with $y_1^k = 100$ used in the study.

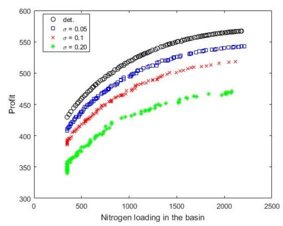

Figure 9: The deterministic and robust frontiers on the 10-farm problem.

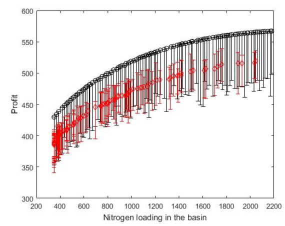

Figure 10: deterministic vs. robust optimization.

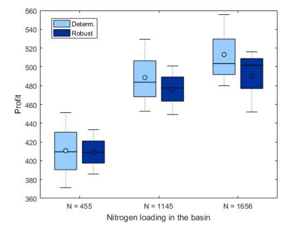

Figure 11: Robust solutions have smaller variations.

Overall, we demonstrated that, in general, the problem of targeting agricultural policy is a multiobjective bilevel problem that is characterized by (i) a large number of decision variables, (ii) at least two inter-dependent levels of optimization, and (iii) uncertainty in decision variables and problem parameters. We tested three different optimization implementations – DEAP, NSGA-II, and m-BLEAQ – and showed that m-BLEAQ is well suited for handling all of the difficulties associated with this problem. We also demonstrated that the choice of recombination operators is important when dealing with complex problems that have large numbers of decision variables. Our results showed that PCX is superior to SBX when handling complex problems, and that NSGA-II-PCX achieved similar results to m-BLEAQ when applied to the 1,175-farm problem. Finally, we also demonstrated a simple method for incorporating uncertainty into the multiobjective bilevel problem for targeting agricultural policy. We found that robust policy solutions can be achieved but only at the expense of upper level objectives. This is especially important for real-world policy planning, which oftentimes require assurances for robustness before implementation.

A large amount of research remains with regards to creating an effective optimization framework for targeting agricultural policy. Further studies should test different genetic algorithms and evolutionary methods to determine if better solutions are possible, especially those that increase convergence time. Also, these methods should be extended to include other alternative management practices including tillage choices and land retirement in addition to fertilizer reduction. The spatial heterogeneity of soil quality and weather will likely produce optimal spatial patterns for allocating different types of agricultural management practices throughout a watershed. Therefore, the spatial distribution of solution policy incentives should be further investigated to find patterns in optimal policy targeting. Also, in our study, we only investigated farm-targeted policies. For future work, mid-level aggregations should be utilized (e.g., zipcode, county-level) to accommodate more easily implementable policies.

References

[1] Jaqueline S Angelo, Eduardo Krempser, and Helio JC Barbosa. Differential evolution assisted by a surrogate model for bilevel programming problems. In *Evolutionary Computation (CEC), 2014 IEEE Congress on*, pages 1784–1791. IEEE, 2014.

[2] Jeffrey G Arnold, Raghavan Srinivasan, Ranjan S Muttiah, and Jimmy R Williams. Large area hydrologic modeling and assessment part i: Model development1, 1998.

[3] J. Bard and J. Falk. An explicit solution to the multi-level programming problem. *Computers and Operations Research*, 9:77–100, 1982.

[4] J. Bard and J. Moore. A branch and bound algorithm for the bilevel programming problem. *SIAM Journal on Scientific and Statistical Computing*, 11:281–292, 1990.

[5] Elias G Bekele and John W Nicklow. Multiobjective management of ecosystem services by integrative watershed modeling and evolutionary algorithms. *Water Resources Research*, 41(10), 2005.

[6] W. Bialas and M. Karwan. Two-level linear programming. *Management Science*, 30:1004–1020, 1984.

[7] Moriah Bostian, Gerald Whittaker, Brad Barnhart, Rolf Färe, and Shawna Grosskopf. Valuing water quality tradeoffs at different spatial scales: An integrated approach using bilevel optimization. *Water Resources and Economics*, 11:1–12, 2015.

[8] Whittaker G. Barnhart B. Färe R. Grosskopf S. Bostian, M. B. Valuing water quality tradeoffs at different spatial scales: An integrated approach using bilevel optimization. *Water Resources and Economics*, 11: 1–12, 2015.

[9] J. Bracken and J. McGill. Mathematical programs with optimization problems in the constraints. *Operations Research*, 21:37–44, 1973.

[10] Luce Brotcorne, Martine Labbe, Patrice Marcotte, and Gilles Savard. A bilevel model for toll optimization on a multicommodity transportation network. *Transportation Science*, 35(4):345–358, 2001.

[11] G. Brown, M. Carlyle, D. Diehl, J. Kline, and K. Wood. A Two-Sided Optimization for Theater Ballistic Missile Defense. *Operations Research*, 53(5):745–763, 2005.

[12] Peter A. Clark and Arthur W. Westerberg. Bilevel programming for steady-state chemical process design-i. fundamentals and algorithms. *Computers & Chemical Engineering*, 14(1):87–97, 1990.

[13] K. Deb and A. Sinha. An efficient and accurate solution methodology for bilevel multi-objective programming problems using a hybrid evolutionary-local-search algorithm. *Evolutionary Computation Journal*, 18(3):403–449, 2010.

[14] Kalyanmoy Deb, Dhiraj Joshi, and Ashish Anand. Real-coded evolutionary algorithms with parent-centric recombination. In *Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on*, volume 1, pages 61–66. IEEE, 2002.

[15] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.

[16] Gabriele Eichfelder. Multiobjective bilevel optimization. *Mathematical Programming*, 123(2):419–449, June 2010. ISSN 0025-5610.

[17] Hongli Feng, Manoj Jha, and Phil Gassman. The allocation of nutrient load reduction across a watershed: assessing delivery coefficients as an implementation tool. *Applied Economic Perspectives and Policy*, 31(1): 183–204, 2009.

[18] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner, Marc Parizeau, and Christian Gagné. Deap: Evolutionary algorithms made easy. *Journal of Machine Learning Research*, 13(Jul):2171–2175, 2012.

[19] N Gadhi and Stephan Dempe. Necessary optimality conditions and a new approach to multiobjective bilevel optimization problems. *Journal of Optimization Theory and Applications*, 155(1):100–114, 2012.

[20] Manoj K Jha, Philip Walter Gassman, and Jeffrey G Arnold. Water quality modeling for the raccoon river watershed using swat. *Transactions of the ASAE*, 50(2):479–493, 2007.

[21] Manoj K Jha, Calvin F Wolter, Keith E Schilling, and Philip W Gassman. Assessment of total maximum daily load implementation strategies for nitrate impairment of the raccoon river, iowa. *Journal of Environmental Quality*, 39(4):1317–1327, 2010.

[22] C.L. Kling. Economic incentives to improve water quality in agricultural landscapes: Some new variations on old ideas. *Journal of Agricultural Economics*, 93(2):297–309, 2011.

[23] L.A. Kurkalova. Cost-effective placement of best management practices in a watershed: lessons learned from Conservation Effects Assessment Project. *Journal of the American Water Resources Association*, 51(2):359–372, 2015.

[24] M. Labbé, P. Marcotte, and G. Savard. A Bilevel Model of Taxation and Its Application to Optimal Highway Pricing. *Management Science*, 44 (12):1608–1622, 1998.

[25] Shortle J. Wilen J. Zilberman D. Lichtenberg, E. Natural resource economics and conservation: Contributions of agricultural economics and agricultural economists. *American Journal of Agricultural Economics*, 92(2):469–486, 2010.

[26] Athanasios Migdalas. Bilevel programming in traffic planning: Models, methods and challenge. *Journal of Global Optimization*, 7(4):381–405, 1995.

[27] M.G. Nicholls. Aluminium Production Modeling - A Nonlinear Bilevel Programming Approach. *Operations Research*, 43(2):208–218, 1995.

[28] Per Kristian Nilsen and Nicolas Prcovic. Parallel optimisation in the scoop library. In *International Parallel Processing Symposium*, pages 452–463. Springer, 1998.

[29] Campbell T. Jha M. Gassman P. W. Arnold J. Kurkalova L. A. Secchi S. Feng H. Kling C. L. Rabotyabov, S. S. Least cost control of agricultural nutrient contributions to the Gulf of Mexico hypoxic zone. *Ecological Applications*, 20(6):1542–1555, 2010.

[30] Sergey S Rabotyagov, Manoj Jha, and Todd D Campbell. Nonpoint-source pollution reduction for an iowa watershed: An application of evolutionary algorithms. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie*, 58(4):411–431, 2010.

[31] CW Richardson, DA Bucks, and EJ Sadler. The conservation effects assessment project benchmark watersheds: synthesis of preliminary findings. *Journal of soil and water conservation*, 63(6):590–604, 2008.

[32] Silvia Secchi, Lyubov Kurkalova, Philip W Gassman, and Chad Hart. Land use change in a biofuels hotspot: the case of iowa, usa. *Biomass and Bioenergy*, 35(6):2391–2400, 2011.

[33] A. Sinha, P. Malo, and K. Deb. Efficient evolutionary algorithm for single-objective bilevel optimization. *arXiv preprint arXiv:1303.3901*, 2013.

[34] A. Sinha, P. Malo, A. Frantsev, and K. Deb. Multi-objective stackelberg game between a regulating authority and a mining company: A case study in environmental economics. In *2013 IEEE Congress on Evolutionary Computation (CEC-2013)*. IEEE Press, 2013.

[35] A. Sinha, P. Malo, and K. Deb. An improved bilevel evolutionary algorithm based on quadratic approximations. In *2014 IEEE Congress on Evolutionary Computation (CEC-2014)*, pages 1870–1877. IEEE Press, 2014.

[36] A. Sinha, P. Malo, and K. Deb. Transportation policy formulation as a multi-objective bilevel optimization problem. In *2015 IEEE Congress on Evolutionary Computation (CEC-2015)*. IEEE Press, 2015.

[37] A. Sinha, P. Malo, and K. Deb. Towards understanding bilevel multi-objective optimization with deterministic lower level decisions. In *Proceedings of the Eighth International Conference on Evolutionary Multi-Criterion Optimization (EMO-2015)*. Berlin, Germany: Springer-Verlag, 2015.

[38] A. Sinha, P. Malo, and K. Deb. Solving optimistic bilevel programs by iteratively approximating lower level optimal value function. In *2016 IEEE Congress on Evolutionary Computation (CEC-2016)*. IEEE Press, 2016.

[39] A. Sinha, P. Malo, K. Deb, P. Korhonen, and J. Wallenius. Solving bilevel multi-criterion optimization problems with lower level decision uncertainty. *IEEE Transactions on Evolutionary Computation*, 20(2):199–217, 2016.

[40] A. Sinha, P. Malo, and K. Deb. Evolutionary algorithm for bilevel optimization using approximations of the lower level optimal solution mapping. 2016 (In press), journal=European Journal of Operational Research, pages=, volume=, number=, doi=10.1016/j.ejor.2016.08.027, publisher=Elsevier.

[41] W. R. Smith and R. W. Missen. *Chemical Reaction Equilibrium Analysis: Theory and Algorithms*. John Wiley & Sons, New York, 1982.

[42] H. Stackelberg. *The theory of the market economy*. Oxford University Press, New York, Oxford, 1952.

[43] H. Tuy, A. Migdalas, and P. Värbrand. A global optimization approach for the linear two-level program. *Journal of Global Optimization*, 3:1–23, 1993.

[44] L. Wein. Homeland Security: From Mathematical Models to Policy Implementation: The 2008 Philip McCord Morse Lecture. *Operations Research*, 57(4):801–811, 2009.

[45] D. White and G. Anandalingam. A penalty function approach for solving bi-level linear programs. *Journal of Global Optimization*, 3:397–419, 1993.

[46] G. Whittaker, R. Confesor Jr, S. M. Griffith, R. FÃd're, S. Grosskopf, J. J. Steiner, G. W. Mueller-Warrant, and G. M. Banowetz. A hybrid genetic algorithm for multiobjective problems with activity analysis-based local search. *European Journal of Operational Research*, 193(1):195–203, 2009.

[47] Y. Yin. Genetic algorithm based approach for bilevel programming models. *Journal of Transportation Engineering*, 126(2):115–120, 2000.

[48] Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4):257–271, 1999.