Towards Practical Models of Complex Agricultural Systems

By

Yuji Saikai

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy
(Agricultural and Applied Economics)

at the
UNIVERSITY OF WISCONSIN–MADISON
2020

Date of final oral examination: 05/04/2020

The dissertation is approved by the following members of the Final Oral Committee:
Paul D. Mitchell, Professor, Agricultural and Applied Economics
Vivak Patel, Assistant Professor, Statistics
Sheldon Du, Associate Professor, Agricultural and Applied Economics
Shawn P. Conley, Professor, Agronomy
Guanming Shi, Professor, Agricultural and Applied Economics

To Kanan

Abstract

Modern agriculture faces some of the most pressing problems in the 21st century: farm profitability, food security, and environmental sustainability. To address these problems, it is essential to raise productivity in a sustainable manner, an overarching goal known as sustainable intensification. In efforts to increase productivity through improved management of agricultural systems, a fundamental challenge is enormous complexity arising from both biological and social aspects of agricultural systems. However, despite the apparent need for research in modeling such socio-ecological systems without trivializing their complexity, most of the existing agricultural research only focuses on basic component sub-processes, making itself largely irrelevant for practical decision making for sustainable intensification. My contribution to the research community is twofold: exemplify practical models for management of agricultural systems and lay the foundation for some specific problems. In particular, two distinct models are constructed to support decision making at different levels of agricultural systems. Both models are predominantly characterized by their computational approaches, which capitalize on the ever increasing data and computational capacity. At an individual level, adaptive experimental designs based on Bayesian optimization techniques help individual farmers to efficiently learn complex management practices through on-farm experiments. In contrast, agent-based models help policy makers to gain insights into complex socio-ecological systems and design effective mitigation policies for insect resistance management at an aggregate landscape level. Although these models are necessarily ad hoc solutions to the specific problems, their modeling techniques (Bayesian optimization and agent-based modeling) are very general and applicable to many other practical problems.

Acknowledgements

I owe my sincere gratitude to my two advisors, Paul Mitchell and Vivak Patel, for their support. I am an introverted, risk-taking, rebellious student who possesses only average intellectual capacity but cannot help being different. Despite the naysaying within the department, Paul always stood behind me no matter what. As a nonconformist, I might have been merely doing nonsense. Vivak patiently helped me with technical credibility and professional communication important regardless of academic disciplines. When needed, he also sheltered me from my circumstances. Without Paul and Vivak, my time in Madison would have been far more difficult and perhaps quite miserable. I will pay it forward by doing the same to my students in the future.

Contents

Chapter 1

Introduction

"All models are wrong, but some are useful."

—George E. P. Box (1976, 1979)

Modern agriculture faces some of the most pressing problems in the 21st century: farm profitability, food security, and environmental sustainability. To address these problems, it is essential to increase agricultural productivity in an environmentally sustainable manner, an overarching goal known as *sustainable intensification* (Tilman et al., 2011). Moreover, as indicated in Sustainable Development Goals (Food and Agriculture Organization of the United Nations, 2016), the productivity increase also needs to happen in terms of nutritional quality beyond bulk yield and development of multifunctional landscape (Jones et al., 2017). Since society-wide productivity is an aggregate result of individual productivity, each farming practice at an individual level must be improved and, to support the improvement, public policies have roles to play (Garnett et al., 2013).

In efforts to improve individual farming practices, a fundamental challenge is the enormous complexity of agricultural systems, which are often regarded as *managed ecosystems* (Antle and Capalbo, 2001, 2002; Swinton et al., 2007) where farmers intervene in the natural environment to produce food, fiber and energy. It is important to bear in mind that this complexity arises from both ecological and social aspects of agricultural systems. On the one hand, to manage ecosystems, farmers try to control complex bio-physical and chemical reactions towards desired economic outcomes. On the other hand, each farmer's management decisions are intrinsically influenced by and influence others' decisions through social interactions and economic conditions, which are endogenously influenced by their collective decisions (Arthur, 2013). Thus, at the center of sustainable intensification is development of software systems that support decision making in complex agricultural systems (Lindblom et al., 2017).

Despite the apparent need for research in modeling such socio-ecological systems without trivializing their complexity, most of the existing agricultural research is not designed for management of complex agricultural systems (Antle et al., 2017a; Capalbo et al., 2017) and largely irrelevant for practical decisions (McCown et al., 2009; McCown, 2002). Instead of confronting complex processes involved in practical agriculture and policy intervention, academic research tends to focus on scientific understanding of basic component processes, which are "more easily studied in a laboratory or institutional setting, and may result in more publishable findings. Producing useful decision tools for farmers or policy decision-makers is at best a secondary consideration in many academic settings" (Antle et al., 2017a, p.256). As a result of this "researcher-centricity", despite the major advances in digital and farming technology, agricultural models have not progressed much over the last 30 years; in other words, there are significant untapped resources that can help dramatically advance farming practices (Lindblom et al., 2017; Rose et al., 2018).

My contribution to the research community of agricultural sciences and economics is to build practical models that are intended to help practitioners make decisions on management of complex agricultural systems. All the models are predominantly characterized by their computational approaches, which capitalize on the ever increasing data and computational capacity. Chapter 2 and 3 are closely connected and both concern efficiency and efficacy of agricultural field experiments to learn optimal input use and management practices. In many branches of the agricultural sciences and commercial enterprises, field experimentation is the fundamental technique to gain insights into the mechanisms of complex biophysical systems in agriculture and how farmers should act upon such systems to maximize profits. My conjecture and presumption for the remainder of the dissertation is that gaining insights into the former (the mechanisms) is certainly helpful but not necessary and often inefficient for gaining insights into the latter (practical management). In other words, simple models for scientific understanding of agricultural systems are hardly useful for practical profit maximization in the complex systems. Given the current state of knowledge about the complexity of agricultural systems, pursuing scientific understanding is not the best route to help farmers and policy makers today. This is a key realization to bear in mind when addressing the issue of researcher-centricity and facilitating practical decision making.

Chapter 2 specifically deals with cost-efficient experimental designs of field trials which are conducted to learn optimal management choices. In most cases, traditionally, field experiments are based on factorial design. It is intuitive, easy to use, and can be effective for some problems. However, what is hidden in these advantages is inefficiency in choosing factor levels to investigate. The inefficiency becomes evident when the objective of experiments is to estimate an optimal choice under some response function, rather than to estimate the overall shape of the response function. Moreover, it is particularly problematic in realistic scenarios, where response functions of practical interest involve many management variables to optimize. In this case, factorial design is likely infeasible, requiring too many combinations of factor levels to investigate. As a solution, Chapter 2 presents a machine learning algorithm for an adaptive experimental design based on a numerical optimization technique called Bayesian optimization.

Chapter 3 extends the algorithm developed in Chapter 2 and applies it to profit maximization problems in precision agriculture. Thanks to advanced farming technologies, precision agriculture holds the promise of increasing productivity through optimizing site-specific management. In reality, however, these technologies are too sophisticated to fully exploit.

Indeed, decisions on site-specific management are still based on simple rule-of-thumb and not remotely optimal (Capalbo et al., 2017). In precision agriculture, there are not only many management factors but also many different growing environments, in each of which the management variables need to be optimized. To help farmers solve larger and more complex optimization problems than those dealt with in Chapter 2, the algorithm is modified so that individual farmers can efficiently learn their own optimal site-specific management for each site in a field.

In contrast to the previous two chapters, Chapter 4 investigates effectiveness of public policies for mitigating development of insect resistance to control technologies at a landscape level. While modeling impacts of pests and diseases on crop yield itself continues to be a significant challenge for the agricultural research community (Donatelli et al., 2017), Chapter 4 sheds light on another important yet often-neglected aspect of pest control problems—significance of social factors. To control insect damage, modern crop production relies on genetic engineering technologies including Bt (*Bacillus thuringiensis*) crops. Although insects eventually develop resistance to any technology and make it obsolete, its lifespan considerably depends on the rate of insect evolution, which in turn depends on how the technology is utilized by farmers. Since farmers' management decisions can be influenced by many social factors and regulations, there is room for designing effective public policies. Thus, to facilitate policy experiments in silico, an agent-based model is constructed and used to integrate key social factors with insect population dynamics. Instead of directly modeling a macro phenomenon of interest, the model is constructed from the bottom up; a macro phenomenon emerges as a result of interaction among heterogeneous agents (farmers) in a spatiotemporally explicit environment.

This dissertation is based on the humility to accept our inadequate understanding of agricultural systems to manage their complexity as well as the sense of urgency to sustainable intensification. To take immediate actions for sustainable agriculture and society at large, the prevailing researcher-centricity needs to be rectified. The dissertation attempts to shift the balance between researcher-centric scientific pursuit and practicality so that we can divert more research resources into practically useful models for sustainable intensification in the near future. To both exemplify the point and lay the foundation for some specific problems, the dissertation presents models that capitalize on the ever increasing data and computational capacity.

Chapter 2

Adaptive experimental design using Bayesian optimization to improve the cost efficiency of field trials

Yuji Saikai
Vivak Patel
Shawn P. Conley
Paul D. Mitchell

2.1 Introduction

In many branches of the agricultural sciences, field experimentation is the fundamental technique to determine optimal input use and management practices across different production environments (Mead et al., 2002; Petersen, 1994). For example, field trials have been used for determining optimal fertilizer application rates (Buresh et al., 2019; Cela et al., 2011; Huang et al., 2008; Jin et al., 2019a; Li et al., 2018; Rens et al., 2018; Storer et al., 2018; Wang et al., 2017; Wang et al., 2012), fungicide dose rates (Lynch et al., 2017), plant varieties (Huang et al., 2008; Lynch et al., 2017; Storer et al., 2018), seeding rates (Dai et al., 2013), and plant densities (Khan et al., 2017; Ren et al., 2017). Unfortunately, since any experiment is subject to limited resources and incurs economic costs, exhaustive field trials are infeasible. Indeed, most studies including those cited above involve only a handful of levels for each factor, and the total number of factors is very small. When involving many factors, each factor usually allows for only two levels (Orlowski et al., 2016). Thus, in determining optimal management, field experiments must be carefully designed to maximize scarce resources.

In most cases, field experiments are based on factorial designs, in which all possible combinations of the levels of every factor are investigated and replicated over several years (Montgomery, 2017). As an example, Jin et al. (2019a) conducted field experiments based on a factorial design of four biochar rates (0,5,20,40) and four nitrogen fertilizer rates (0,60,90,120), which were replicated from 2011 to 2016. As another example, Lynch et al. (2017) conducted field experiments based on a factorial design of two fungicide types (azole, azole+SDHI), five applications rates (0, 0.25, 0.50, 1.0, 2.0), and three varieties (SR5, SR7, SR8), which were replicated from 2012 to 2015. Although these and many other field experiments that employ factorial designs provide qualitative insights into dependencies between factors and clues about optimal levels, in terms of quantitatively maximizing scarce resources, they are problematic for several reasons:

1. With sufficient domain knowledge, hand-picked factor levels may work. But, they are less likely to work if growing environments and/or factors are considerably different from the existing ones on which such domain knowledge is based.
2. Even with well-chosen factor levels, the fixed number of levels limits exploration of factors that incorporate continuous values.
3. Finally, static experimental designs by definition preclude adaptation of designs that reflects the information obtained in the preceding years to explore more promising values.

Recently, Chen et al. (2019) address the inefficiency issue of factorial design by using orthogonal experimental design. However, their approach is still a static design over three years and not particularly designed for efficient optimization. Indeed, the study amasses a large number of samples (>40,000).

In addition to the inefficiency in data collection, the standard approaches to agricultural experimentation suffer from model misspecification that can potentially have a negative impact on profitability. After collecting data, researchers next estimate some statistical model of yield/profit function and maximize this estimated function to determine optimal management choices. The standard approaches typically use relatively simple statistical models of limited flexibility. Therefore, if the underlying function takes a complex shape, which is quite possible in cases involving many variables as we show in our experiments, such models will fail to capture the function's major structure and provide poor guidance on profit maximization. In contrast, our approach uses a sufficiently flexible statistical model to avoid model misspecification.

In summary, field experiments that use factorial designs for optimizing input use and management practices make inefficient use of scarce resources dedicated to experimentation. Moreover, simple statistical models that typically accompany factorial designs may suffer from model misspecification, further reducing efficiency. In this work, we propose an adaptive, flexible experimental designs for optimization. Our approach, based on Bayesian optimization (BO) (Brochu et al., 2010; Shahriari et al., 2016), reflects information from the preceding years in the current design, so as to make the most of scarce resources, and has the ability to capture potentially complex profit functions.

We demonstrate the advantage of our approach over the traditional factorial designs by using two distinct simulation environments: one constructed from real field-trial data and the other from the existing cropping systems simulator. As a result of our methodology, field experiments can be conducted more efficiently, allowing farmers to determine optimal input use and management practices with fewer experimental resources and, thereby, reducing the costs.

2.2 Materials and methods

We compare two regimes of experimental designs, BO design and factorial design, in six different simulation environments. Each simulation environment is characterized by an "oracle," a conceptual device that tells us the crop yield for a set of specific levels of production factors. Using the yield obtained from the oracle, we then calculate a profit using an output price and input costs. Once we set up the simulation environments, we carry out simulations over 1- to 10-year time horizons and compare their performances on profit maximization based on both years of experiment and number of sampling plots required to reach a certain level of profit.

2.2.1 Oracle and simulation environment

An oracle returns a yield when we query a set of specific values of production factors. It is a yield response function corrupted by random noise. Mathematically, yield (y) is defined as follows:

$$y = f(x) + \varepsilon,$$

where $f : \mathcal{X} \to \mathbb{R}_+$, a function that takes a set of production factors $x \in \mathcal{X}$ and returns a nonnegative yield $y \in \mathbb{R}_+$. Note that we assume $\varepsilon \sim N(0, \sigma^2)$, that is, normally distributed noise with standard deviation of σ. For each yield (y), profit (π) is defined as follows:

$$\pi = py - c \cdot x,$$

where p is an output price and c is a set of input prices. To characterize the effects of simulation environments on the performance of designs, we construct two distinct types of simulation environments: low-dimensional with discrete choice and high-dimensional.

Scenario A (low-dimensional with discrete choice)

This environment is meant to capture a traditional field study that involves a small number of both continuous and discrete factors. We base this simulation study on the following story.

> Imagine a researcher who tries to determine optimal seeding rates and seed treatments of a new seed variety under 16 different soil pHs $\{6.0, 6.1, \ldots, 7.5\}$. There are two types of seed treatments, F (fungicide) and FI (fungicide and insecticide), in addition to UTC (untreated control), where UTC is cheapest and FI is most expensive. Optimality is defined as maximizing the expected profit. Overall, through field trials, the researcher tries to determine 16 optimal combinations of seeding rate and seed treatment, one for each soil pH. Given the economic costs associated with each trial, the researcher wants to make efficient use of the limited experimental resources. Which combinations in what order should the researcher use to proceed with the experiment?

We consider this to be a reasonable setting, albeit stylized, because the ability to adjust soil pH is usually constrained both financially and physically such as buffer pH (Camberato, 2014). Hence, it is valuable to have a recommendation for each soil pH.

Let x_1 denote seeding rate (in 1,000 seeds/ac), x_2 denote soil pH, and τ denote seed treatment. Then, we have $x_1 \in [0, 150]$, $x_2 \in \{6.0, 6.1, \ldots, 7.5\}$, and $\tau \in \{UTC, F, FI\}$. The oracle for this setting is a collection of three yield response functions (f_τ for each τ). In other words, when we query a combination (τ, x_1, x_2), the oracle returns a noise-corrupted yield,

$$y = f_\tau(x_1, x_2) + \varepsilon.$$

We construct these three response functions using real-world data from soybean field trials (Gaspar et al., 2015). Specifically, for each seed treatment (τ), we fit a nonparametric regression, namely a local linear regression (Hastie et al., 2016). Then, for each yield, we calculate the corresponding profit (π) using a soybean price and seed cost for that treatment. Formally,

$$\pi = py - c_\tau x_1,$$

where p is the soybean price and c_τ is the seed cost for treatment τ, all of which are found in Gaspar et al. (2015).

The following figures depict three profit response surfaces (without noise). As you can see, each surface is not globally concave.

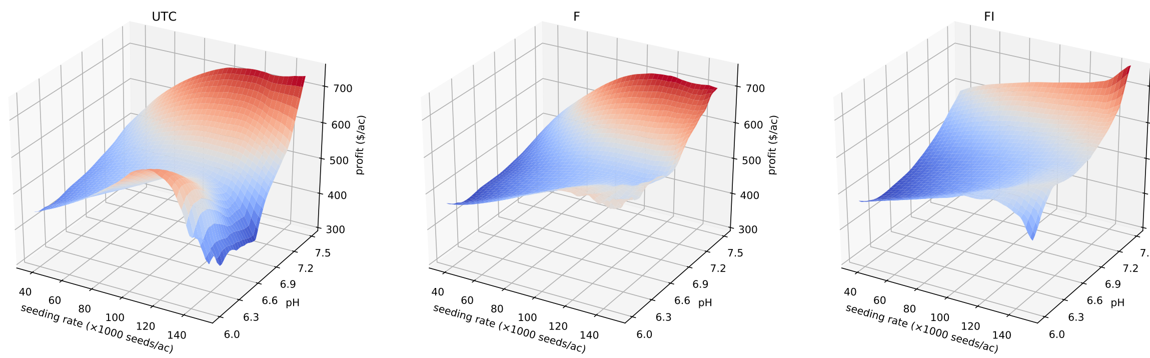

Figure 2.1: Profit response surfaces

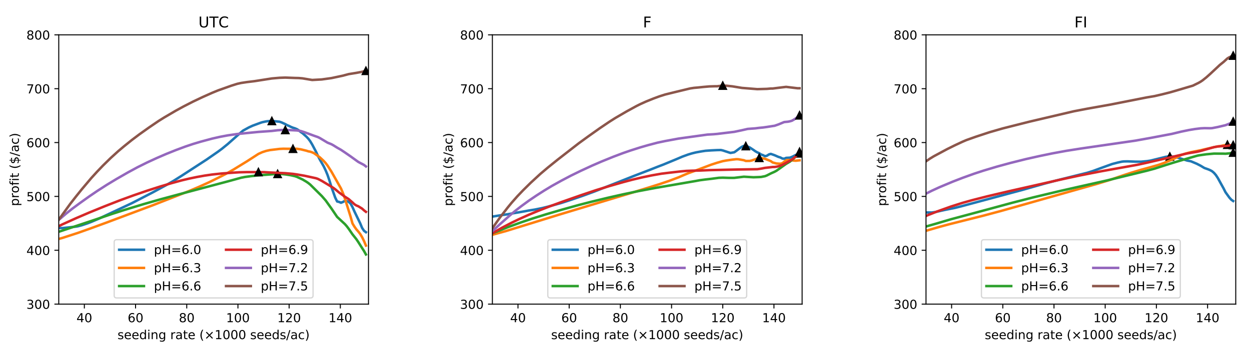

Figure 2.2: Profit response curves at the selected pHs

▲ indicates the maximum profit for each pH. The following table contains specific values.

pH	6.0	6.1	6.2	6.3	6.4	6.5	6.6	6.7	6.8	6.9	7.0	7.1	7.2	7.3	7.4	7.5
Treatment	UTC	UTC	UTC	FI	FI	FI	FI	FI	FI	FI	F	F	F	F	FI	FI
Seeding rate	113	119	112	148	148	147	150	150	150	150	150	150	150	150	150	150

Table 2.1: Optimal seed treatments and seeding rates for each pH

To investigate the effects of noise, we implement three different noise levels, $\sigma \in \{10, 30, 50\}$, which corresponds to respectively about 10%, 30%, and 50% of the maximum possible yield.

Scenario B (high-dimensional)

This environment is meant to serve as a more realistic scenario in which farmers make many decisions on input use and management practices. In particular, there are more continuous choices than the low-dimensional case. Scenario B is relevant not only in practice but also has attracted attention in the literature (Orlowski et al., 2016). Since there rarely exist datasets densely covering a high-dimensional continuous space, to create a suitable oracle, we make use of the Agricultural Production Systems sIMulator (APSIM), a highly advanced simulator of agricultural systems (Holzworth et al., 2014) being widely used for generating hypothetical datasets (Jin et al., 2018, 2019b, 2017a; Lobell et al., 2013, 2014, 2015).

We simulate a maize production system in Ames, Iowa, using the weather data from 2013. In the APSIM maize module, we identify six production factors that are suitable for our purpose:

- x_1 : sowing density (seeds/m^2)
- x_2 : sowing depth (mm)
- x_3 : row spacing (m)
- x_4 : N fertilizer amount before sowing (kg/ha)
- x_5 : N fertilizer amount at sowing (kg/ha)
- x_6 : N fertilizer amount for top dressing (kg/ha)

Thus, this oracle returns a yield (y) based on the levels of six factors $x = (x_1, \ldots, x_6)$

$$y = f(x) + \varepsilon.$$

Profit is calculated in a similar way:

$$\pi = py - c \cdot x.$$

The output price (p) is obtained from Duffy (2013), and the input costs (c) are obtained from Johanns (2019). We assume no cost for sowing depth and row spacing, which implies the cost vector $c = (c_1, 0, 0, c_4, c_5, c_6)$. It turns out that the maximum profit is \$1,288/ha, achieved by the management choices $x = (7, 120, 0.5, 0, 100, 20)$. Due to the six-dimensional space, it is impossible to directly plot the profit surface. We instead plot six pairs of interest, (x_1, x_5), (x_1, x_6), (x_2, x_3), (x_4, x_5), (x_4, x_6), and (x_5, x_6), while fixing the other four factors at the optimal levels. As seen below, the profit function takes a complex shape as a mixture of flat, convex and concave portions.

To investigate effects of noise, we implement three different noise levels, $\sigma \in \{500, 1000, 1500\}$, which corresponds, respectively, to about 5%, 10%, and 15% of the maximum possible yield. Notice that this is a much smaller variation than seen in scenario A. The reason is that, due to the higher dimensionality, the effect of noise is stronger and 15% of noise is large enough to make learning very difficult.

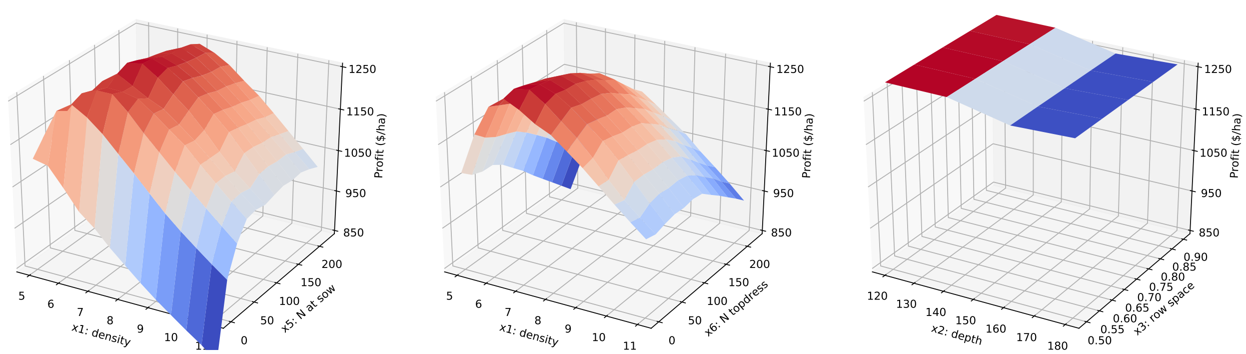

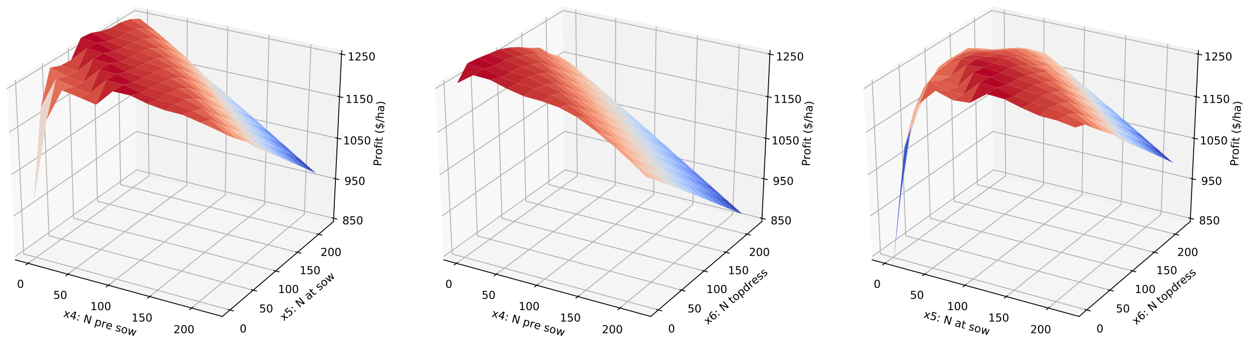

Figure 2.3: Selected profit surfaces of the high-dimensional oracle around the maximum

2.2.2 Experimental designs

Every experimental design either explicitly or implicitly assumes an accompanying statistical model that researchers estimate using the data collected by implementing the design (Montgomery, 2017). In agricultural experiments, it is common to use simple concave functions including quadratic (Bachmaier and Gandorfer, 2009; Meyer-Aurich et al., 2010b; Whelan et al., 2012), negative exponential (Edwards and Purcell, 2005; Gaspar et al., 2015), and piecewise linear (Ouedraogo and Brorsen, 2018; Park et al., 2018). In the BO literature, (adaptive) designs are called “acquisition functions,” and typically accompanied by a Gaussian process statistical model (Rasmussen and Williams, 2006).

Each experiment proceeds on a batch-sampling basis over years. Let M denote a batch size. Then, every year a researcher employs M plots and collect M samples. For scenario A, an optimal seeding rate for each pH is searched for over $x_1 \in [30, 150]$. For scenario B, an optimal level of the six factors is searched for over $x_1 \in [5.0, 11.0]$, $x_2 \in [120, 180]$, $x_3 \in [0.5, 0.9]$, $x_4 \in [0, 220]$, $x_5 \in [0, 220]$, and $x_6 \in [0, 220]$.

Factorial design

For scenario A, a factorial design consists of $\{45, 75, 105, 135\}$ in seeding rate and $\{6.0, 6.5, 7.0, 7.5\}$ in pH for each of three seed treatments, making up $M = 4 \times 4 \times 3 = 48$. For scenario B, we choose two levels for each factor, making up $M = 2^6 = 64$. Specifically, $x_1 \in \{7.0, 9.0\}$,

$x_2 \in \{140, 160\}$, $x_3 \in \{0.6, 0.8\}$, $x_4 \in \{60, 140\}$, $x_5 \in \{60, 140\}$, and $x_6 \in \{60, 140\}$. Intuitively, two levels for each factor may be too coarse to detect good optimal values over the continuous search space. However, with three levels, it would be $3^6 = 729$ and too large to be practical. This already indicates a fundamental problem of factorial design in a high-dimensional space, a problem dubbed the "curse of dimensionality" (Bellman, 2015). As is standard in the literature, the design in each scenario is fixed over the years. Finally, for an accompanying statistical model, we assume a quadratic model, which is in the case of $x = (x_1, x_2)$:

$$f(x) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_1^2 + \beta_4 x_1 x_2 + \beta_5 x_2^2,$$

and in the case of $x = (x_1, x_2, x_3, x_4, x_5, x_6)$:

$$f(x) = \sum_{k_1+k_2+\cdots+k_6 \leq 2} \beta_k \prod_{j=1}^{6} x_j^{k_j},$$

where $k_j \in \{0, 1, 2\}$ for all $j \in \{1, 2, \ldots, 6\}$, and $k = (k_1, k_2, k_3, k_4, k_5, k_6)$ denotes one of 28 indices for the coefficient β jointly defined by $k_1, k_2, \ldots, k_6$.

Bayesian optimization design

In contrast, the BO design adaptively chooses any combinations recommended by the algorithm. Each time after obtaining M samples, it updates the Gaussian process statistical models and reflects these updates on the next year's sampling recommendation. Hence, the adaptive experimental design. The algorithm is named "batch expected improvement," which is an extension of the standard expected improvement (EI) algorithm (Jones et al., 1998; Moćkus et al., 1978). The details of the algorithms for both scenario A and B are found in Appendix. The basic mathematics for and the specifications of Gaussian process are also found in Appendix.

While the standard EI algorithm is a strictly sequential algorithm designed to recommend only a single combination for the next sampling occasion, in agricultural experiments, we test multiple combinations at each occasion (typically, a year). Hence, we need to extend it so as to recommend multiple combinations at a time. To this end, we run the standard EI algorithm M times within a year by assuming $M-1$ hypothetical samples, as if they were sequentially observed. We adopt the idea from Ginsbourger et al. (2010) and assume each hypothetical sample to be equal to the lowest value observed so far (see line 7 in Algorithm 1). In other words, we assume within a year the algorithm makes a bad recommendation (remember the objective of the profit maximization—searching for a combination that returns the highest value). This heuristic allows the EI algorithm to do more exploration than exploitation; otherwise, the EI algorithm would tend to over-exploit suboptimal choices without sufficient exploration (Berk et al., 2018).

2.2.3 Performance metric

To compare alternatives, we need a performance metric. For this purpose, we use a notion of "regret" defined as a percentage of missed profit—a relative difference between the maximum possible profit and a profit computed using a statistical model of the environment estimated with the samples collected by following a design.

$$\text{Regret} = \frac{\text{True maximum profit} - \text{Estimated maximum profit}}{\text{True maximum profit}} \times 100.$$

For scenario A, since a regret can be calculated for each of 16 pHs, we use their average as the regret of a design.

2.2.4 Overall procedure

Our simulation study places the researcher in two distinct scenarios, each of which has three subscenarios characterized by different noise levels. In each subscenario, experiments last for T years. In each year, a design does the following:

1. Picks M testing combinations
2. Queries each combination to the oracle
3. Receives corresponding M yield observations and computes their profits
4. Adds them to the sample pool
5. Updates the statistical model for the underlying profit function

Note that step 5 is meaningful only for the BO design because the factorial design is static and independent of the statistical model. After T years, a regret is calculated as a performance metric of the design. To reduce the effect of the randomness involved in both yield noise and the algorithms themselves, we run multiple Monte Carlo simulations for each subscenario and report their sample mean and standard deviation as a final result.

2.3 Results

2.3.1 Scenario A

The study for scenario A was conducted with 1,000 Monte Carlo simulations. Besides the fact that the factorial design had $M = 48$, the BO version also used the same number for the batch size. The effects of noise are illustrated separately for the factorial and BO designs in the figures below. As expected, the greater the noise level, the more difficult it is to learn profit functions and optimize them. With a greater noise level, clearly, the regret in any time horizon T is higher. In addition, at noise level 50, the BO version hardly learns anything even after 10 years, but the factorial does no better, nonetheless.

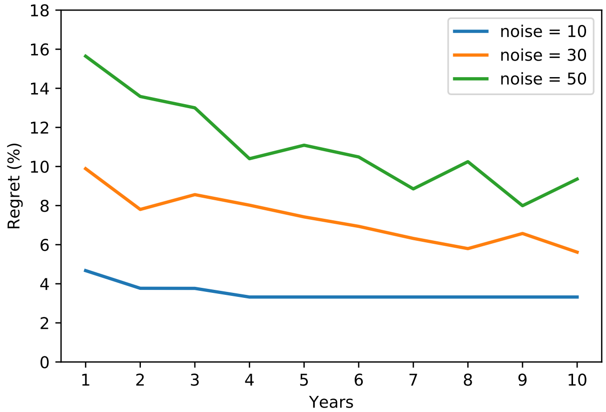

Figure 2.4: Effect of noise on learning (Factorial)

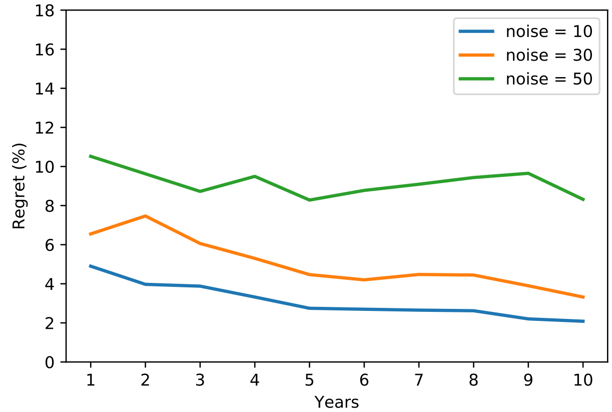

Figure 2.5: Effect of noise on learning (BO)

To show more detail, the following three figures plot the results separately for each noise level. Solid lines trace the mean regret at each time horizon (T), and shaded areas represent 1 standard deviation around the means. Overall, on average, the BO design demonstrates higher performance for most time horizons in all the subscenarios. However, the differences may not be seen as economically significant, and moreover, their distributions are largely overlapped, indicating even weaker appeal of the higher mean performances.

10% noise

Years	1	2	3	4	5	6	7	8	9	10
Factorial	6.8	5.4	4.8	4.1	3.8	4.0	3.6	3.5	3.5	3.4
(std)	(5.0)	(4.6)	(3.9)	(2.4)	(1.4)	(2.1)	(1.2)	(1.2)	(1.2)	(1.3)
BO	7.4	4.9	4.2	3.7	3.1	3.1	2.7	2.6	2.2	2.2
(std)	(6.0)	(2.3)	(1.6)	(1.4)	(1.4)	(1.7)	(1.5)	(1.5)	(1.0)	(1.2)
Difference	-0.6	0.5	0.6	0.4	0.7	0.9	0.9	0.9	1.3	1.2

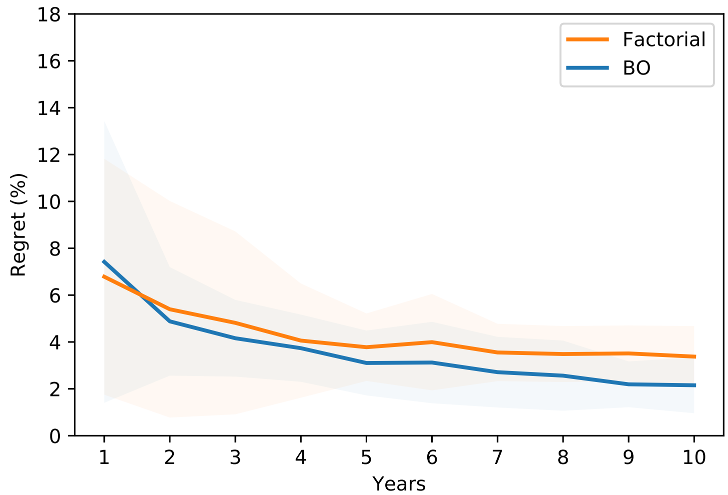

Figure 2.6: Regret under noise = 10%, M = 48

30% noise

Years	1	2	3	4	5	6	7	8	9	10
Factorial	11.9	9.6	9.7	8.9	8.2	8.1	7.6	6.8	7.2	6.7
(std)	(7.3)	(5.7)	(6.4)	(5.1)	(5.7)	(5.1)	(5.0)	(4.4)	(4.7)	(4.0)
BO	9.9	9.5	8.5	8.7	6.7	6.7	6.3	5.5	5.5	4.2
(std)	(7.0)	(6.0)	(6.8)	(7.7)	(6.4)	(6.4)	(6.4)	(3.5)	(4.7)	(2.7)
Difference	2.0	0.1	1.2	0.2	1.5	1.4	1.3	1.3	1.7	2.5

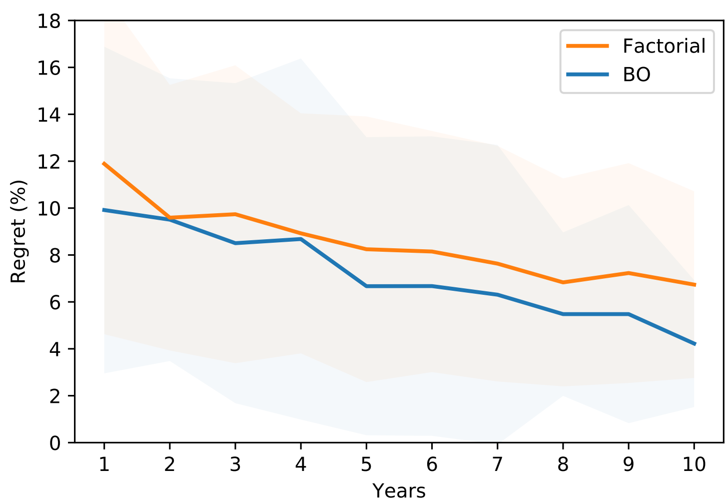

Figure 2.7: Regret under noise = 30%, M = 48

50% noise

Years	1	2	3	4	5	6	7	8	9	10
Factorial	15.7	14.5	13.0	11.9	10.8	11.3	10.5	10.2	9.3	10.6
(std)	(7.9)	(7.7)	(7.5)	(6.9)	(6.5)	(6.4)	(6.3)	(5.7)	(5.6)	(6.1)
BO	10.7	10.6	9.3	8.9	9.1	9.7	9.3	10.9	9.7	8.3
(std)	(3.6)	(5.4)	(5.0)	(2.9)	(3.7)	(4.2)	(2.0)	(4.5)	(2.5)	(1.7)
Difference	5.0	3.9	3.7	3.0	1.7	1.6	1.2	-0.7	-0.4	2.3

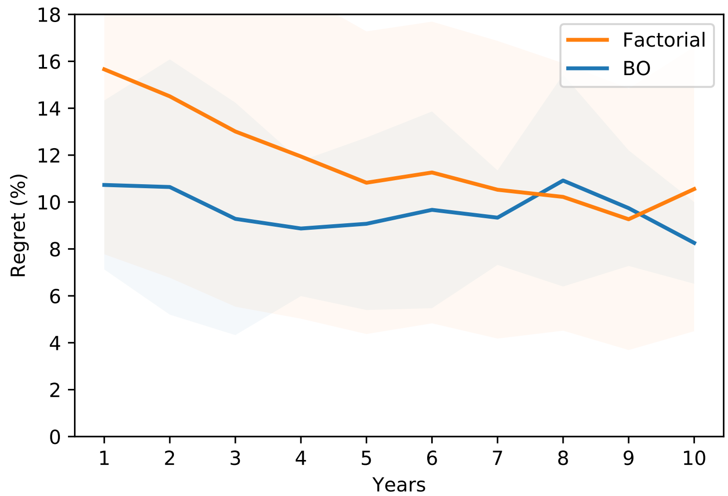

Figure 2.8: Regret under noise = 50%, M = 48

2.3.2 Scenario B

The study for scenario B was conducted with 1,000 Monte Carlo simulations. As opposed to $M = 64$ of the factorial design, the BO design here uses only $M = 30$ for the batch size. The reason for the smaller batch size is because it turned out $M = 30$ is sufficient to produce convincing results, and the computational costs grow increasingly with batch size. First, the effects of noise are illustrated separately for the factorial and BO designs in the figures below. Similar to scenario A, there is a pattern that the greater the noise level, the more difficult it is to learn profit functions and optimize them.

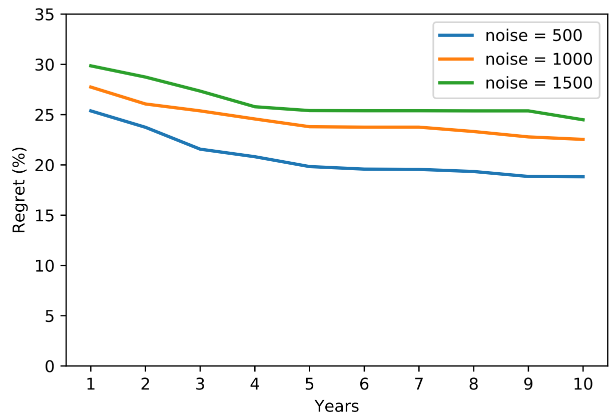

Figure 2.9: Effect of noise on learning (Factorial)

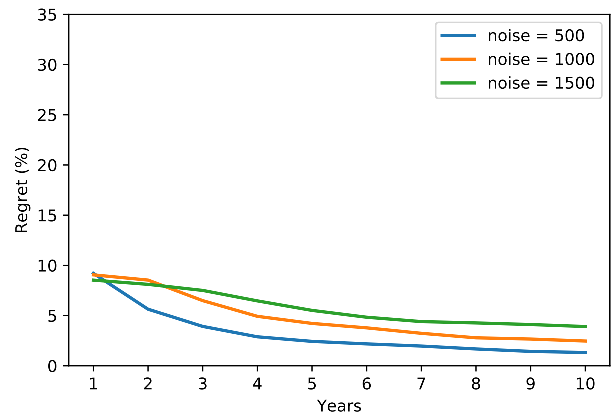

Figure 2.10: Effect of noise on learning (BO)

To show this in more detail, the following three figures plot the results separately for each noise level. Solid lines trace the mean regret at each time horizon (T), and shaded areas represent the corresponding 1 standard deviations around the means. It is evident that, on average, the BO design outperformed the factorial design by a wide margin, at least 12.9 percentage points and larger with longer time horizons. Besides, as indicated by the two largely separated shaded areas in each figure, the mean regrets of the BO design is also statistically distant from that of the factorial. Moreover, it is probably a more practical question to ask how often the BO design "beat" the factorial design. So, at the bottom of each table, we also show the percentage of simulation runs when the BO had lower regrets than the factorial at each time horizon. To be fair, we gave both designs the same random realization of noise at each simulation run (and of course different randomness across 1,000 runs). As implied by the large mean differences, the BO beat the factorial at least 83.3% of times and generally more as the horizon became longer.

5% noise

Years	1	2	3	4	5	6	7	8	9	10
Factorial	25.5	23.6	22.4	21.7	21.4	20.9	20.6	20.4	20.3	19.9
(std)	(9.3)	(7.7)	(6.2)	(5.7)	(5.1)	(4.8)	(4.4)	(4.1)	(4.0)	(3.9)
BO	12.6	9.6	6.9	4.8	3.7	2.9	2.5	2.1	1.8	1.7
(std)	(11.0)	(9.2)	(7.8)	(6.3)	(4.6)	(3.8)	(3.1)	(2.8)	(2.1)	(2.0)
Difference	12.9	14.0	15.5	16.9	17.7	18.0	18.1	18.3	18.5	18.2
%wins	83.3	87.4	92.5	95.8	97.9	98.2	99.0	99.3	99.6	99.8

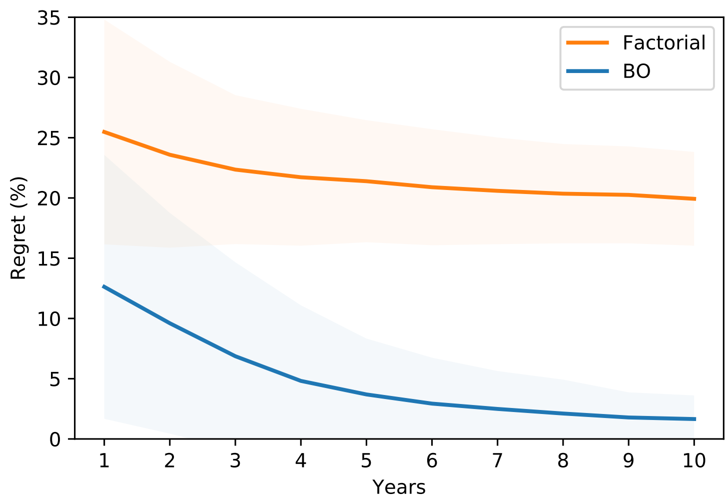

Figure 2.11: Regret under noise = 5%, M = 30, 64

10% noise

Years	1	2	3	4	5	6	7	8	9	10
Factorial	29.4	27.8	26.0	25.3	24.9	24.1	24.1	23.7	23.6	23.3
(std)	(12.2)	(11.0)	(9.7)	(9.1)	(8.4)	(8.1)	(8.0)	(7.2)	(6.8)	(6.8)
BO	11.3	11.9	10.3	8.2	6.4	5.4	4.7	3.8	3.5	3.1
(std)	(8.7)	(10.4)	(9.8)	(8.5)	(6.6)	(5.8)	(5.8)	(4.5)	(4.0)	(3.2)
Difference	18.1	15.9	15.7	17.1	18.5	18.7	19.4	19.9	20.1	20.2
%wins	89.0	88.4	89.4	91.3	95.2	96.9	96.3	98.5	98.8	99.0

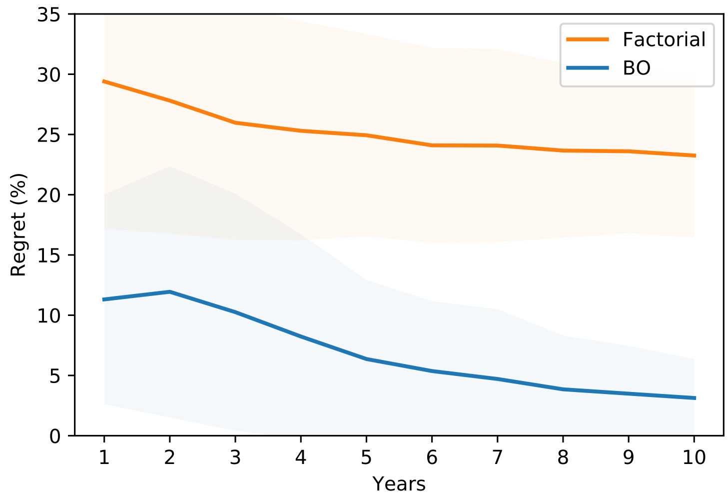

Figure 2.12: Regret under noise = 10%, M = 30, 64

15% noise

Years	1	2	3	4	5	6	7	8	9	10
Factorial	31.7	30.0	28.8	27.7	26.8	26.5	26.4	25.9	26.0	25.4
(std)	(13.1)	(12.5)	(11.9)	(10.9)	(10.4)	(10.2)	(10.1)	(9.8)	(9.2)	(8.9)
BO	10.6	10.5	10.5	9.6	8.6	7.8	6.6	5.8	5.5	5.2
(std)	(8.3)	(8.3)	(8.9)	(8.6)	(8.3)	(7.8)	(6.3)	(5.5)	(5.7)	(5.3)
Difference	21.1	19.5	18.3	18.1	18.2	18.7	19.8	20.1	20.5	20.2
%wins	92.8	92.2	90.2	90.0	92.1	93.2	95.6	95.9	96.6	96.8

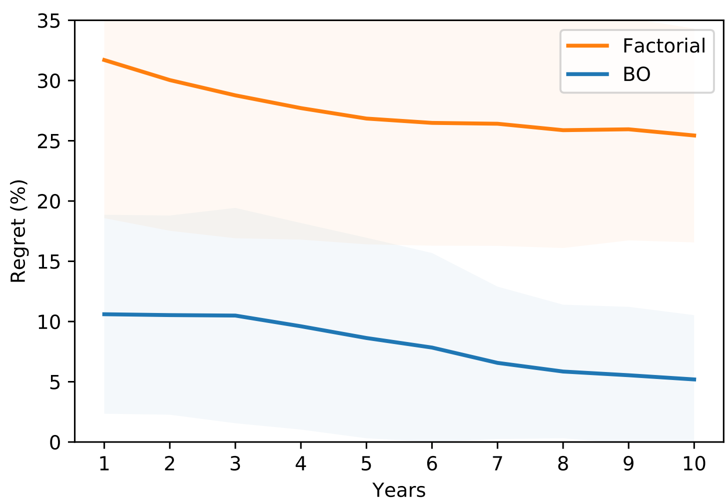

Figure 2.13: Regret under noise = 15%, M = 30, 64

2.4 Discussion

We created two distinct simulation environments and tested optimization performances of factorial designs and adaptive designs based on BO. Two environments were constructed to highlight conditions in which the BO design could offer the most benefits—high-dimensional environments involving many management and environmental variables. The common results in both designs were that the noisier the production systems, the more difficult it is to learn and optimize. These were intuitive results and reiterated the difficulty involved in optimizing agricultural systems, where a considerable amount of noise is typically present. In terms of relative performance, the BO designs outperformed the factorial designs on average, indicated by the blue curves located below the orange curves for most time horizons in Figures 6-8 and 11-13. Recall that one of the reasons for replication is to reduce the variability of estimator or standard error (Mead et al., 2002). However, as the similar sizes of the shaded areas in each figure indicated, the variability of the factorial design in both scenarios is by no means smaller than that of the BO design at large and even larger in many cases (e.g. Figures 8, 12, 13). So, the factorial design in this case actually sacrificed the quality (i.e., average regret) of the estimator without gaining anything. Besides the common characteristics of the results between the two scenarios, there are notable differences between scenario A and scenario B.

In scenario A, while the overall performance of the BO design is higher in most time horizons, the performance differences are small, ranging from -0.7 to 5.0 percentage points across all the subscenarios (Figures 6-8). As a result, these differences may not be seen as economically significant. Moreover, the distributions of the two designs at each time horizon are largely overlapped, indicating even weaker statistical appeal of the BO design's higher average performances. This was attributed to the low dimensionality of the scenario. Specifically, while there were technically three dimensions in scenario A—seed treatment, seeding rate, and pH—the seed treatment dimension was discrete containing only three elements $\{UTC, F, FI\}$. Therefore, it was still "easy" to cover the dimension by the factorial design, reducing the relative advantage of the BO version.

In contrast, in scenario B, the BO design overwhelmingly outperformed the factorial design by a wide margin, ranging from 12.9 to 21.1 percentage points. With the maximum possible profit of $1,288/ha, 20 points of difference in regret are equivalent to $258/ha, which is of considerable economic benefits. To examine this in more detail, consider a five-year experimental horizon. The differences in regret across three subscenarios are 17.7 (5% noise), 18.5 (10% noise), and 18.2 (15% noise), corresponding to $228/ha, $238/ha, and $234/ha respectively. Remember that the BO design used only 30 samples each year, less than half of the 64 samples used by the factorial design. Thus, the addition of experimental cost saving made the economic benefits of the BO version even more significant. Furthermore, as clearly seen under the 5% noise subscenario (Figure 11), the performance difference in some cases widened as the time horizon became longer, indicating that the BO design learned more than the factorial.

Despite the promising ability and performance of BO experimental designs, there are several clarifications and limitations to note for real-world implementations and future research.

- We do not claim that the constructed simulation environments are particularly accurate or realistic in every detail. Instead, what we tried to create are independent and impartial devices that allow us to systematically test and compare alternative experimental designs.
- We implicitly assumed that, within a subscenario characterized by an oracle and a noise level, the noise distribution is stationary over the time horizon. For example, given the strong influence of weather on agricultural yield, it may be necessary to reflect weather patterns over multiple years.
- We assumed homoskedasticity—a common noise level across all values of x within a subscenario. With appropriate data, we can test that. Then, if violated, we can estimate heteroskedastic errors and reflect them in simulations.

- Our BO technique is limited to one-shot optimization, in other words, that researchers choose all the factor levels at the beginning of the year and wait to see results until the end of the year. In agriculture, many management choices are made within the year. To handle such more realistic situations, we need to construct dynamic models in which feedback information and learning take place both within and across years.

- When variables were theoretically continuous (e.g., seeding rate), we assumed that it was possible to choose any arbitrary level and run the BO algorithm accordingly. In reality, however, choices of continuous variables are constrained for various practical reasons. For example, for a technical reason, we cannot change a fertilizer amount precisely by a fraction of a kilogram. Even though technically possible, it is economically infeasible to change row spacing every time an algorithm recommends a different value. Therefore, for real-world implementations, we need to modify those continuous factors into discrete ones.

2.5 Conclusions

The goal of our work is to introduce BO techniques to agricultural audiences so that they can apply them to their own problems. To this end, we have described an experimental design based on BO and demonstrated its potential by conducting simulation studies. In general, the benefits of BO will likely be significant when the dimensionality of management and environmental variables is high, and the nature of production factors is continuous. As we try to optimize more variables, the number of combinations to examine increases exponentially. Consequently, fixed factorial designs or any brute-force search strategy is less effective, relatively increasing the advantage of sophisticated designs such as those based on BO. In addition, the more variables involved in defining the yield response function, the higher the chances of interactions between them (Bullock and Bullock, 2000; Ruffo et al., 2006), leading to a complex shape of the profit function, which is difficult to optimize by using the traditional experimental designs. Thus, it is safer and likely beneficial for researchers to use more flexible statistical models (e.g., Gaussian process used in BO) and guard against potential model mispecification and, thereby, substantial loss. In summary, experimental designs based on BO address these issues and potentially create significant economic benefits.

Appendices

Algorithms for the BO designs

In terms of notation, T is the total number of years used for experimentation, $\mathcal{S}_\tau$ is a set of samples for seed treatment τ, GP_τ is a Gaussian process regression model for the profit response function of seed treatment τ, $\min_y\{\mathcal{S}_{\tau^*}\}$ implies the minimum observed yield for τ^*, and α_τ is an acquisition function defined based on GP_τ. A pair of braces $\{\ \}$ indicates a collection or set that contains elements, for example, $\{\alpha_\tau\}$ is a set of three acquisition functions: α_{UTC}, α_F, and α_{FI}. Notice that, in the first year, $M/3$ number of samples are randomly collected for each seed treatment and stored in $\mathcal{S}_\tau$.

Algorithm 1 Batch expected improvement for scenario A

require: T, M, $\{\mathcal{S}_\tau\}$, $\{GP_\tau\}$, $\{\alpha_\tau\}$
for $t \in \{2, 3, \ldots, T\}$ **do**
  $\mathcal{I}_\tau \leftarrow \{\ \}$ for all $\tau \in \{UTC, F, FI\}$
  $\widehat{GP}_\tau \leftarrow GP_\tau$ for all $\tau \in \{UTC, F, FI\}$
  for $m \in \{1, 2, \ldots, M\}$ **do**
    $(\tau^*, x^*) \leftarrow \text{argmax}_{\tau,x}\, \alpha_\tau(x)$
    $\mathcal{I}_{\tau^*} \leftarrow \mathcal{I}_{\tau^*} \cup \{(x^*, \underline{y})\}$ where $\underline{y} = \min_y\{\mathcal{S}_{\tau^*}\}$
    Update $\widehat{GP}_{\tau^*}$ with $\mathcal{S}_{\tau^*} \cup \mathcal{I}_{\tau^*}$
  for $\tau \in \{UTC, F, FI\}$ **do**
    for $(x, y) \in \mathcal{I}_\tau$ **do**
      $y \leftarrow Oracle(\tau, x)$
      $\mathcal{S}_\tau \leftarrow \mathcal{S}_\tau \cup \{(x, y)\}$
    Update GP_τ with $\mathcal{S}_\tau$
return $M \times T$ number of samples

Notice that, despite the larger factor space in scenario B, it is algorithmically simpler because of the lack of discrete choice of seed treatment that require separate GPs. Note also that the difference between the two algorithms is not because of ad hoc adaptation to the specific scenarios; instead, it is due simply to the structural difference of the two scenarios. Hence, our BO design is consistent and general.

Algorithm 2 Batch expected improvement for scenario B

require: T, M, $\mathcal{S}$, GP, α
for $t \in \{2, 3, \ldots, T\}$ **do**
  $\mathcal{I} \leftarrow \{\ \}$
  $\widehat{GP} \leftarrow GP$
  for $m \in \{1, 2, \ldots, M\}$ **do**
    $x^* \leftarrow \text{argmax}_x\, \alpha(x)$
    $\mathcal{I} \leftarrow \mathcal{I} \cup \{(x^*, \underline{y})\}$ where $\underline{y} = \min_y\{\mathcal{S}\}$
    Update $\widehat{GP}$ with $\mathcal{S} \cup \mathcal{I}$
  for $(x, y) \in \mathcal{I}$ **do**
    $y \leftarrow Oracle(x)$
    $\mathcal{S} \leftarrow \mathcal{S} \cup \{(x, y)\}$
  Update GP with $\mathcal{S}$
return $M \times T$ number of samples

Gaussian process

Gaussian process is a Bayesian nonparametric model, and its behavior is largely dependent on a kernel (Rasmussen and Williams, 2006). A kernel is a function that returns a similarity measure $k(x, x')$ between two points x and x'. We use the Matérn kernel—a popular class of isotropic stationary kernels.

$$k_\nu(x, x') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{d}{\rho} \right)^\nu B_\nu \left(\sqrt{2\nu} \frac{d}{\rho} \right),$$

where Γ is the gamma function, B_ν is the modified Bessel function of the second kind, and d is a metric often induced by the Euclidean norm, i.e. $d = \|x - x'\|$. The Matérn kernel is characterized by two parameters ν and ρ, which control, respectively, the smoothness and the scaling of distance. As standard in applied work, we do not estimate but rather handpick ν and write as Matérn$_\nu$ or $k_\nu(x, x')$. To simplify the notation, let r denote the scaled distance, $r = d/\rho$. An important property of the Matérn kernel is that when $\nu = p + 1/2, p \in \mathbb{N}$, it can be written as a product of an exponential and a polynomial of order p:

$$k_{p+1/2}(x, x') = \sigma^2 \exp\left(-\sqrt{2p+1}r \right) \frac{p!}{(2p)!} \sum_{i=0}^{p} \frac{(p+i)!}{i!(p-i)!} (2\sqrt{2p+1}r)^{p-i}.$$

Common choices of ν are $1/2, 3/2, 5/2$ and ∞, with each of which the kernel reduces to, respectively,

$$\begin{aligned}
k_{1/2}(x, x') &= \sigma^2 \exp(-r) \\
k_{3/2}(x, x') &= \sigma^2 \exp\left(-\sqrt{3}r \right) \left(1 + \sqrt{3}r \right) \\
k_{5/2}(x, x') &= \sigma^2 \exp\left(-\sqrt{5}r \right) \left(1 + \sqrt{5}r + \frac{5}{3}r^2 \right) \\
k_\infty(x, x') &= \lim_{\nu \to \infty} k_\nu(x, x') = \sigma^2 \exp\left(-\frac{1}{2}r^2 \right).
\end{aligned}$$

We use $\nu = 3/2$ for the BO design. Matérn$_\infty$ is also known as squared exponential kernel or radial basis function. The following figure plots $k_\nu(x, x')$ with $\sigma^2 = \rho = 1$ for $\nu \in \{1/2, 3/2, \infty\}$.

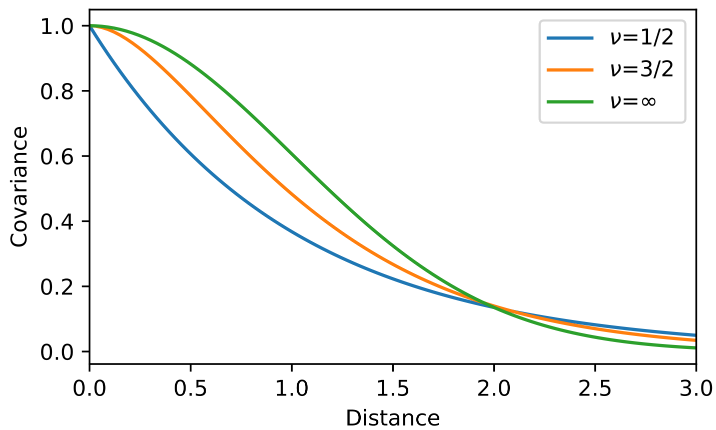

Figure 2.14: Matérn$_\nu$ kernels

Expected improvement acquisition function

In Bayesian optimization, an algorithm prescribes the next sampling point x based on how we value the mean and variance at x estimated by the accompanying GP. Specifically, the recommendation x_t for the next round t is determined by maximizing an acquisition function $\alpha(x|D_{t-1})$:

$$x_t = \underset{x}{\operatorname{argmax}}\, \alpha(x|D_{t-1}),$$

where D_{t-1} is the data used to fit the GP at round $t-1$. The acquisition function is a reflection of the underlying utility of the next sample or our preference in selecting the next sampling point. It is heuristic and designed to trade off exploration of the search space and exploitation of the current promising areas. There are a number of acquisitions functions proposed in the literature. One of the popular acquisition functions is called expected improvement, which is constructed based on the following intuitive idea. Let y^* be the maximum value observed up until round $t-1$, i.e. $y^* = \max\{y_1, \ldots, y_{t-1}\}$. Then, we may define "improvement" at point x at round t to be

$$\max\{0, GP(x) - y^*\},$$

which is random as $GP(x)$ is a random function. Thus, the expected improvement acquisition function is defined to be:

$$\alpha_{EI}(x|D_{t-1}) = \mathbb{E}[\max\{0, GP(x) - y^*\}|D_{t-1}].$$

When using Gaussian process, at each point x in the domain, we have $GP(x) \sim \mathcal{N}(\mu(x), \sigma(x))$, which allows the expected improvement to have a closed form (Jones et al., 1998; Moćkus et al., 1978):

$$\alpha_{EI}(x|D_{t-1}) = \begin{cases} (\mu(x) - y^*)\Phi\left(\frac{\mu(x)-y^*}{\sigma(x)}\right) + \sigma(x)\phi\left(\frac{\mu(x)-y^*}{\sigma(x)}\right) & \text{if } \sigma(x) > 0 \\ 0 & \text{if } \sigma(x) = 0 \end{cases},$$

where Φ is the standard normal cumulative distribution function and ϕ is the standard normal probability density function.

Chapter 3

Machine learning for optimizing complex site-specific management

Yuji Saikai
Vivak Patel
Paul D. Mitchell

3.1 Introduction

Modern agriculture faces some of the most challenging problems of the 21st century, including food security, farm profitability and environmental sustainability, all of which require increasing agricultural productivity. To increase the productivity, it is crucial to exploit advanced farming (together known as *precision agriculture*) technologies such as yield monitors, remote sensing, and variable rate application, from which site-specific management (SSM) emerges as an effective management strategy (Auernhammer, 2001; Bongiovanni and Lowenberg-DeBoer, 2004; Cassman, 1999; Gebbers and Adamchuk, 2010). This is because SSM can optimize a production system at the subfield level, which amounts to finer-scale optimization than at the field level. Therefore, for both individual profitability and collective societal benefits, SSM has been advocated to farmers over the past two decades (Cook and

Bramley, 1998).

Despite the potential benefits for farmers, the adoption of SSM has been slower than expected (Bramley, 2009; Schimmelpfennig, 2016), which is attributed to a lack of relative advantage over the current management strategies (Pathak et al., 2019), particularly a lack of greater profitability (Castle et al., 2016; Gandorfer and Meyer-Aurich, 2017). In principle, those advantages of SSM can be realized by exploiting advanced technologies (e.g., adjusting fertilizer application rates across the field according to varying soil conditions inferred from satellite imagery). In reality, however, these technologies are quite sophisticated and difficult to exploit. Indeed, many farmers express their concerns about the complexity of these technologies (Bramley and Ouzman, 2019) and have yet to be convinced of the value of SSM (Antle, 2019; Antle et al., 2017b; Leonard et al., 2017; Lowenberg-DeBoer and Erickson, 2019). This lack of actionable procedures or decision support systems for SSM is noted as a serious problem in the literature (Lindblom et al., 2017).

The vast majority of SSM research investigates only a handful of management types in particular research environments. Commonly explored types include nitrogen fertilizer (Anselin et al., 2004; Boyer et al., 2011; Jin et al., 2017b; Karatay and Meyer-Aurich, 2019; Lo et al., 2019; Pahlmann et al., 2017; Pannell et al., 2019; Thöle et al., 2013), irrigation water (Cid-Garcia et al., 2014; Haghverdi et al., 2016), and sowing density (He et al., 2019a,b). Besides the fact that other types of management (e.g., tillage, spraying, and harvest) are also important, in practice, what really makes farm management difficult is the total number of decisions farmers have to make. Note that the total number of decisions increases exponentially as types of management increases, because each type of management typically involves decisions on its amount, frequency, and timing. Since crop yield and hence the overall effectiveness of SSM are determined by the totality of these management decisions and environmental factors (Bullock and Bullock, 2000; Bullock et al., 2002; Ruffo et al., 2006), independent studies of few management types involving a small number of decisions provide only partial knowledge and can be misleading as complex systems typically involve significant nonlinearity (Altieri, 2018; Gliessman, 1990).

Besides the small number of management decisions investigated in the existing research, farmers are also concerned about the generalizability of results obtained under the particular research environments, because such results may not be representative of their unique environments (Cook et al., 2013). Each farm is unique and operates with different resources (e.g., machinery) in different environments (e.g., soil and weather). Consequently, results from studying "representative" cases are useful for only farmers who face similar managerial and environmental conditions (Gomez and Gomez, 1984; Panten et al., 2010). This problem is exacerbated in modern agriculture where the high dimensionality created by advanced sensor technologies makes it even more difficult for selected cases to be representative, the problem called "the curse of dimensionality" in mathematics (Bellman, 2015). As an example of how particular a typical study setting can be, Lo et al. (2019) have studied SSM of nitrogen fertilizer at a university research site in Nebraska that has been "under annual summer corn or soybean production without any tillage and any stover removal" with irrigation water applied through "a center pivot with sprayhead sprinklers positioned every other interrow at a height of 0.6 m above ground" using "GrowSmart Precision Variable Rate Irrigation system". The problem is not how particular existing studies are, as they serve different research purposes. Rather, the problem is a lack of studies that generalize the effects of environmental factors and jointly examine many management variables for the purpose of optimizing SSM in practice.

To address these issues (few management variables only in particular environments) and offer an actionable procedure for SSM, we develop a machine learning algorithm that enables each farmer to efficiently learn unique SSM through on-farm experiments implemented via existing advanced farming technologies. We emphasize the significance of on-farm experimentation (Cook et al., 2013; Griffin, 2018; Meyer-Aurich et al., 2010a; Piepho et al., 2011) to deal with both issues. First, on-farm experimentation allows farmers to adaptively design experiments and efficiently navigate a high-dimensional variable space (Saikai et al., 2020). Second, on-farm experimentation inherently allows each farmer to collect data from his/her unique environment and circumvents the representativeness issue. Particularly for SSM, use of field-scale experimentation is important to deal with spatial, infield variability (Bramley et al., 2011; Panten et al., 2010; Pringle et al., 2004a,b). Notice that our approach is different from the most common type of machine learning (i.e., supervised learning), which constructs an empirical model that assumes specific variables and estimates their associated parameters using a large observational dataset. The constructed model is supposed to be representative of users' production systems and therefore capable of indicating their optimal choices. However, owing to the nature of observational data, such a dataset likely contains insufficient variation in the high-dimensional management space because observed management choices mostly follow the standard recommendations from the existing low-dimensional studies (Bullock et al., 2019). In other words, purposeful experimentation is crucial to include unconventional management choices and discover unexpected optimal choices. Consequently, it is highly unlikely for such supervised learning models to be able to indicate optimal choices. Our approach is, instead, a machine learning algorithm that allows each farmer to construct a unique model. It is sufficiently versatile to be used for learning optimal SSM through on-farm experiments in a wide range of farming scenarios. The algorithm is based on Bayesian optimization (Saikai et al., 2020), and capable of handling an arbitrary number of management and environmental variables and adapt to unexpected interaction effects. Moreover, if the farmer has historical data, large or small, it can be incorporated as prior knowledge.

In the remainder of the paper, we first mathematically formulate the farmer's problem, and then we describe our machine learning algorithm as a solution to the problem. We test the algorithm's performance and versatility in two simulated environments, with either medium- or high-complexity. To highlight the generality of the approach, we report results on a per-hectare basis so that they can be easily scaled. Results suggest that complex SSM can be learned very efficiently through on-farm experiments within a few years, and it can be more profitable and more environmentally sustainable than uniform management.

3.2 Materials and methods

3.2.1 Farmer's problem

Imagine a farmer who has access to precision agriculture equipment for SSM but currently does not implement it due to a lack of knowledge, a typical story about SSM (Lindblom et al., 2017). Suppose that the farmer is now interested in learning optimal SSM through on-farm experiments. To conduct field-scale experiments, the farmer divides an entire field into a grid of sites of equal size according to the capacity of the existing variable rate technologies, yield monitors, and other sensors that monitor soil properties. Specifically, it is assumed that there is no spatial misalignment among these technologies, all of which has the same spatial resolution. As a result, this resolution defines a site as the observational unit (Bullock et al., 2019), and in each site the farmer collects a pair of data $(\boldsymbol{x}, y)$—applying a vector of management $\boldsymbol{x}$ and observing the corresponding scalar yield y. Let M denote the total number of sites. Site $s \in \{1, 2, \ldots, M\}$ is characterized by a state vector $\boldsymbol{z}_s$, to which management $\boldsymbol{x}_s$ is applied. Then, a site-specific profit function for site s is

$$\pi(\boldsymbol{x}_s|\boldsymbol{z}_s) = py(\boldsymbol{x}_s|\boldsymbol{z}_s) - \boldsymbol{c} \cdot \boldsymbol{x}_s,$$

where $y(\boldsymbol{x}_s|\boldsymbol{z}_s)$ is a site-specific yield function for site s, p is an output price, and $\boldsymbol{c}$ is a vector of prices for $\boldsymbol{x}_s$. Notice that this specification is technically a partial profit, as we subtract only the costs for modeled management $\boldsymbol{x}_s$. Nonetheless, it is immaterial because our analysis is based on the difference between profits using SSM and uniform management.

For conventional low-dimensional yield functions, it is common to use simple concave functions such as quadratic (Bachmaier and Gandorfer, 2009; Meyer-Aurich et al., 2010b; Whelan et al., 2012), negative exponential (Edwards and Purcell, 2005; Gaspar et al., 2015), and piecewise linear (Ouedraogo and Brorsen, 2018; Park et al., 2018). These simple functions may serve well for answering isolated questions about the optimality of a single management variable under homogeneous conditions. However, in high-dimensional cropping systems, the yield function is a fundamental source of the challenge because its uncertainty increases with the number of variables entering the site-specific yield function $y(\,\cdot\,|\,\cdot\,)$.

Note that each site need not be recognized as distinct or, equivalently, each $\boldsymbol{z}_s$ need not be distinct. A simple consequence of this assumption is that adjacent sites $\{s_1, s_2, \dots\}$ may have the same values $\boldsymbol{z}_{s_1} = \boldsymbol{z}_{s_2} = \dots$ and form a homogeneous "zone", which receives the same management. The research literature and farmer practice commonly use this type of zone delineation for SSM (Albornoz et al., 2018; Cid-Garcia et al., 2014; Fraisse et al., 2001; Fridgen et al., 2004; Gili et al., 2017; Leroux and Tisseyre, 2019; Li et al., 2007; Scudiero et al., 2013). Our formulation is more general and contains zone delineation as a special case.

Having each site-specific profit defined, field-level profit is simply the sum of the site-specific profits:

$$\sum_{s=1}^{M} \pi(\boldsymbol{x}_s|\boldsymbol{z}_s) = \sum_{s=1}^{M} py(\boldsymbol{x}_s|\boldsymbol{z}_s) - \boldsymbol{c} \cdot \boldsymbol{x}_s.$$

The farmer's objective is to learn optimal SSM $\boldsymbol{x}_s^*$ for all $s \in \{1, 2, \dots, M\}$

$$(\boldsymbol{x}_1^*, \dots, \boldsymbol{x}_M^*) = \underset{\boldsymbol{x}_1,\dots,\boldsymbol{x}_M}{\operatorname{argmax}} \sum_{s=1}^{M} \pi(\boldsymbol{x}_s|\boldsymbol{z}_s).$$

In contrast, under uniform management, a single management $\boldsymbol{x}$ is applied to every site s, so that field-level profit function is

$$\sum_{s=1}^{M} \pi(\boldsymbol{x}|\boldsymbol{z}_s) = \sum_{s=1}^{M} py(\boldsymbol{x}|\boldsymbol{z}_s) - \boldsymbol{c} \cdot \boldsymbol{x},$$

and optimal uniform management $\bar{\boldsymbol{x}}^*$ is

$$\bar{\boldsymbol{x}}^* = \underset{\boldsymbol{x}}{\operatorname{argmax}} \sum_{s=1}^{M} \pi(\boldsymbol{x}|\boldsymbol{z}_s).$$

3.2.2 Solution algorithm

We construct an algorithm based on Bayesian optimization (BO) (Brochu et al., 2010; Shahriari et al., 2016), a class of numerical optimization techniques to find the global optimum of an unknown objective function. As with many other numerical optimization techniques, BO navigates the search space by examining one point at a time until it locates an acceptable point and halts. Since BO tries to optimize an unknown function, it needs a surrogate model to guide its search. A Gaussian process (GP) statistical model is the standard choice in the literature.

BO has two features that make it suitable for agricultural experiments (Saikai et al., 2020). First, GP as a nonparametric Bayesian model is sufficiently flexible so that it can adapt to cases in which the objective function takes a complex shape. In high-dimensional SSM, this complexity will likely happen due to strong interactions among the many variables involved. Second, BO is in general known for its sample efficiency, which means that BO can locate a good enough point with relatively a small number of examinations. Since agricultural experiments take time before obtaining results, typically a year, sample efficiency is a strongly desirable feature.

While the basic BO sequentially processes one sample at a time, we modify it using the "batch BO" proposed by Saikai et al. (2020) in order to process M samples at a time. Each year, the algorithm proceeds as follows:

1. Prescribe $\boldsymbol{x}_s$ for each site s by maximizing the acquisition function $\alpha(\boldsymbol{x}|\boldsymbol{z}_s)$
2. Observe a yield $y(\boldsymbol{x}_s|\boldsymbol{z}_s)$ for each s
3. Compute the corresponding $\pi(\boldsymbol{x}_s|\boldsymbol{z}_s)$ for each s

4. Update the GP with $\{(\boldsymbol{x}_s, \boldsymbol{z}_s, \pi_s)\}_{s=1}^{M}$ and the samples from the preceding years.

After completing the planned number of years of experiments, a candidate for $\boldsymbol{x}_s^*$ for each s can be obtained by maximizing the mean function of the learned GP with fixed $\boldsymbol{z}_s$. Below is the complete algorithm.

Algorithm 3 Batch Bayesian optimization for site-specific management

1: **require:** T, M, $\mathcal{S}$, GP, α
2: **for** $t \in \{1, 2, \ldots, T\}$ **do**
3: $\quad \mathcal{I} \leftarrow \{\ \}$
4: $\quad \widehat{GP} \leftarrow GP$
5: $\quad$ **for** $s \in \{1, 2, \ldots, M\}$ **do**
6: $\quad\quad \boldsymbol{x}_s \leftarrow \text{argmax}_{\boldsymbol{x}}\, \alpha(\boldsymbol{x}|\boldsymbol{z}_s)$
7: $\quad\quad \mathcal{I} \leftarrow \mathcal{I} \cup \{(\boldsymbol{x}_s, \boldsymbol{z}_s, \underline{\pi})\}$ where $\underline{\pi} = \min\{\mathcal{S}_\pi\}$
8: $\quad\quad$ Update $\widehat{GP}$ with $\mathcal{S} \cup \mathcal{I}$
9: $\quad$ **for** $s \in \{1, 2, \ldots, M\}$ **do**
10: $\quad\quad y_s \leftarrow Oracle(\boldsymbol{x}_s|\boldsymbol{z}_s)$
11: $\quad\quad \pi_s \leftarrow py_s - \mathbf{c} \cdot \boldsymbol{x}_s$
12: $\quad\quad \mathcal{S} \leftarrow \mathcal{S} \cup \{(\boldsymbol{x}_s, \boldsymbol{z}_s, \pi_s)\}$
13: $\quad$ Update GP with $\mathcal{S}$
14: **return** $M \times T$ number of samples

In terms of notation, T is the total number of years used for experimentation, $\mathcal{S}$ is a set of samples, $\min\{\mathcal{S}_\pi\}$ implies the minimum realized profit, $\alpha(\cdot|\boldsymbol{z}_s)$ is the acquisition function for site s defined based on the Gaussian process GP, and $Oracle(\boldsymbol{x}|\boldsymbol{z})$ returns an observed yield when $\boldsymbol{x}$ is applied to a site characterized by $\boldsymbol{z}$. As in Saikai et al. (2020), for the acquisition function, we use the expected improvement (Jones et al., 1998; Moćkus et al., 1978). Also, for the oracle function, we use a crop simulator as described in the next section. Notice that in Line 5-8 the interim $\widehat{GP}$ is kept updated with hypothetical observation ($\underline{\pi}$) so that we can prescribe M sampling points $\{\boldsymbol{x}_s\}_{s=1}^{M}$ without any real sample. As a small detail, in Line 5, site s is chosen in a random order to avoid a systematic bias arising from how sites are numbered. Another detail is that, when updating GPs, we fit the hyperparameters of the GP only to observed data (Line 13) and not to hypothetical data (Line 8).

3.2.3 Simulation experiments

To construct simulation environments, we make use of the Agricultural Production Systems sIMulator (APSIM), an advanced simulator of cropping systems (Holzworth et al., 2014) widely used for various purposes, including generating synthetic datasets (Burke and Lobell, 2017; Jin et al., 2018, 2019b, 2017a; Lobell et al., 2013, 2014, 2015). In each environment, we run the algorithm to learn optimal SSM over T years and compare the profit resulting from implementing the learned SSM against the benchmark profit resulting from uniform management. This analysis assumes that uniform management follows university extension recommendations.

Depending on specific scenarios, we can customize the algorithm in many ways. An interesting modification is to incorporate observational data collected prior to beginning on-farm experimentation, the case for many farmers. When using uniform management as a benchmark for comparison, a natural dataset incorporated is the data from implementing uniform management, as it represents the existing knowledge. Here, we initialize the GP embedded in the algorithm as follows: let $\{(\bar{\boldsymbol{x}}, \boldsymbol{z}_s, \bar{\pi}_s)\}_{s=1}^{M}$ be a set of the uniform management ($\bar{\boldsymbol{x}}$), site characteristics ($\boldsymbol{z}_s$), and the corresponding profits ($\bar{\pi}_s$) and then, we fit the GP to these M data points before the algorithm starts a learning process. Note that the use of uniform management for both benchmark and prior knowledge is a useful case for illustration. In practical applications, farmers may use any benchmark management of interest (either uniform or not) and any existing dataset. Finally, since the algorithm itself involves some randomness, we conduct Monte Carlo simulations and present averaged results over 100 Monte Carlo samples.

Though the algorithm's applicability is by no means restricted to the scenarios described in this section, we rely on the APSIM simulator and construct illustrative test beds within its capability. We simulate a maize production system in Ames, Iowa using the daily weather data for 2013, the most recent annual dataset available in APSIM. In terms of management variables, we follow Saikai et al. (2020) and identify six variables $\boldsymbol{x} = (x^1, \ldots, x^6)$ in the

APSIM maize module:

- x^1 : Sowing density (seeds/m^2)
- x^2 : Sowing depth (mm)
- x^3 : Row spacing (m)
- x^4 : Amount of N fertilizer applied before sowing (kg/ha)
- x^5 : Amount of N fertilizer applied at sowing (kg/ha)
- x^6 : Amount of N fertilizer applied for top dressing (kg/ha)

Based on Iowa State University extension recommendations, we specify uniform management ($\bar{\boldsymbol{x}}$) as:

$$(\bar{x}^1, \bar{x}^2, \bar{x}^3, \bar{x}^4, \bar{x}^5, \bar{x}^6) = (8, 50, 0.76, 67, 67, 67).$$

$\bar{x}^1 = 8$, $\bar{x}^2 = 50$, and $\bar{x}^3 = 0.76$ follow from Elmore (2013) and Farnham (2001). The recommended total nitrogen amount is identified by using the Corn Nitrogen Rate Calculator (Sawyer, 2019), which gives $\bar{x}^4 + \bar{x}^5 + \bar{x}^6 = 201$. We evenly split it into $\bar{x}^4 = \bar{x}^5 = \bar{x}^6 = 67$. Finally, for calculating profits, the output price is $p = \$0.177$/kg (Duffy, 2013), and input costs are $c^1 = \$3.64/1000$ seeds and $c^4 = c^5 = c^6 = \$1.29$/kg (Johanns, 2019). We assume no cost for sowing depth and row spacing, which implies the cost vector $\boldsymbol{c} = (c^1, 0, 0, c^4, c^5, c^6)$.

Based on these same sources, when the algorithm searches for the optimal management, the search space is restricted to the following:

- $x^1 \in [6.0, 10.0]$ (seeds/m^2)
- $x^2 \in [25, 150]$ (mm)
- $x^3 \in [0.4, 1.0]$ (m)
- $x^4, x^5, x^6 \in [0, 200]$ (kg/ha)

Finally, data points resulting from the benchmark uniform management $\{(\bar{\boldsymbol{x}}, \boldsymbol{z}_s, \bar{\pi}_s)\}_{s=1}^{M}$ are the only existing dataset incorporated prior to beginning on-farm experimentation. Since initial uniform management provides no variation in $\boldsymbol{x}$, to build up smoothly, in the first year, the algorithm randomly chooses $\boldsymbol{x}$ for each s from the range defined by $\pm 50\%$ of the uniform management.

Scenario A (medium complexity)

This scenario assumes that a square field is divided into a grid of 25 sites ($M = 25$). All the sites are distinct, each characterized by a state vector $\boldsymbol{z}_s = (z_s^1, z_s^2)$ where z_s^1 is plant available water capacity (mm) and z_s^2 is organic carbon (%). We set $z_s^1 \in \{231, 259, 288, 317, 346\}$ and $z_s^2 \in \{2.56, 2.88, 3.2, 3.52, 3.84\}$ (0%, $\pm$10% and $\pm$20% from the default values in the APSIM soil module used).

231, 2.56	259, 2.56	288, 2.56	317, 2.56	346, 2.56
231, 2.88	259, 2.88	288, 2.88	317, 2.88	346, 2.88
231, 3.2	259, 3.2	288, 3.2	317, 3.2	346, 3.2
231, 3.52	259, 3.52	288, 3.52	317, 3.52	346, 3.52
231, 3.84	259, 3.84	288, 3.84	317, 3.84	346, 3.84

Figure 3.1: Simulated maize field divided into a grid of 25 distinct sites. The first number in each grid indicates plant available water capacity (mm) and the second indicates organic carbon (%) at that site.

Scenario B (high complexity)

This scenario imagines that the farmer possesses more precise equipment that can operate at a more granular scale, and so divides a field into more granular sites: $16 \times 9 = 144$ sites. We also assume that the farmer has conducted more exhaustive soil tests, measuring four state variables (z^1, z^2, z^3, z^4) in each site:

- z^1: plant available water capacity (mm)
- z^2: organic carbon (%)

- z^3: initial nitrate-N (kg/ha)
- z^4: initial ammonium-N (kg/ha)

The addition of z^3 and z^4 is because of their significance for nitrogen management (Camberato and Nielsen, 2017) and availability in APSIM. Notice that nitrogen is also supplied by soil organic matter (z^1) through N-mineralization, creating stronger interactions among management and environmental variables (Sawyer, 2008), so that SSM in scenario B is more complex than in scenario A. We generate a state vector for each site in a random but spatially correlated fashion, namely, a random walk (see Appendices for details). Below are summary statistics of the generated state vectors for the 144 sites.

- z^1: mean = 296, std = 32, min = 199, max = 365
- z^2: mean = 3.19, std = 0.31, min = 2.59, max = 3.90
- z^3: mean = 9.1, std = 0.98, min = 7.0, max = 11.5
- z^4: mean = 10.6, std = 1.6, min = 7.7, max = 14.1

Instead of reporting four numbers at each site, we illustrate the in-field variability using a yield map arising from applying the uniform management to the generated field. Since each site receives the same management, the variability in yield indicates the variability in the underlying growing conditions.

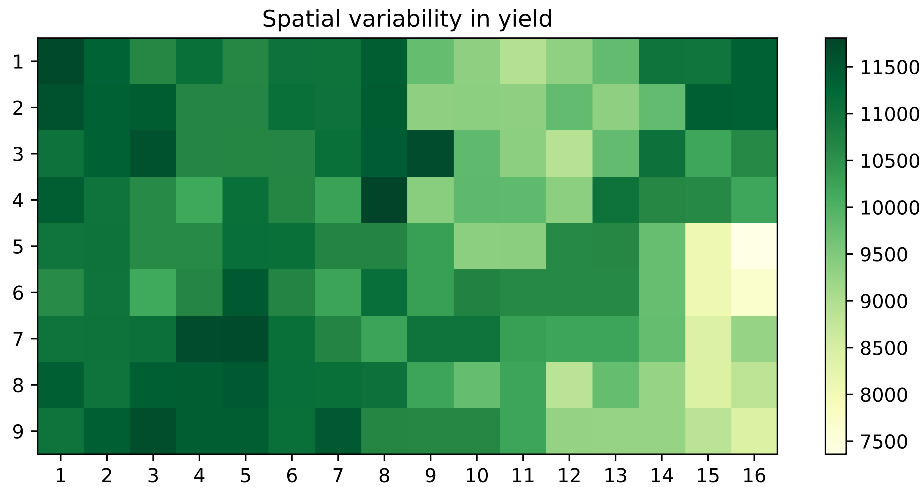

Figure 3.2: Yield map (kg/ha) resulting from the uniform management in scenario B. Axis ticks indicate site coordinates.

3.3 Results

3.3.1 Scenario A (medium complexity)

Table 3.1 and Figure 3.3 report field-level profits ($/ha) from implementing the SSM learned after conducting experiments for T years. Specifically, the value for each $T \in \{1, 2, \ldots, 10\}$ is the profit if the farmer terminates the experiments after T years and implements the learned SSM without further improvement.

Years (T)	1	2	3	4	5	6	7	8	9	10
Learned	1103	1237	1266	1274	1277	1280	1283	1284	1285	1285
Uniform	1234	1234	1234	1234	1234	1234	1234	1234	1234	1234
Difference	-131	3	32	40	43	46	49	50	51	51
										($/ha)

Table 3.1: Field-level profits ($/ha) from implementing the learned SSM and uniform management in scenario A.

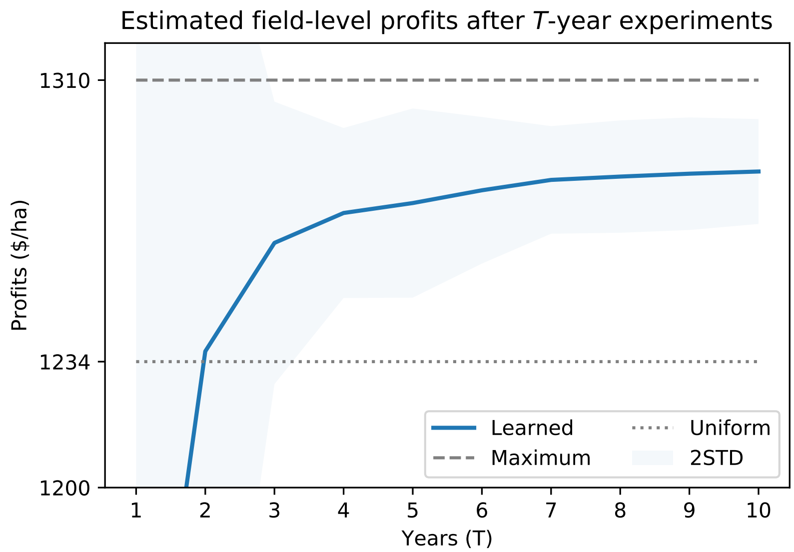

Figure 3.3: Learning curve of the algorithm with profits plotted against years of experiments in scenario A. The shaded areas indicate two standard deviations around each mean field-level profit over 100 Monte Carlo samples.

The shaded areas indicate two standard deviations around each mean field-level profit over 100 Monte Carlo samples. The dashed line indicates the maximum possible profit obtained by implementing the optimal SSM, while the dotted line indicates the profit from uniform management.

Since a field-level profit is the sum of the site-specific profits, we next examine profit at each site. Figure 3.4 illustrates profits from the learned SSM after five years and uniform management. $T = 5$ is chosen because the learning mostly levels off and the deviation from the mean prediction becomes small after four or five years.

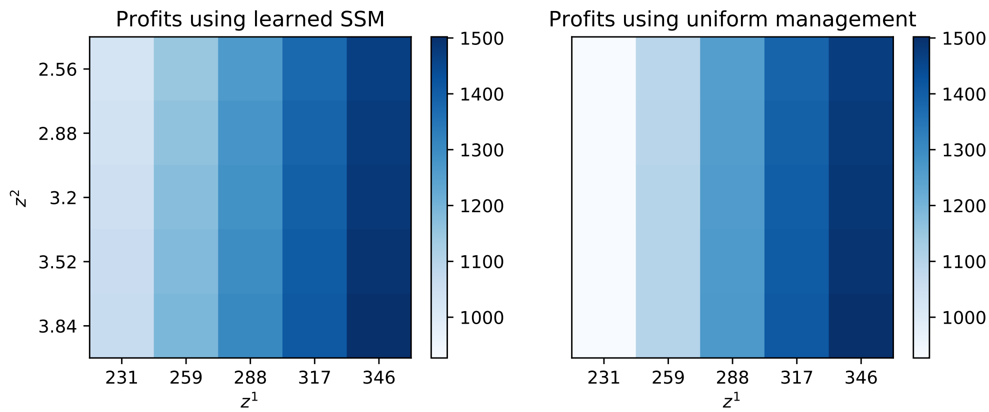

Figure 3.4: Site-specific profits ($/ha) for the learned SSM after 5 years and uniform management in scenario A

As indicated in the panel for uniform management, plant available water capacity (z^1) has a much stronger influence on profits than organic carbon (z^2). However, as both variables increase, the site becomes more fertile (though hard to see for organic carbon z^2). Since panels for both management systems look quite similar, to highlight their difference, Figure 3.5 illustrates the difference at each site.

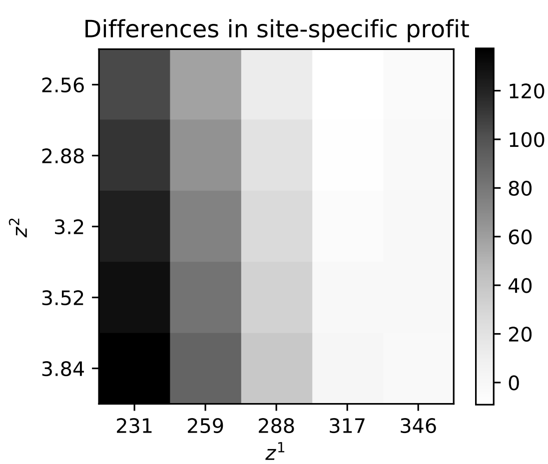

Figure 3.5: Differences in site-specific profit ($/ha) for SSM versus uniform management in scenario A. The maximum difference is $138/ha at site (231,3.84) and the minimum difference is $-9.4/ha at site (317,2.56).

Finally, Figure 3.6 reports the learned SSM ($x^1, \ldots, x^6$) after 5 years.

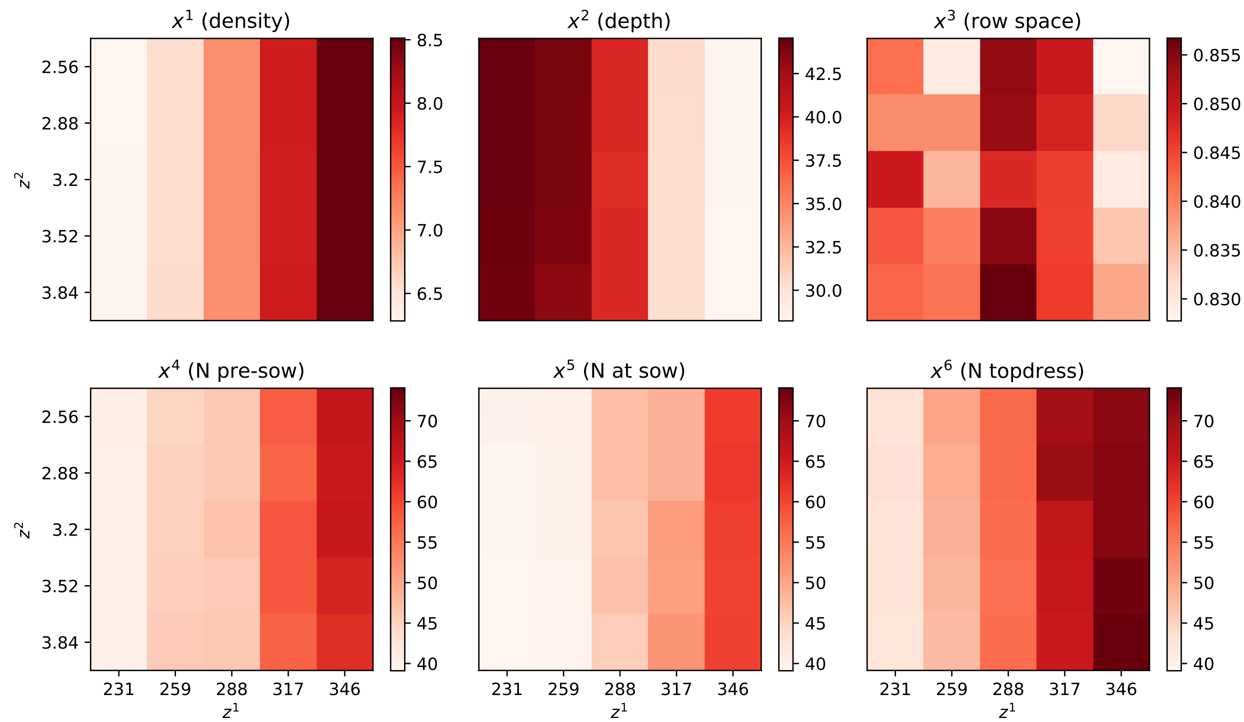

Figure 3.6: Learned SSM after 5 years in scenario A. The average sowing density is 7.3 seeds/m^2, and the average amount of total nitrogen is 156 kg/ha.

The average sowing density is

$$\frac{1}{25}\sum_{s=1}^{25} x_s^1 = 7.3 \text{ seeds/m}^2,$$

and the average amount of total nitrogen fertilizer is

$$\frac{1}{25}\sum_{s=1}^{25}\sum_{i=4}^{6} x_s^i = 156 \text{ kg/ha}.$$

As a result, \$43/ha higher profit is achieved by using 0.7 fewer seeds/m^2 and 45 kg/ha less nitrogen than uniform management. To further emphasize the generality and robustness of our algorithmic approach, results from other years than 2013 are also provided in Appendices.

3.3.2 Scenario B (high complexity)

Table 3.2 and Figure 3.7 report field-level profits (\$/ha) from implementing the SSM learned after conducting experiments for T years. Again, the value for each $T \in \{1, 2, \ldots, 10\}$ is the profit if the farmer terminates the experiments after T years and implements the learned SSM without further improvement.

Years (T)	1	2	3	4	5	6	7	8	9	10
Learned	1319	1324	1331	1333	1335	1335	1336	1337	1338	1339
Uniform	1295	1295	1295	1295	1295	1295	1295	1295	1295	1295
Difference	24	29	36	38	40	40	41	42	43	44
										(\$/ha)

Table 3.2: Field-level profits (\$/ha) from implementing the learned SSM and uniform management in scenario B.

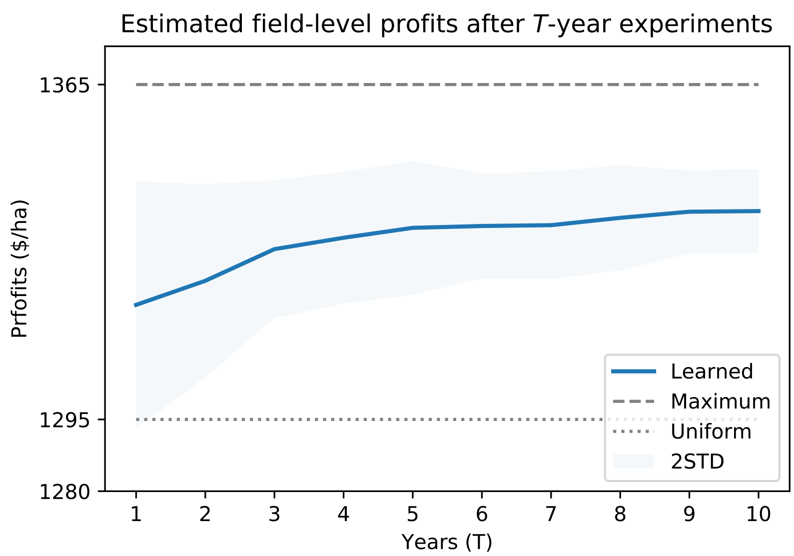

Figure 3.7: Learning curve of the algorithm with profits plotted against years of experiments in scenario B. The shaded areas indicate two standard deviations around each mean field-level profit over 100 Monte Carlo samples.

The shaded areas indicate two standard deviations around each mean field-level profit over 100 Monte Carlo samples. Again, the dashed line indicates the maximum possible profit obtained by implementing the optimal SSM, while the dotted line indicates the profit from uniform management.

The following heatmaps (Figure 3.8) compare the site-specific profits from the learned SSM at $T = 5$ and uniform management.

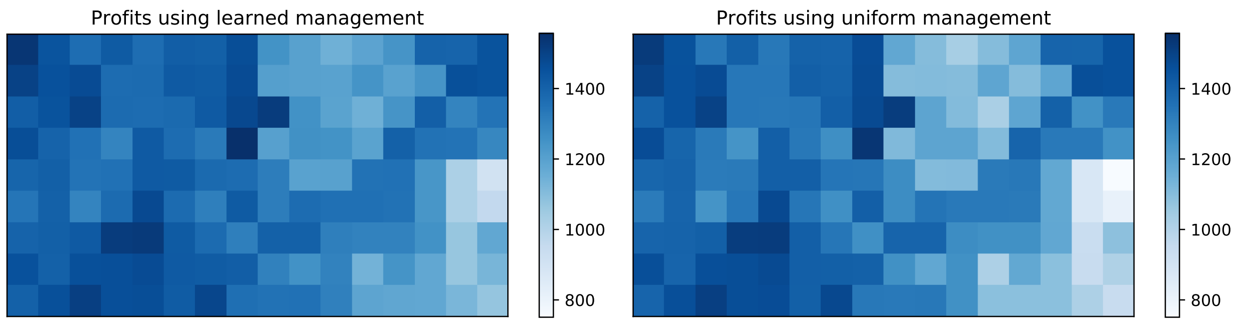

Figure 3.8: Site-specific profits (\$/ha) for the learned SSM after 5 years and uniform management in scenario B

As in scenario A, the difference between the two management systems is difficult to discern. The SSM, however, has higher profits (darker colors) in low-yielding sites. Figure 3.9 illustrates the difference at each site.

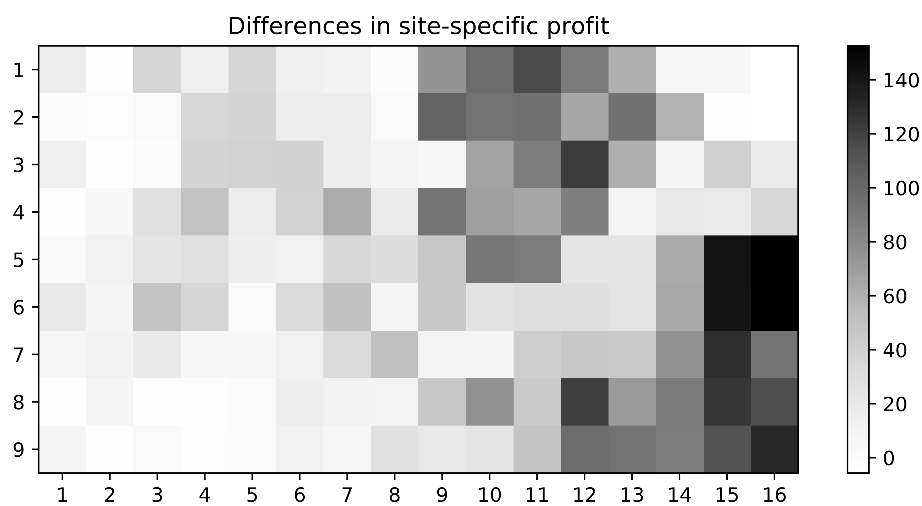

Figure 3.9: Differences in site-specific profit (\$/ha) for SSM versus uniform management in scenario B. The maximum difference is \$153/ha at site (16,6) and the minimum difference is \$-5.7/ha at site (8,1).

Finally, Figure 3.10 illustrates the learned SSM $(x^1, \ldots, x^6)$ after 5 years.

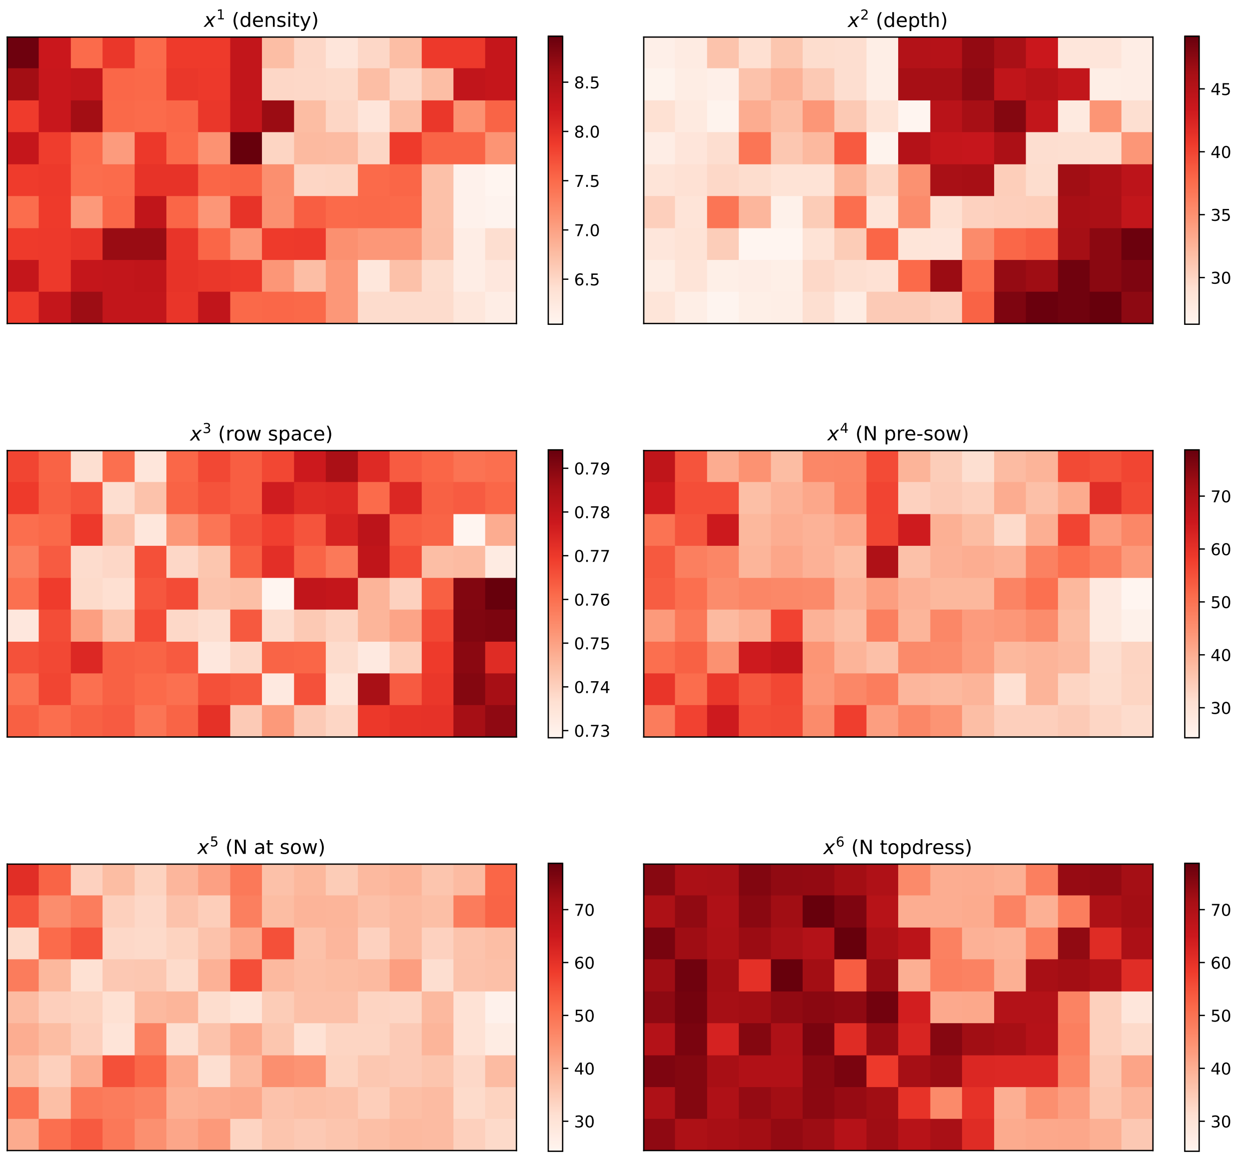

Figure 3.10: Learned SSM after 5 years in scenario B. The average sowing density is 7.4 seeds/m^2, and the average amount of total nitrogen is 146 kg/ha.

The average sowing density is

$$\frac{1}{144}\sum_{s=1}^{144} x_s^1 = 7.4 \text{ seeds/m}^2,$$

and the average amount of total nitrogen fertilizer is

$$\frac{1}{144}\sum_{s=1}^{144}\sum_{i=4}^{6} x_s^i = 146 \text{ kg/ha}.$$

As a result, \$40/ha higher profit is achieved by using 0.6 fewer seeds/m^2 and 55 kg/ha less nitrogen than uniform management.

3.4 Discussion

The results for field-level profit presented in Tables 3.1 and 3.2 and Figures 3.3 and 3.7 are well aligned with our common notion about learning—the longer an algorithm learns, the higher profit the learned SSM generates. When the algorithm starts with little existing data to incorporate, the algorithm has difficulty identifying good management. This is particularly the case in scenario A, in which the profit from the learned SSM after one year is far below that for uniform management. In contrast to scenario A, good management is found in scenario B even after one year because, although most sites are technically distinct, some sites are similar to each other and provide mutual information. As a result, with a greater number of sites in the field, more information is collected each year. Despite the low performance in the first year in scenario A, the algorithm quickly learns and its performance surpasses the performance with uniform management after two years. In both scenarios, the algorithm continues to learn and widen the performance gap. With five years of learning, the estimated profit reaches \$1,277 or 97% of the maximum possible profit (\$1,310) in scenario A and \$1,335 or 98% of the maximum possible profit (\$1,365) in scenario B.

The random sampling used in year 1 creates the large shaded area formed by two standard deviations around the mean predictions over the first few years in scenario A (Figure 3.3), indicating that imprecise prediction of mean profits. However, performance dramatically improves after four years, and thereafter its spread around the mean profits continues to shrink. Combined with the increasing mean profits, this is a desirable feature because it implies that no matter how the algorithm starts off, after several years, the algorithm consistently learns good SSM. Scenario B exhibits far less imprecision due to the larger sample size used right from the beginning (Figure 3.7).

As seen in Figures 3.4,3.5,3.8, and 3.9, the higher field-level profits for the learned SSM are due mainly to their higher profits from the low-yielding sites (e.g., sites with $z^1 \in \{231, 259\}$ in scenario A and sites around (11,2) and (15,7) in scenario B). These results imply that uniform management is excessively tailored to the high-yielding conditions—$z^1 \in \{317, 346\}$ in scenario A and the left half of the field in scenario B—leading to the decrease in profitability in the low-yielding sites where it is optimal to put less inputs. Overall, albeit not necessarily true in other environments, the algorithm discovers that it is profitable to put more inputs in the high-yielding sites and less in the low-yielding sites as indicated in panels for x^1, x^4, x^5, and x^6 in Figure 3.6 and 3.10.

We dismiss the patchy look of row spacing (x^3) in Figures 3.6 and 3.10 as an artifact of numerical optimization, which strictly distinguishes two values whenever one results in even a minuscule amount greater than the other. Indeed, the color bar for x^3 has a very small range (0.830–0.855 in scenario A and 0.73–0.79 in scenario B), indicating little practical significance for learning SSM. Nonetheless, we have included row spacing in the learning of SSM because we do not assume its insignificance before running the algorithm. In general, with little prior knowledge about which management variables are significant and should be included for learning their optimal management choices, we should include them and let the algorithm learn.

In both scenarios, the learned SSM is evidently more efficient in input use for generating profits than the benchmark uniform management. After five years, for example, the learned SSM generates \$43/ha higher profits with 45 kg/ha less nitrogen in scenario A and \$40/ha higher profits with 55 kg/ha less nitrogen in scenario B. In terms of yield, the learned management produces 210 kg/ha less maize in scenario A and 278 kg/ha less maize in scenario B. While the SSM optimization is guided by profit maximization, it turns out to be environmentally more sustainable as well because both costs of fertilizer (i.e., to profitability and to the environment) are aligned so that less is better. However, higher yield does not necessarily coincide with higher profit, though yield increases with more inputs, substantially higher input costs can reduce profit.

Although the focus of this paper is on development and demonstration of the algorithm, we mention some implications and implementation of the algorithm in practice for future empirical studies. Imagine a corn farmer in Iowa whose farming situation is well captured by scenario B. Then, implications of the higher projected profits largely depend on whether the farmer needs to invest in new equipment and if so, how large the operation scale is. As mentioned in section 2.1, a typical story about SSM is that, despite the existing access to required equipment, a lack of actionable procedure for SSM prevents the farmer from implementing it (Lindblom et al., 2017). In this case, extra profits may be sufficient to cover costs for experiment and overcome psychological barriers to change (e.g. accepting lower yield for higher profit). In contrast, if new investment is necessary, further financial analysis is required. Suppose that the farm size is 100 ha and extra \$40/ha is projected. Then, the analysis involves comparing a stream of extra \$4,000/year against equipment and other financial costs.

To implement the algorithm in practice, we suggest the following design for a web application that seamlessly works with common desktop software (e.g., Ag Leader SMS™) used to communicate with a wide range of precision farming equipment. For the sake of simple illustration, it is assumed that data is organized and contained in a single table, implying no spatial misalignment across all the variables. In addition, for all site $s \in \{1, 2, \dots, M\}$, a state vector $\boldsymbol{z}_s$ remains the same over time. First of all, as an initial setup, the farmer enters management and soil information in the web-app interface. For the former, specify management variables $(x^1, x^2, \dots)$ to be investigated and a range of values that each variable can take. For the latter, enter $\boldsymbol{z}_s$ for all s, which should be surveyed in advanced. Then, at the beginning of each year, the farmer uploads the existing data from the preceding years. Based on this information, the algorithm runs on the web server and returns a prescription $\boldsymbol{x}_s$ for all s in a file format (e.g., shapefile) that can be downloaded and imported into the desktop software. The desktop software exports the prescription data into a file format (e.g., tgt file) that can be imported into the tractor display. Now, following the loaded prescription, the farmer implements each management and, after harvesting, exports observed yield data $\{y_s\}$ out of the display. Finally, using price information, the farmer calculates profits $\{\pi_s\}$ and, along with the corresponding $\{(\boldsymbol{x}_s, \boldsymbol{z}_s)\}$, add them to the existing data.

Despite the promising results, there are several clarifications and limitations to note before real-world implementations, as well as future research needs.

- While we have formulated a farmer's problem as profit maximization and developed an algorithm to solve it, in reality, many farmers are concerned with maximizing their yields. In fact, our optimization framework is flexible and can be applied to solving yield maximization problems as well, which will be the focus in future work. In this paper, we concentrate on the profit maximization formulation for the following reasons. First and foremost, since the main objective of this paper is to introduce the novel optimization method to the literature, we try to highlight its features and usefulness for precision agriculture. To this end, profit maximization formulation is technically simpler than yield maximization. In addition, as extensively studied over the past two decades, the adoption of PA (or the lack thereof) is explained by many and complex determinants, among which profitability is identified as one of the key factors in the literature (Castle et al., 2017; Gandorfer and Meyer-Aurich, 2017; Pathak et al., 2019; Schimmelpfennig, 2016). Even for those who try to maximize yields, as part of commercial enterprises, yield maximization is rarely unconstrained. Indeed, there are typically implicit upper bound for input use (e.g., avoiding unnecessary fertilizer application). In this case, constrained yield maximization becomes similar to profit maximization.

- This analysis assumes no costs for switching management from site to site, which can be unrealistic for some inputs and management types. For example, varying types of fertilizer, seed treatments or hybrids may require equipment modifications, multiple field passes or additional labor. As another example, changing seeding rates too frequently may put an excessive strain on and damage an electric motor. If these are costs to consider, optimization of SSM will be even harder due to the switching frequencies of management as another set of control variables to optimize.

- The algorithm assumes choice variables are continuous, even though choices of continuous variables can be constrained for various practical reasons. For example, though the algorithm may recommend fertilizer application rates that differ by less than 0.1 kg/ha, such small differences are impossible to implement practically with current equipment. Real-world implementation will require modifying the algorithm to convert continuous choice variables into appropriate discrete variables.

- The simulations use a single season's weather pattern, output price, and input prices that prevail over T years. In reality, these differ from year to year, and such fluctuations may have strong implications for the algorithm's performance.

- The current algorithm assumes risk neutrality of farmers when the acquisition function prescribes the next sampling choices. In the real world, however, many farmers are risk averse (Monjardino et al., 2015), for whom it will be difficult to accept and implement some strongly explorative prescriptions made by the algorithm that reflects expected profit but not its variance. To be more realistic and useful for practical applications, future work needs to modify the acquisition function and parameterize a level of risk aversion so that the algorithm can be uniquely adjusted to each farmer.

- This approach is limited to one-shot optimization in which the farmer makes all the management decisions at the beginning of the year, implements them, and waits to see results at the end of the year. In practice, many management choices are sequentially made throughout the year, while they affect site characteristics and are affected by evolving site characteristics. To handle these more realistic situations requires dynamic models with information feedback and learning that take place both within and across years.

- The simulations assume no spatial interaction between input choices across sites (e.g., fertilizer use on one site does not affect adjacent sites). Depending on the grid granularity, these assumptions may be excessively strong under some circumstances. Thus, besides the temporal dynamics mentioned above, future work focuses on developing spatiotemporally explicit models that handle dynamics over space and time.
- As with many other machine learning studies on large and complex datasets, it is extremely beneficial to have access to a simulator of high fidelity. We find in APSIM only six management variables and four environmental variables suitable for this study. Since the algorithm is very flexible and capable of dealing with an arbitrary number of variables, demonstration of the algorithm would be more realistic and compelling with a simulator in which a greater number of variables and resulting yield are intricately interdependent, better representing the complexity of biophysical systems. Thus, further advancement of crop simulators is crucial for validation and improvement of the current algorithm.

3.5 Conclusions

We have proposed an algorithmic approach to optimizing complex site-specific management with many management and environmental variables. The proposed algorithm enables individual farmers to efficiently learn their own site-specific management through on-farm experiments. We have demonstrated its performance using simulated environments. The results have provided a positive answer to both the learnability of complex site-specific management and the higher profitability possible relative to uniform management. While these results are promising, we do not know in general under which environment the algorithm works. Thus, we need more validation studies, especially field experiments with farmer co-operators. Long-term, the results suggest that on-farm experimentation implemented with precision agriculture equipment can help to realize the benefits of precision agriculture—more profitable management, greater food security, and improved environmental sustainability.

Appendices

APSIM configuration

As a basis, we use the Continuous Maize module in APSIM. Then, to simulate maize production in Ames, Iowa, we modify its default settings as follows.

- Metfile: USA_Iowa_Ames.met
- Calendar: Jan 1, 2013 - Dec 31, 2013
- Cultivar: Pioneer 3394
- Sowing window START data: 15-apr
- Sowing window END data: 2-may
- Soil: Iowa Nicollet soil series
- Initial nitrogen: 0 kg/ha for both NO3 and NH4 for scenario A
- Initial water: 80% filled from top

Sensitivity analysis

In addition to two simulated environments with medium- and high-complexity, to further emphasize the generality and robustness of our algorithmic approach, we conducted simulation experiments in different years than 2013. Since output price, input prices, and weather are all dependent on a particular year, the differences in year provide different environments for profit maximization. The price information for each year was obtained from the same sources (Duffy, 2013; Johanns, 2019).

Year	2007	2008	2009	2010	2011	2012
Output price ($/kg)	0.17	0.16	0.14	0.21	0.24	0.27
Seed price ($/1000 seeds)	1.82	2.10	3.13	3.44	3.25	3.40
Nitrogen price ($/kg)	0.69	1.02	1.51	0.73	1.13	1.40

Note that all sensitivity analysis was conducted under the environments with medium complexity, because of the significantly greater computational resources required in environments with high complexity. While there were considerable variations in both the growing and economic conditions across the different years, overall, the algorithm is quite versatile and able to learn good SSM within a few years in every environment. Similar to Figure 3.3, for each environment, we plot estimated field-level profits after T-year experiments.

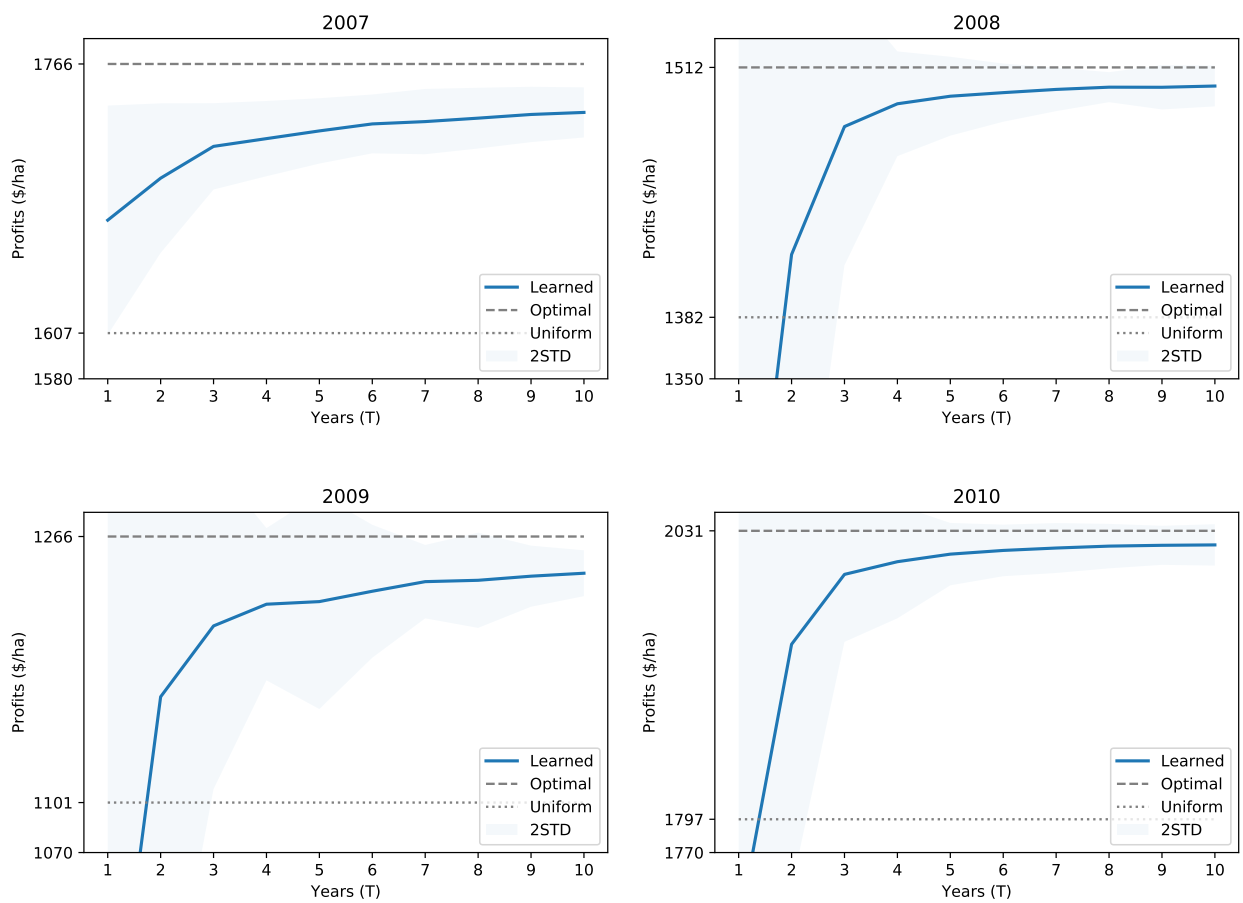

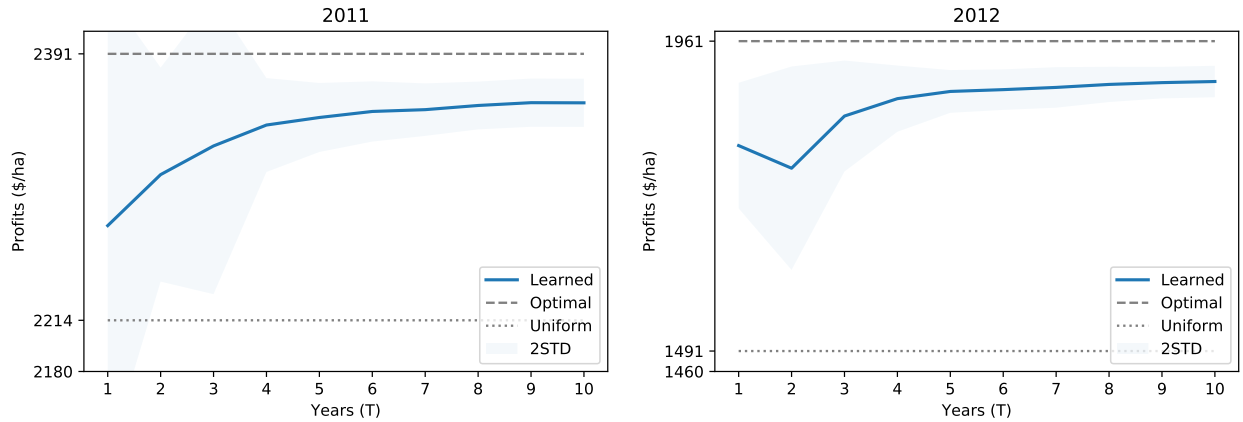

Constructing scenario B

To generate a state vector $z = (z^1, z^2, z^3, z^4)$ for each site, we need to choose which site s and what values for $(z_s^1, z_s^2, z_s^3, z_s^4)$. For both purposes, we use random walk. Start from the mid site $(9, 5)$ with the initial values $(288, 3.2, 10, 10)$ assigned. Then, with probability of $1/3$, randomly either move right, move left, or stay. Similarly, with probability of $1/3$, randomly either move up, move down, or stay. This gives us the next site to consider. If the move means hitting a boundary, it stays at the site. Once moving into the new site, see if the site has already been assigned a state vector. If not, with probability of $1/3$, randomly perturb the state vector at the originating site by either -5%, 0%, or 5%. Continue the process until all sites are assigned a state vector.

Chapter 4

An agent-based model of insect resistance management and mitigation for Bt maize: A social science perspective

Yuji Saikai
Terrance M. Hurley
Paul D. Mitchell

4.1 Introduction

Globally, farmers have planted more than 2.3 billion hectares of genetically engineered crops since their commercial introduction in 1996, including a new maximum of 190 million hectares in 2017 (ISAAA, 2017). Focusing on maize (*Zea mays*), the world's leading grain crop with annual production exceeding a billion metric tons, the United States, Brazil and Argentina together produced almost half of the world's supply in 2017 (USDA, 2018b). Bt maize—maize genetically engineered to produce *Bacillus thuringiensis* (Bt) toxins in plant tissues for insect control—accounted for more than 80% of the maize planted in each of these three nations in 2017 (ISAAA, 2017). After more than two decades of commercial use of genetically engineered crops, insect resistance to Bt toxins continues to be a major concern around the world (Tabashnik, 2015). A high-dose/refuge resistance management strategy continues to be the primary policy in multiple nations for delaying resistance to these Bt toxins (Huang et al., 2011; Pardo-López et al., 2013; Tabashnik et al., 2013). Nevertheless, field-evolved resistance to some of these Bt toxins has been documented for populations of western corn rootworm (*Diabrotica virgifera virgifera*) in the United States and various lepidopteran species in multiple locations (Gassmann et al., 2011).

The commercialization of Bt crops has generated a variety of research, including bioeconomic models that integrate population genetics and pest ecology with farmer economic returns (Crowder et al., 2005; Hurley et al., 2001; Mitchell and Onstad, 2005). Though these models contributed to the development of insect resistance management policies, little other work exists on the role of social factors in the evolution of insect resistance to commercialized toxins. Insect resistance to these toxins evolves in response to human management activities, activities driven by a variety of social factors that include not only economic considerations, but also sociological, psychological, cultural, historical and political considerations (Hurley and Mitchell, 2014). As a result, examining genetic and ecological processes in isolation from these broader social factors driving human behavior potentially misses key determinants of the evolution of insect resistance. Hence, a broader, complex systems model of insect resistance management that incorporates both biological and social processes can potentially provide new insights (Rebaudo and Dangles, 2013).

In the United States (US), the Environmental Protection Agency (EPA) required companies commercializing Bt crops to develop resistance mitigation plans as a condition for product registration (USDA, 2010). Once a resistant population has been officially documented according to the EPA process, these resistance mitigation plans generally restrict the availability of the technology (Bt seed) in and around the region where the resistant population emerges. Though resistant insect populations and field failures in the US have been documented in the scientific literature (Gassmann et al., 2011, 2014), the official EPA criteria have yet to trigger implementation of these mitigation plans for any pest. Instead, the EPA has required a more generalized response by Bt crop registrants (Andow et al., 2016). Interestingly, little research exists that evaluates and compares the mitigation plans that have been filed or other mitigation policies, particularly from an economic perspective. Given the length of time that Bt crops have been in use in the US and elsewhere, insect resistance is likely to become an increasing problem, making more research on mitigation responses and strategies especially timely.

This paper has two goals. First, we develop an agent-based model of insect resistance to Bt maize that incorporates farmer adoption behavior. We then use the model to compare different mitigation policies in order to inform policymakers and other stakeholders of the types of programs that are likely to generate the largest economic benefits for society. Second, focusing specifically on the impact of social networks on farmer adoption behavior, we show that social factors can also play a key role in the evolution of insect resistance to Bt toxins in agricultural cropping systems.

Agent-based modeling has become more widely-used for studying complex systems and emergent behavior, including socio-ecological modeling of insect resistance management (Miller and Page, 2009; Peck, 2004). In agent-based models, an observed macroscopic phenomenon emerges as a result of interaction among heterogeneous agents in a dynamically evolving environment. Agents typically follow simple decision rules and influence each other either directly or indirectly through the environment, which itself evolves according to its own rules and agent actions. Because the processes being explicitly modeled are complex, researchers use computer simulations to examine outcomes over a wide range of parameter values. In short, agent-based models are laboratory experiments conducted in silico (Epstein, 2006; Peck, 2004). Despite the remaining challenges to overcome, such as ad hoc assumptions and lack of relevant data for validation (Durlauf, 2012; Feola and Binder, 2010; Filatova et al., 2013), agent-based modeling can provide insights into complex systems that would be difficult to study otherwise. Given the merits, applications of agent-based models to pest and resistance management in agricultural systems have been developed (Gay et al., 2017; Renton, 2013; Renton et al., 2014).

Although agent-based models can integrate many factors, they still face the fundamental tradeoff in modeling: fidelity to the phenomenon being examined and abstraction for ease of analysis and interpretation (Peck, 2004). This paper focuses on deriving new insights into policy options for mitigating insect resistance once it has evolved, and emphasizes the significance of social factors for questions relevant to policymakers (Renton, 2013). As a result, social components are richer than existing models that use individual-based modeling to incorporate social factors (Milne et al., 2015), while the biological aspects of the model are simpler than other models focusing on biological processes (Ives et al., 2017; Onstad and Meinke, 2010; Storer, 2003).

We extend existing work (Milne et al., 2015) on insect resistance management for Bt crops by more fully leveraging the power of agent-based modeling. First, we explicitly model the local influence that neighbors have on farmers through social networks as they make decisions regarding adoption of Bt maize, creating a hybrid decision process that mixes both individual profit considerations and a desire to mimic neighbors. Second, we allow the additional cost of planting Bt seed to vary over time, because this cost influences adoption decisions and companies have reduced the cost of single-toxin Bt seed to encourage farmers to continue to plant Bt maize in the face of pest population suppression (Shi et al., 2010). With this pricing flexibility, we calibrate the farmer decision model using historical data that reflects these decreasing prices, and then can examine the impact of a tax on Bt seed as a policy option for mitigating resistance.

For this analysis, we parameterize a bioeconomic model of maize production with the option to use high-dose Bt maize to manage European corn borer (*Ostrinia nubilalis*). We calibrate the Bt maize adoption model using aggregate historical adoption data for farmers in the state of Wisconsin. Through the calibration process, we emphasize the significant role that a social factor—the local influence of social networks on Bt maize adoption (Kaup, 2008)—can play in the evolution of insect resistance. Using the calibrated model, we then simulate a number of mitigation policies implemented either over the entire landscape or around the areas where resistance develops. In particular, we consider combinations of an increased refuge requirement and a tax on the sale of Bt seed for all farms, and a ban on the use of Bt maize and areawide use of an additional insecticide to control the pest in the area around where resistance emerges. To assess the relative performance of each policy, we use economic surplus as a monetary measure of the social value generated by the use of Bt maize and conduct sensitivity analysis of key parameters to explore the robustness of model results.

4.2 Materials and methods

4.2.1 Landscape

The spatially explicit model used a 30×70 grid space representing the cropland in Wisconsin. Modeled farmers mimic the Wisconsin crop landscape (USDA, 2012) and plant 44% of the fields to maize, the host crop for the pest. Fields maintain their initial random assignment to maize or non-maize production during a simulation but are reassigned at the start of each new simulation. During a simulation, maize farmers decide adoption of Bt maize each period. Figure 4.1 depicts a typical model landscape, in this case with 59% Bt adoption and a resistance allele frequency of 41% for the total population. A circle ($\circ$) represents a farmer who plants conventional (non-Bt) maize, whereas a black dot ($\bullet$) represents a farmer who adopts Bt maize. A light-gray background (■) indicates that the pest population in an individual field before adult dispersal has a resistance allele frequency of more than 50%, the criterion used for declaring that a population is resistant (Tabashnik, 2008). To avoid boundary effects, top fields wrap to corresponding bottom fields and left-most fields to corresponding right-most fields, creating a torus, implying that the model space is part of a larger landscape with comparable dynamics occurring for the pest population and its genetic structure (Storer, 2003).

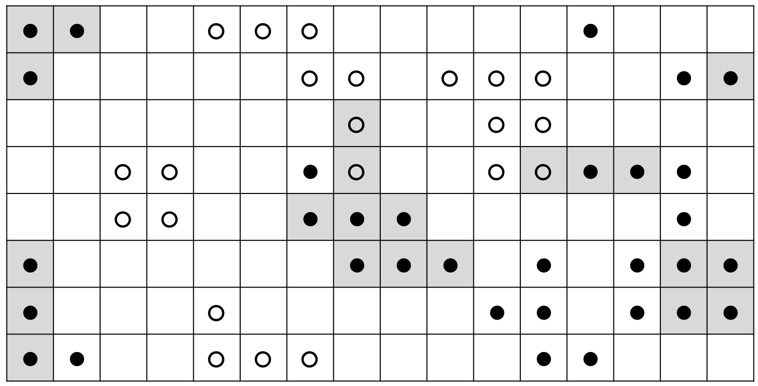

Figure 4.1: Example model landscape. Circle (○) represents a field planted to non-Bt maize, black dot (●) represents a field planted to Bt maize, and light-gray (■) indicates a field with a pest population with a resistance allele frequency of more than 50% before adult dispersal.

4.2.2 Pest Population Genetics

The pest population-genetics model is parametrized for European corn borer (*Ostrinia nubilalis*), a major pest for Midwestern maize and the primary target for initial commercial releases of Bt maize beginning in 1996 (Hutchison et al., 2010). The insect model uses discrete time steps corresponding to distinct generations and consistent with the seasonality of many types of crop production and pest life cycles. *O. nubilalis* typically has two generations per year in the major US maize production region, though northern regions may have only one generation per year and southern regions may have three or more (Mason et al., 1996). The model simplifies these dynamics to one discrete time step per year that aggregates population dynamics and genetic selection across these generations. Hutchison et al. (2010) used a comparable empirical approach to estimate annual population growth rates for *O. nubilalis* using annual observations of second-generation adult population densities in Minnesota and Wisconsin.

Historically, the *O. nubilalis* population in the Midwestern US has oscillated with an approximately seven year cycle (Onstad and Maddox, 1989) largely due to the entomopathogenic parasite *Nosema pyrausta* (Bell et al., 2012). Field data for second-generation populations in Wisconsin over 1944-1995 show an average peak and trough for the oscillation of about 1.2 and 0.2 larvae per plant (Bell et al., 2012; Hutchison et al., 2010). The population model approximates these dynamics using a lagged logistic growth model:

$$N_{t+1} = gN_t \left(1 - \frac{N_{t-1}}{K}\right) \tag{4.1}$$

where N_t is the second-generation larval population (larvae per maize plant), g is the annual growth rate and K is the carrying capacity. Using $g = 2.15$ and $K = 1.4$ generates a reasonable approximation of historical *O. nubilalis* population dynamics in Wisconsin, with a similar range of population minimums and maximums as observed and six or seven years between peaks.

The genetics model assumes two alleles, R for resistant and S for susceptible, creating three genotypes, homozygous resistant RR, homozygous susceptible SS, and heterozygous RS, with respective Bt toxin larval survival rates of 1.0, 0.0, and 0.18. Each period, after Bt toxin mortality and density-dependent mortality for larvae, random mating occurs among the adult population within each field before adult dispersal. Note that, with random mating, RR and RS genotypes both contribute R alleles, but RS genotypes do so half as often on average, also contributing S alleles just as often. Let α_t, β_t, and γ_t respectively denote the relative frequencies in period t of RR, SS, and RS genotypes, which by definition sum to 1. Random mating then implies

$$\begin{aligned}
1 &= \alpha_{t+1} + \beta_{t+1} + \gamma_{t+1} \\
&= (\alpha_t + \beta_t + \gamma_t)^2 \\
&= (\alpha_t + 0.5\gamma_t)^2 + (\beta_t + 0.5\gamma_t)^2 + 2(\alpha_t + 0.5\gamma_t)(\beta_t + 0.5\gamma_t)
\end{aligned}$$

so that

$$\begin{aligned} \alpha_{t+1} &= (\alpha_t + 0.5\gamma_t)^2 \\ \beta_{t+1} &= (\beta_t + 0.5\gamma_t)^2 \\ \gamma_{t+1} &= 2(\alpha_t + 0.5\gamma_t)(\beta_t + 0.5\gamma_t) \end{aligned} \tag{4.2}$$

Adults disperse uniformly within a radius of maize fields (i.e., not onto non-maize fields). While the literature provides a range of observations for dispersal due to various influencing factors including season, gender, and mating status (Dorhout et al., 2011, 2008; Showers et al., 2001), it takes place within 20km in most cases. To capture the effects of dispersal yet remain computationally tractable, the model uses a dispersal radius of 15km, which corresponds to 3 fields in the model grid space. We assume that the natal field is always available as a destination, which implies that if no neighboring fields exist within the dispersal range, adults stay in the same field.

4.2.3 Farmer behavior

Individual farmers manage each field and decide each period whether to plant Bt or non-Bt maize. A number of economic and social factors influence farmers' adoption decisions for Bt maize (Useche et al., 2009). Rather than explicitly enumerating and modeling these multiple factors, agent-based models combine simple behavioral models with suitable random components and let complex phenomenon emerge (Epstein, 2006). Though expected profitability greatly influences farmer management decisions in commercial agriculture, their local social networks also significantly influence their behaviors, not just by providing additional information regarding the relative profitability of different practices (Kaup, 2008; Lubell and Fulton, 2007; McAllister et al., 2015). Therefore, we model the Bt adoption process as a hybrid of individual profit maximization and local imitation to capture the effect of social networks.

4.2.4 Farmer profit

The profit-based component of farmer behavior uses the following switching function (Milne et al., 2015):

$$\Pr(\text{Switch C to A}) = \begin{cases} 1 - \exp[-\rho(\pi_A - \pi_C)] & \text{if } \pi_A > \pi_C \\ 0 & \text{otherwise} \end{cases} \tag{4.3}$$

Here, π_A is the profitability of the alternative choice and π_C is the profitability of the current choice, with both profits calculated using the pest population density for the previous period in the field. The function determines the probability that the farmer switches from the current choice to the alternative (Switch C to A), with the probability increasing as the alternative becomes relatively more profitable than the current choice. We use a "soft" probabilistic switching decision to capture the effect of other unobserved individual factors (Mason et al., 1996). The parameter ρ captures the responsiveness of farmer adoption to profitability differences, with a greater ρ increasing the probability farmers use the more profitable alternative. As explained in the Calibration section, we calibrate ρ against the Bt seed adoption data for Wisconsin to derive $\rho = 0.0023$. The negative-exponential function implies that the switching probability is the farmers' expected utility gain from switching when the gain is uncertain, assuming constant absolute risk aversion for the farmer, a commonly used assumption for empirical analysis (Chavas, 2004; Mitchell and Hutchison, 2008).

Farmer profit for a field (π) is crop revenue minus cost, where revenue declines as the pest population increases and cost varies with the scenario:

$$\pi = PY\,(1 - Loss(N)) - Cost \tag{4.4}$$

Here P is crop price (\$Mg^{-1}), Y is potential or pest-free crop yield (Mg ha^{-1}), $Loss$ is proportional crop loss, which depends on N, the average pest population density (larvae per plant), and $Cost$ is the production cost (\$ ha^{-1}). To focus on factors other than annual variability in crop prices and yields, crop price and potential yield are fixed at reported averages in 2017 for Wisconsin farmers: $P = \$129.91$ ha^{-1} and $Y = 10.92$ Mg ha^{-1} (USDA, 2019). These values imply constant potential revenue of \$1,418.62 ha^{-1} across fields and seasons. The proportion of potential revenue lost due to pest damage depends on the average larval population density based on an empirical model (Hurley et al., 2004): $Loss(N) = 0.1186N^{0.5146}$.

$Cost$ consists of a base cost C (\$ ha^{-1}) that does not vary by policy scenario and costs that do:

$$Cost = \begin{cases} C + T(1-\theta)(1+\tau) + C_s & \text{if Bt maize} \\ C + C_s & \text{if non-Bt maize} \end{cases} \tag{4.5}$$

Based on US Department of Agriculture crop budgets (USDA, 2016), the base cost C is set to \$1,202.51 ha^{-1}, the reported average for 2017 in the region containing Wisconsin for all costs except opportunity costs for land and operator labor and management. T (\$ ha^{-1}) is the additional seed cost for Bt maize ("technology fee"), which varies over time based on the function estimated with Wisconsin market data (Onstad and Meinke, 2010). Specially, $T = \$17.49$ ha^{-1} from 1996 to 2003, and then declines to $T = \$17.45$ ha^{-1} for 2004, \$15.78 ha^{-1} for 2006, \$13.75 ha^{-1} for 2007, \$11.41 ha^{-1} for 2008, \$9.18 ha^{-1} for 2009, \$8.29 ha^{-1} for 2010, \$7.82 ha^{-1} for 2011, \$7.39 ha^{-1} for 2012, and then remains at a base of \$7.04 ha^{-1} for years 2013 and afterward. The remaining cost parameters vary with the policy scenario: θ is the proportion of refuge (non-Bt maize) planted with Bt maize, τ is the tax rate for Bt maize, and C_s is the cost (\$ ha^{-1}) for a foliar insecticidal spray as part of areawide management of adults. Each scenario sets these cost parameters at appropriate values as described in the Policy Experiments section. For example, a refuge only scenario sets $\tau = C_s = 0$ and sets θ at 0.05, 0.20 or 0.50; a ban only scenario sets $\tau = C_s = 0$ and $\theta = 1$ (100% refuge) in fields where a ban is in effect; and an areawide spray policy sets $\tau = 0$ and imposes the cost C_s for all affected fields. The cost of an insecticidal spray C_s is \$33.51 ha^{-1} based on published survey averages for insecticide active ingredients used in maize and application costs, adjusted for inflation to 2017 equivalents (Mitchell, 2017; USDA, 2018a).

Lastly, a refuge policy is implemented as a fixed proportion θ of non-Bt maize with complete compliance by farmers, the so-called "refuge in a bag" (Hodgson, 2010) in which the company mixes Bt and non-Bt maize seeds before purchase. In our model, the refuge requirement has two effects. First, the effective seed cost is the proportion $(1 - \theta)$ of the technology fee T at that period. Second, the effective survival rate of each genotype is calculated as the weighted average: $\theta + (1 - \theta)s$, where s is the original survival rate. That is, with probability θ, any genotype survives due to the non-Bt maize, and with probability $1 - \theta$, each genotype survives according to its Bt toxin survival rate. Initially, we assume $\theta = 0.05$, which is the lowest refuge requirement already in place, and later increased refuge levels are examined as resistance mitigation policies.

4.2.5 Social network

Network analysis has been widely applied to understand the diffusion of innovations as a social phenomenon, including in agriculture (Easley and Kleinberg, 2010; Jackson, 2010). Neighboring farmers have been shown to create a local environment that affects individual farmer adoption decisions, both for hybrid maize seed and for Bt maize (Kaup, 2008; Ryan and Gross, 1943). To capture this social network effect, the model assumes each farmer in a field is connected to farmers in neighboring fields, with the size of the neighborhood determined by a "radius". Figure 4.2 shows an example of a size-2 neighborhood for a farmer with nine neighbors who plant maize, either Bt or non-Bt.

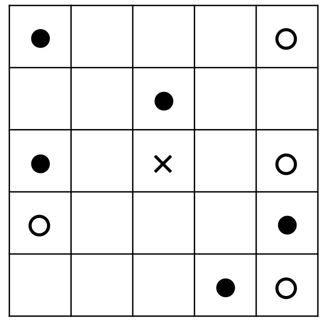

Figure 4.2: Example size-2 neighborhood. It is centered on a farmer (×) with nine neighbors who plant maize, either Bt maize (•) or non-Bt maize (◦).

Those neighbors themselves have their own neighborhoods, with each connection undirected so that the local social networks are tightly overlapped. The number of neighbors for a size-n neighborhood can range from 0 to a maximum of $4n(n+1)$. With no data for social network sizes for farmers and considering that $n = 3$ gives up to 48 neighbors (implying a substantial computational burden), the model randomly assigns a neighborhood size to each farmer for all seasons using a uniform distribution over $\{0, 1, 2\}$. Given this local social network, each maize farmer chooses each season to grow either Bt or non-Bt maize for a field. A parameter q defines the impact of social networks on farmer adoption decisions. With probability q, a farmer focuses solely on individual profits using the switching function and with probability $1-q$ follows the majority choice of his neighbors in the previous season. For example, the farmer in Figure 4.2 follows the majority and plants Bt maize next season because his neighborhood has 5 Bt maize adopters and 4 non-adopters.

4.2.6 Running the model

Each model run begins with initialization, including randomly placing farmers across the landscape. Since corn fields occupy roughly 44% of total farmland in Wisconsin (represented by 30×70 fields configured as a torus), the total number of maize farmers for a run is approximately $0.44 \times 30 \times 70 = 924$. After initialization, the run proceeds period by period, with a period corresponding to a growing season or year. Before introducing Bt maize into the model, the insect module runs for 11 periods, which corresponds to the pre-Bt periods and helps stabilize the model's biological dynamics. Thereafter, the model simultaneously updates the pest population density of each field for each period. First, Bt toxin effects reduce each fields pest population based on the survival rates of the genotypes established there the previous season. Second, mating determines the genotype composition of the next generation based on random mating of the population in the field. Third, reproduction determines the pest population density based on the lagged logistic growth model. Fourth, the population locally re-mixes across fields based on the dispersal model. Finally, maize farmers simultaneously make planting decisions (whether to plant Bt or non-Bt maize) for the next period based on the farmer behavioral model. In short, during a growing season the Bt toxin (if present) reduces the natal population in a field, survivors randomly mate and produce the next generation, which then disperses locally across fields, and then farmers make maize planting decisions for the following spring.

4.2.7 Calibration

We used aggregate Bt maize adoption data for Wisconsin to calibrate the model. Our calibration minimized the average of the mean squared error (MSE) of prediction for the simulated landscape compared to the observed data. Specifically, the MSE for a run was the squared deviation of the simulated Bt adoption rate from the annual Wisconsin adoption data, averaged across all periods with adoption data ($t = 11$ to 32). Since runs were random, the MSEs were averaged across 1,000 runs. The two calibration parameters were the responsiveness of farmers to expected profit differences between alternatives (ρ) and the probability (q) that farmers focus solely on profit differences to make adoption decisions, rather than their neighbors' choices. To avoiding both over-fitting and excessive computational requirements, a grid search was used with increments of 0.0002 for ρ and of 0.1 for q. To highlight the significance of local networks, we also calibrated the model by fixing $q = 1$ and using only ρ, which "shut off" all social network effects on adoption.

Plotting the Wisconsin Bt maize adoption data and both calibration fits shows the superior fitting of the two-parameter hybrid model relative to the one-parameter model (Figure 4.3). Using the same random seeds for both models, the optimum solutions are $\rho = 0.0036$ and $q = 0.3$ for the hybrid model and $\rho = 0.0022$ for the single parameter model. These optimal values for the two-parameter model imply that 70% of the years, farmers follow the majority choice of their neighborhood, suggesting that network effects are important for understanding farmer adoption dynamics for Bt maize.

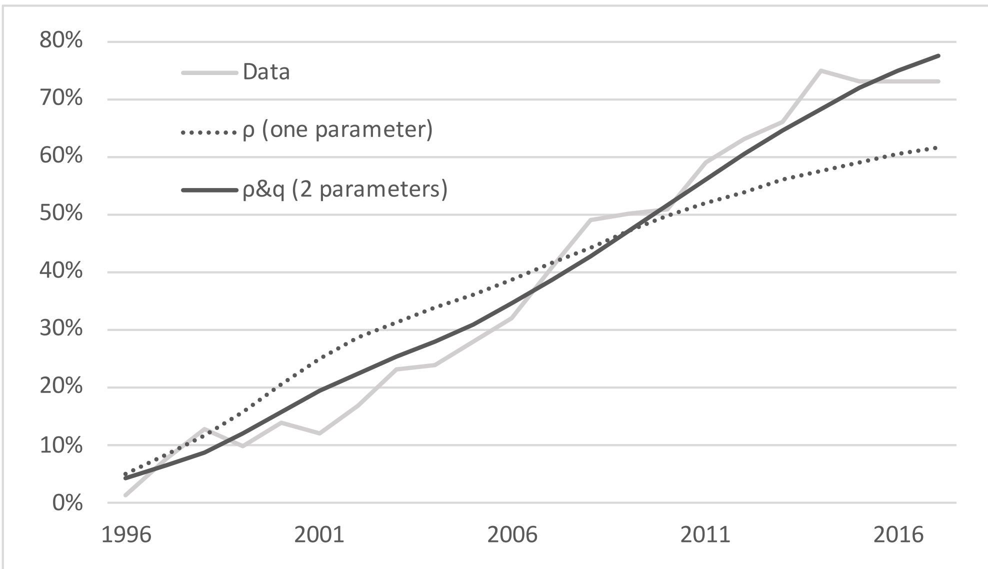

Figure 4.3: Aggregate adoption of Bt maize in Wisconsin and two simulated results. The simulated results are generated by calibrating one parameter and two-parameters.

4.3 Results

4.3.1 Baseline results

Running the calibrated model 1,000 times with different random seeds and averaging over these iterations gave baseline results for the insect population, the Bt seed adoption rate, and the resistance (R) allele frequency at the landscape level. In the model, periods 0 to 10 were an initialization phase, periods 11 to 32 were a calibration phase corresponding to years

1996 to 2017, and periods 33 to 60 were projections (see Materials & Methods). The baseline model captured the aggregate Bt adoption rate of Wisconsin farmers by calibrating two parameters that determined Bt maize adoption—farmer responsiveness to profit incentives and the farmer tendency to mimic the adoption decisions of neighbors due to social network effects. The calibrated model reproduced the previously noted oscillation of the European corn borer population before the advent of Bt maize (Bell et al., 2012), and the documented suppression of the pest population due to the widespread farmer adoption of Bt maize in Wisconsin and other states (Hutchison et al., 2010). As expected, the calibrated model projected a surge in the R-allele frequency as the insect resistance developed, resulting in the eventual recovery of the pest population. Baseline results suggested that period 33 was the beginning of a significant increase in the R-allele frequency. In period 33, the R-allele frequency was 4.1%, but rose quickly, exceeding 10% in period 36, 20% in period 38, 30% in period 39, 40% in period 40 and 50% in period 41. The pest population did not recover until later, with the average density not exceeding 0.5 larvae per plant until period 50.

4.3.2 Policy experiments

We simulated policies to mitigate resistance to the Bt toxin once it emerged. Refuge requirements have been the lynch pin of resistance management, and so the mitigation policies we examined began with increasing refuge requirements. In addition, building on the model's capacity for capturing the complexity from the interaction of biological and social factors, we experimented with combinations of three other types of mitigation policies: localized bans on the use of Bt maize around areas where resistance emerged, areawide applications of other insecticides to control the pest around areas where resistance emerged, and a uniform tax on the sale of Bt seed for all farmers buying it. Refuge policies and localized bans directly regulate the use of Bt maize, areawide spray policies directly manage resistant pest populations, and the Bt seed tax adjusts farmer incentives to use Bt maize. The simulation of resistance mitigation policies was a combination of different assumptions for these four policy parameters: the refuge requirement, localized bans, areawide management, and a Bt seed tax.

The simulated landscape consisted of a grid of fields, with 44% of fields assigned randomly to maize production initially and the remainder to non-maize. Fields remained in their initial allocation throughout a simulation, but were reassigned for each simulation. During a simulation, insect resistance was declared when the R-allele frequency exceeded 50% in the pest population in a field after Bt toxin mortality and before pest dispersal occurred. The 50% threshold was chosen because at this level the landscape-average pest population began to increase (Figure 4.4), implying that higher population densities were occurring in some fields due to resistance. For resistance mitigation, the refuge requirement was increased from the baseline of 5% to either 20% or 50% for all farmers on the landscape planting Bt maize, with complete compliance achieved using seed mixtures. The localized ban was imposed only on farms within a radius r of any field where resistance was declared, again with complete compliance assumed. We considered two radii: once and twice the distance of adult dispersal from the natal field ($r = 1 \times$ dispersal, $r = 2 \times$ dispersal). Conceptually, this ban was a 100% refuge requirement applied locally and dynamically imposed and lifted according to the situation in the previous period. For areawide management, a non-Bt insecticide was applied in the period when resistance was declared, either covering only the field of resistance or all maize fields in a neighborhood around the field within the distance of adult dispersal ($r = 0 \times$ dispersal, $r = 1 \times$ dispersal). We assumed 100% compliance with the insecticide application for all fields within this area and that the application reduced the pest population by 80% after Bt toxin mortality and increased farmer costs by \$33.51 ha^{-1}. This cost was based on published survey averages for active ingredient and application costs and adjusted for inflation to 2017 equivalents (Mitchell, 2017; USDA, 2018a). Finally, the tax policy increased the Bt seed cost by 25% or 50% for all farmers on the landscape for all periods after resistance was declared. In brief, each policy parameter had the following three levels: refuge requirement (5%, 20%, 50%), localized ban (none, $r = 1 \times$ dispersal,

$r = 2 \times$ dispersal), areawide spray (none, $r = 0 \times$ dispersal, $r = 1 \times$ dispersal), and Bt seed tax (0%, 25%, 50%). Three levels for each of these four policy parameters created $3^4 = 81$ mitigation policy combinations to simulate.

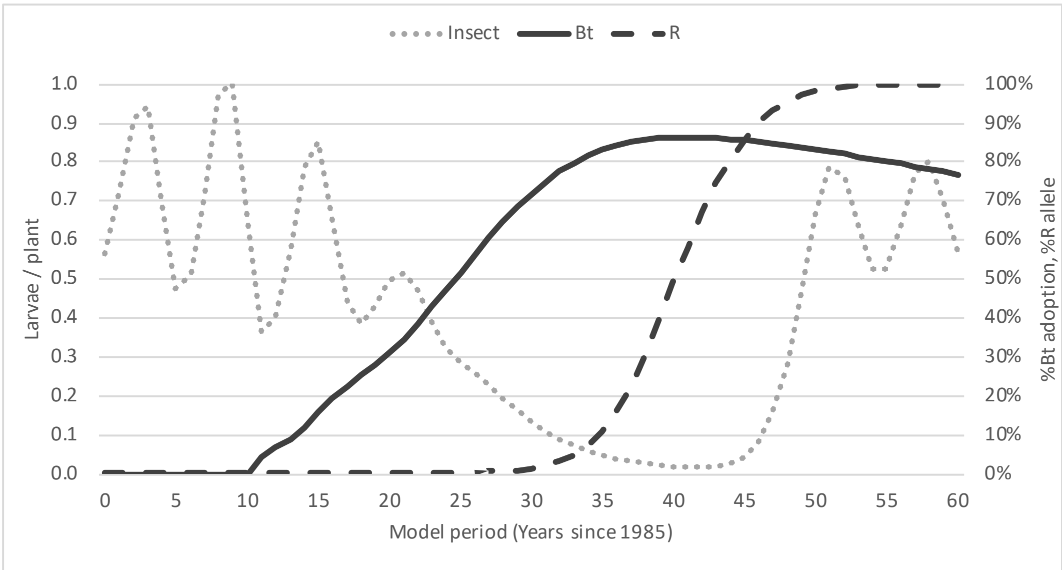

Figure 4.4: Baseline results from the calibrated model. It contains the insect population density (Insect), Bt seed adoption rate (Bt), and the resistance allele frequency (R) (results for each period are averages over 1,000 simulations).

The calibrated model was run 1,000 times for each policy and, just as for the baseline, the following three results variables were averaged over all 1,000 iterations for each period: aggregate farmer adoption of Bt maize, population-level R-allele frequency, and average pest population density for the landscape. In addition, as a performance metric to compare each policy, we approximated economic surplus each period as the sum of farmer profits and the technology fees collected by the seed company, divided by the total number of farmers.

Costs for spraying insecticides were subtracted from farmer profits for those making applications, while collected taxes were subtracted from farmer profits, but added to the economic surplus (i.e., the tax was a surplus transfer, not a surplus loss). To simplify the analysis, we did not discount future surpluses. Each policy scenario began after the calibration phase (i.e., at period 33), and the cumulative surplus was evaluated for each length of planning horizon ranging from 1 to 25 years (i.e., periods 33 to 57).

To build intuition about the nature of each policy treatment (i.e. refuge, tax, spray, and ban), we first report results for each policy individually (not combinations of policies) by plotting the dynamics for Bt adoption, the R-allele frequency, and the pest population density (Figure 4.5–4.7). In Figure 4.5 (Bt adoption), results for the ban policy (Ban 1x) are plotted with a separate vertical axis due to its qualitatively different and much stronger effect than for the other policies. Also, results for the spray policy with $r = 1 \times$ dispersal (Spray 1x) and the ban policy with $r = 2 \times$ dispersal (Ban 2x) are omitted as they were very similar to those with smaller radii. In total, Figure 4.5 plots the Bt adoption rate against the planning horizon for the following policies: baseline (Baseline), 20% refuge (20% Refuge), 50% refuge (50% Refuge), 25% seed tax (25% Tax), 50% seed tax (50% Tax), areawide spray in the field with resistance (Spray 0x) and a localized ban on Bt seed within one pest dispersal radius of the field with resistance (Ban 1x). Consistent with Figure 4.4, the baseline policy showed a continuing increase in Bt maize adoption from planning horizon year 0 (period 32 in Figure 4.4), with a peak of almost 86.5% in planning horizon year 9 (period 41 in Figure 4.4). All policies showed this same general trend (with one exception), but with a lower adoption peak occurring sooner for the seed tax policies, a higher adoption peak occurring later for the increased refuge policies (especially for 50% refuge), and a slightly higher and later peak occurring for the areawide spray policies. The one exception were localized bans the sale of Bt seed, for which implementation caused a rapid decline in the use of Bt maize, with almost complete dis-adoption by the end of the simulation in horizon period 25.

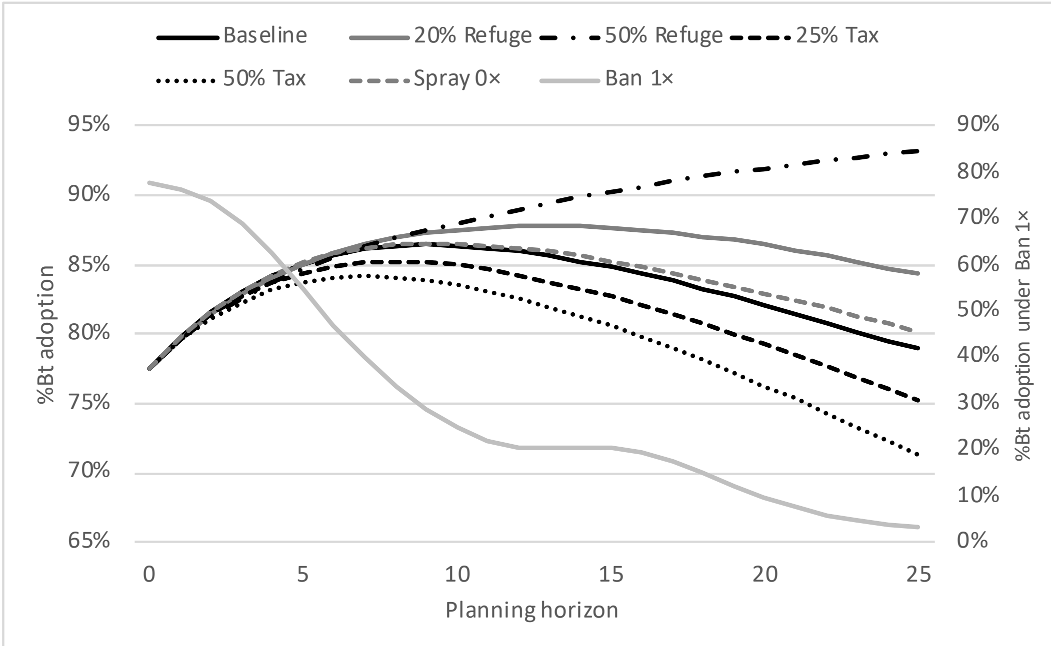

Figure 4.5: Bt adoption rate under single policies plotted against the planning horizon. The results for each period are averages over 1,000 simulations.

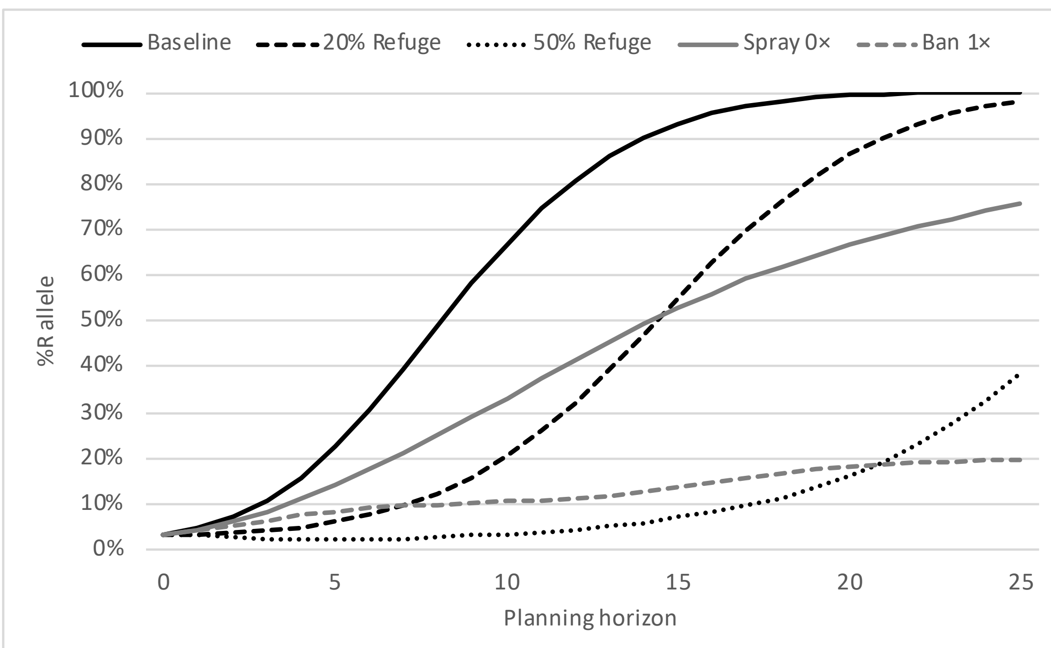

Figure 4.6: R-allele frequency under single policies plotted against the planning horizon. The results for each period are averages over 1,000 simulations.

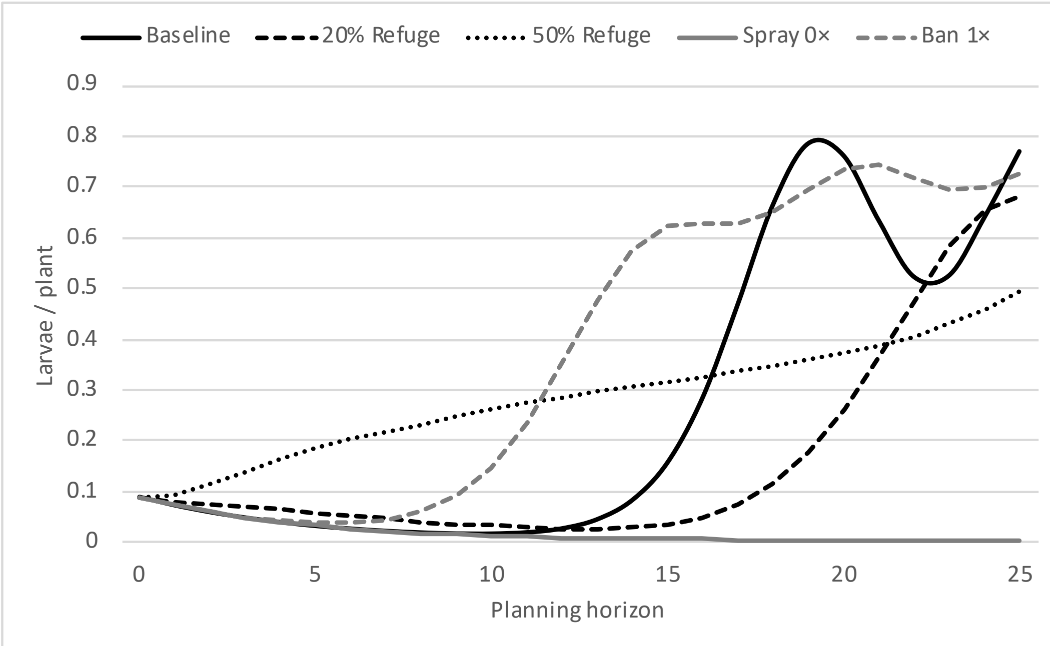

Figure 4.7: Insect population density under simple policies plotted against the planning horizon. The results for each period are averages over 1,000 simulations.

In Figure 4.6 (R-allele frequency), results for both tax policies are omitted as they were almost identical to the results for the baseline, suggesting low policy efficacy. This result was surprising because Bt adoption differs noticeably for these policies (Figure 4.5). All mitigation policies plotted in Figure 4.6 slowed the development of resistance compared to the baseline. The most effective mitigation policies were the 50% refuge for all farms (50% Refuge) and a localized ban on and around fields with resistance (Ban 1x), both of which kept the R-allele frequency below 20% for more than 20 years. By horizon period 25, however, the 50% refuge policy showed a rapid increase in the R-allele frequency, suggesting its failure, while the ban policy kept the frequency below 20%, suggesting that it was the most effective policy for mitigating resistance over the long-run (> 25 years). The 20% refuge for all farms (20% Refuge) effectively mitigated the resistance for about 10 years, and then the R-allele frequency began a rapid increase, reaching the baseline level by horizon period 25. The spray policies (Spray 0x) were not particularly effective for mitigating resistance, showing a steady increase in the R-allele frequency, though slower than for the baseline and tax policies.

In Figure 4.7 (pest population density), results for both tax policies are again omitted as they were almost identical to results for the baseline. The areawide spray policy (Spray 0x) kept the pest population density low over all 25 years, even with a radius = 0, due to the efficacy of the insecticide spray. The baseline with no intervention to mitigate resistance kept the pest population density low for about 15 years, and then the population increased and began to oscillate as expected. Surprisingly, the refuge policies showed distinctly different patterns over the 25 years. The 20% refuge policy (20% Refuge) kept the pest population low for about 20 years (about 5 years longer than the baseline), while the 50% refuge (50% Refuge) showed a long slowly increasing pest population density over all 25 years, exceeding the baseline in year 17 and the 20% refuge policy in year 23. Interestingly, the ban policy (Ban 1x) only kept the pest population low for about 10 years (about 5 years longer than the baseline).

These results showed the tradeoffs inherent in the mitigation of resistance. For example, the 50% refuge and ban policies were both the most effective at reducing the frequency of resistance alleles (Figure 4.6), but came at the cost of reduced adoption of Bt maize (Figure 4.5) and higher average pest populations (Figure 4.7), both implying lower benefits. Hence, we used economic surplus as a measure that integrates across costs and benefits in order to compare mitigation policies and to develop recommendations.

Figure 4.8 plots average annual economic surplus against the planning horizon, again omitting results for the tax policies as they were almost identical to baseline results. Each point on the curves in Figure 4.8 is an annualized average of the accumulated surplus over the corresponding planning horizon, i.e., the sum of landscape surplus over the planning horizon, divided by the number of years in the planning horizon and the maize planted area. As seen in Figure 4.8, the baseline (5% refuge, no localized ban, no areawide spray, no Bt seed tax) generated the greatest average annual surplus for all planning horizons up to 15 years. This result occurred because the surplus measure was cumulative, and the baseline policy accumulated more surplus during the early years than the other policies. However, for planning horizons of 16 or more years, the optimal mitigation policy increased the refuge requirement from 5% to 20% for all farms, but did not impose a localized ban, an areawide spray, or Bt seed tax. The areawide spray was suboptimal due to the additional costs incurred by farmers, while the ban policy was sub-optimal due to the loss of Bt maize benefits for farmers and the lost revenue for the seed company. Interestingly, the 50% refuge policy generated the lowest economic surplus—though it was one of the most effective mitigation policies, its cost in terms of lost benefits to farmers was too high. Recall that results for the two omitted tax policies (25% Tax, 50% Tax) were almost identical to the baseline.

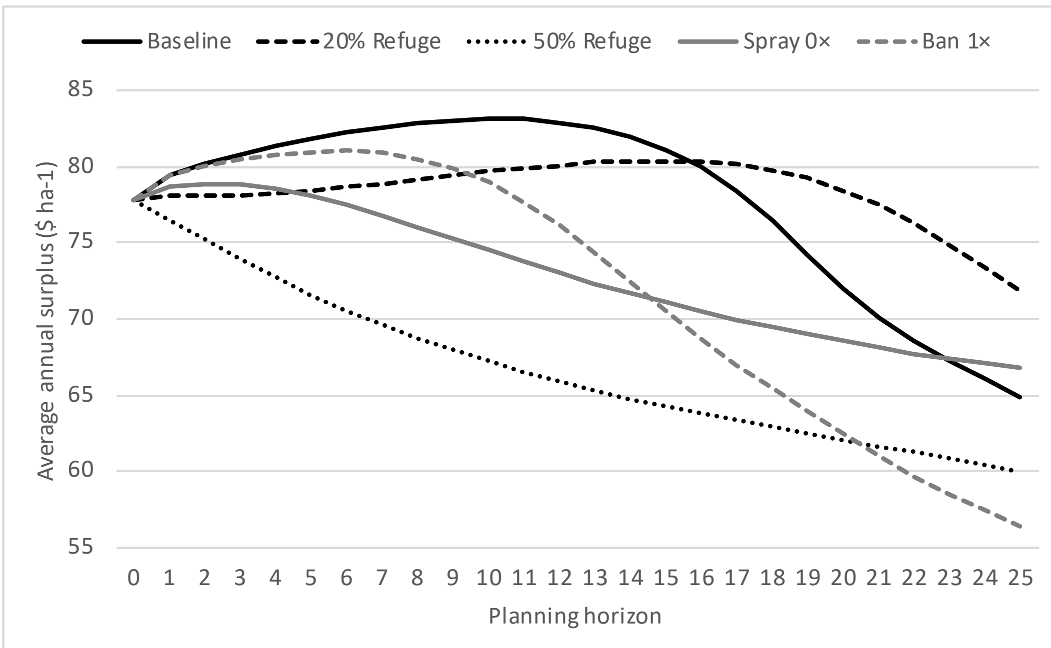

Figure 4.8: Average annual surplus for different mitigation policies plotted against the planning horizon.

Because the results in Figure 4.8 did not include combinations of mitigation policies, Table 4.1 summarizes results over the 81 policy combinations evaluated. Three combinations emerged as optimal for some length of planning horizon. For a planning horizon of 1 to 15 years, the baseline policy (5% refuge, no localized ban, no areawide spray, no Bt seed tax) continued to be optimal even as resistance increased. For a planning horizon of 16 to 22 years, the optimal policy increased the refuge requirement from 5% to 20%, but did not impose a localized ban, an areawide spray, or Bt seed tax. For a planning horizon of 23 to 25 years, technically adding the 50% tax to the 20% refuge requirement was optimal, but the increase in economic surplus was trivial ($< 0.05\%$). Therefore, our economic surplus criterion suggested that the optimal resistance mitigation policy was no intervention if a shorter (≤ 15 years) planning horizon was used and, if a longer (≥ 16 years) planning horizon was used, increasing the required refuge to 20% for all farmers when resistance emerged. Because mitigation policies that increased the required refuge decreased current benefits to achieve increased future benefits, discounting implies that the 20% and 50% refuge policies would have generated less surplus than plotted in Figure 4.8. Calculations showed that, with a 13% or higher discount rate, the no-intervention baseline remained the optimal policy for all planning horizons less than or equal to 25 years.

Table 4.1: Optimal policy combinations by length of planning horizon

Length of Planning Horizon	Refuge Requirement	Localized Ban	Areawide Spray	Bt Seed Tax
1-15	5%	None	None	0%
16-22	20%	None	None	0%
23-25	20%	None	None	50%

We also investigated the distribution of surplus shares under three resistance mitigation policies: the baseline (with a 5% refuge), a 20% refuge for all farmers planting Bt maize, and both a 20% refuge and a 50% Bt seed tax for all farmers planting Bt maize. Recall that economic surplus was the sum of farmer profit, the technology fee collected by the company and tax revenue and that adding the Bt seed tax as a mitigation policy had little impact on surplus with a 20% refuge requirement. For the baseline, farmers and the companies roughly divided the surplus evenly as yield gains and technology fees (Figure 4.11). Increasing the refuge requirement from the baseline of 5% to 20% to mitigate resistance increased the company share of surplus by about 5 to 10 percentage points, with the farmer share falling to about 40% (Figure 4.11). Adding a 50% Bt seed tax on top of the 20% refuge requirement to mitigate resistance, the tax burden was borne more by the companies, with their share declining by about 15 percentage points to 45% of the surplus, while farmers received about 35% of the surplus and tax revenue accounts for about 20% of the surplus.

4.3.3 Role of social networks

To highlight the difference created by incorporating the effects of social networks, Figure 4.9 and Figure 4.10 show results with all parameters the same as for the baseline except that the model was recalibrated with social networks "shut off" by setting the parameter $q = 1$. In this case, farmers made Bt maize adoption decisions based only on their individual expected profitably, giving no weight to their neighbors' decisions. In terms of Bt maize adoption, without the effect of social networks, the farmer adoption rate grew faster first, but then slowed and eventually declined from period 43 onward (Figure 4.9). This result was explained by the lack of social network effects. Without them, profitable adoption by early adopters was not slowed by neighboring non-adopters. Similarly, as the technology became less effective due to resistance, profit-motivated dis-adoption of Bt maize was not slowed by neighbors' inertia. As a consequence of the lower usage of Bt maize, the R-allele frequency reached key levels later than for the baseline. Specifically, the R-allele frequency did not exceed 10% until period 42, 20% until period 45, 30% until period 46, 40% until period 47, and 50% until period 48, or about 7 years later than for the baseline. Hence, not including the effects of social networks on farmer adoption of Bt maize slowed the estimated evolution of resistance by about 7 years.

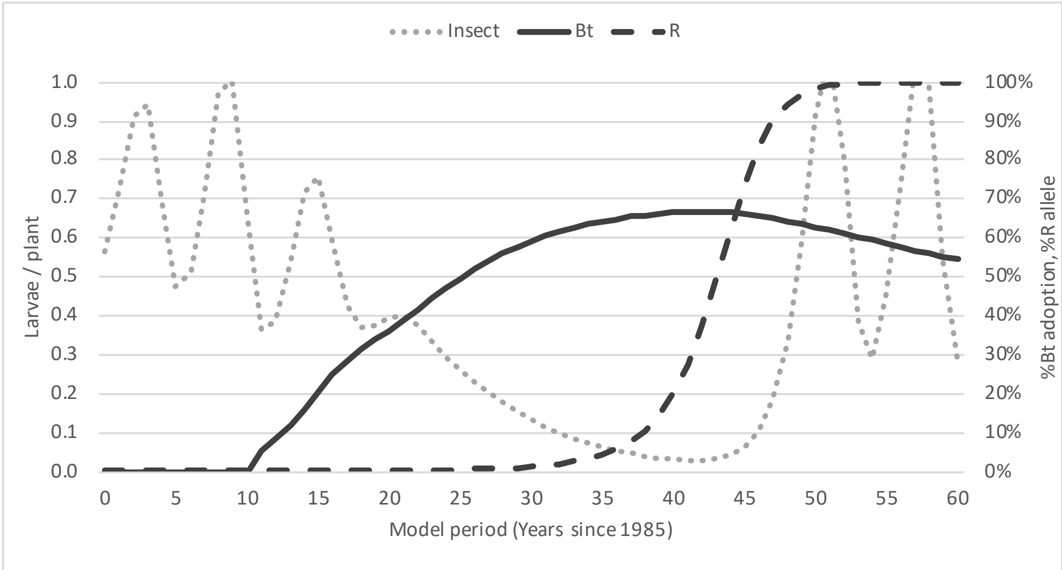

Figure 4.9: Results from the calibrated model without social network effects. It contains the insect population density (Insect), Bt seed adoption rate (Bt), and the resistance allele frequency (R) for the calibrated model without social network effects. The results for each period are averages over 1,000 simulations.

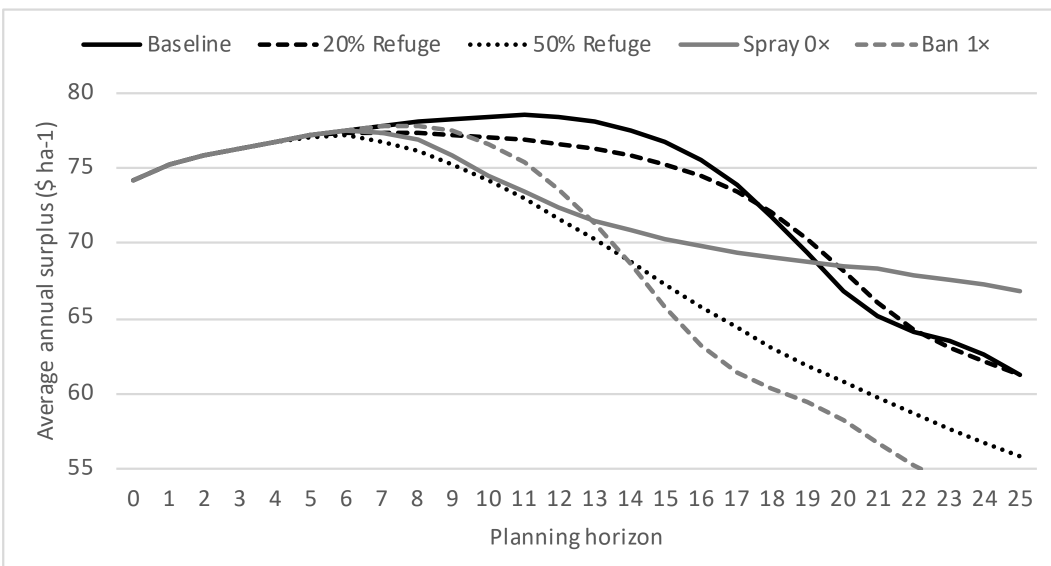

Figure 4.10: Average annual surplus for different mitigation policies. Each is plotted against the planning horizon for the calibrated model without social network effects.

Figure 4.10 plots average annual economic surplus against the planning horizon with the effect of social networks on adoption "shut off." Again, results for the tax policies were omitted as they were almost identical to baseline results. Compared to Figure 4.8, which incorporated the effects of social networks on adoption, Figure 4.10 shows that all mitigation policies generated essentially the same surplus for the first 6 or 7 years. Because mitigation policies were not implemented until the R-allele frequency exceeded a 50% threshold, the slower projected evolution of resistance without social networks effects delayed policy implementation, so that all policies were initially equivalent.

Based on Figure 4.10, the optimal policy depended on the planning horizon and varied from Figure 4.8. The baseline again generated the greatest average annual surplus for all planning horizons up to 17 years (about the same as in Figure 4.8). Again, the 20% refuge for all farmers was the optimal mitigation policy for longer planning horizons, but only for the narrow range from 18 to 20 years. Furthermore, the difference between the baseline (with 5% refuge) and the 20% refuge mitigation policy was much smaller than in Figure 4.8. However, just as in Figure 4.8, the 50% refuge policy generated among the lowest amounts of economic surplus. Interestingly, for planning horizons exceeding 20 years, the areawide spray policy became optimal, which did not occur in Figure 4.8. This result occurred because without social network effects, farmers more quickly dis-adopted Bt maize when resistance developed, thus avoiding the higher costs of the spray policy and lower Bt maize benefits, and so they generated higher surplus. Again, the ban policy generated the lowest economic surplus over many planning horizons, even with the more rapid dis-adoption of Bt maize when resistance developed.

Because the results in Figure 4.10 did not include combinations of mitigation policies, Table 4.2 summarizes results over the 81 policy combinations evaluated, just as Table 4.1 did for Figure 4.8. Without social network effects, farmer adoption of Bt maize only responded to individual profitability, which created some shifts in the optimal mitigation policy. Three policy combinations again emerged as optimal for some length of planning horizon. For a planning horizon of 1 to 17 years, the baseline policy continued to be optimal even as resistance increased, and for a planning horizon of 18 to 19 years, the optimal policy increased the refuge requirement from 5% to 20%. These were the same policies as when the effects of social networks were included, but the planning horizons changed to be slightly longer for the baseline policy and shorter for the 20% refuge policy (Table 4.1). The greatest change without social network effects was for the longest planning horizons. For a 20- to 25-year planning horizon, the optimal resistance mitigation policy was to reduce the refuge requirement back to 5% for all farmers, to add a 50% Bt maize seed tax on all farmers, and to make areawide insecticide applications in areas where resistance emerged. With social network effects and longer planning horizons of 22 to 25 years, the refuge remained at 20% and only the 50% Bt seed tax was added (Table 4.1). The greater responsiveness of farmers to individual profitability without social network effects made more active mitigation policies optimal, but only for longer planning horizons. However, the effect was not large, as again calculations showed that with a 9% or higher discount rate, the no-intervention baseline remained the optimal policy for all planning horizons less than 25 years.

Table 4.2: Optimal policy combination by length of planning horizon without social network effects

Length of Planning Horizon	Refuge Requirement	Localized Ban	Areawide Spray	Bt Seed Tax
1-17	5%	None	None	0%
18-19	20%	None	None	0%
20-25	5%	None	$r = 0$	50%

4.4 Discussion

Research on insect resistance mitigation strategies and empirical applications of agent-based models to pest and resistance management are limited (Gay et al., 2017; Renton, 2013; Renton et al., 2014). Hence, as part of this paper's first goal, we demonstrated the capacity of an agent-based model to produce results of use by policymakers and other stakeholders, specifically examining resistance mitigation policies for Bt maize and the European corn borer and the role of social networks in Bt maize adoption. We evaluated 81 resistance mitigation policies that combined three levels of four policies (non-Bt maize refuge, areawide non-Bt insecticide sprays, localized Bt maize bans, Bt maize seed tax) implemented when and/or where resistance emerged. These combinations showed variation in projected dynamics for Bt maize adoption, resistance allele frequency in the pest population, and the average pest population density.

From a biological perspective focused on keeping the frequency of resistance alleles in the pest population low, the most effective mitigation policies were a 50% refuge requirement for all farms when resistance emerged on the landscape and a localized ban on planting Bt maize within one radius of adult dispersal of farms with resistance. Bringing a broader economic perspective that balanced costs and benefits, we used economic surplus (the sum of farmer profit from maize production, company technology fees from selling Bt maize seed, and any tax revenue collected) to identify recommended mitigation policies. Surprisingly, results showed that when resistance emerged, the optimal response in terms of maximizing economic surplus was making no policy changes, but continuing the current resistance management policy of 5% non-Bt refuge, with no requirement of insecticidal sprays or localized Bt maize bans in and around areas of resistance, or Bt seed taxes when resistance emerges. For planning horizons beyond 16 years it became optimal to increase refuge requirements to 20% for all farmers when resistance developed. Furthermore, for planning horizons beyond 22 years it became optimal to add a 50% tax on all Bt maize seed sold when resistance developed in addition to the 20% refuge requirement. These results show the impacts that incorporating broader social science perspectives into resistance management or mitigation can have on recommended policy responses.

Several caveats apply to these results, as models cannot avoid the fundamental trade-off between fidelity to the phenomenon examined and abstraction for ease of analysis and interpretation. Baseline results assume one single-toxin Bt maize producing a high dose of the toxin. However, multiple single-toxin Bt maize hybrids with different modes of action have been commercialized in the US, and single-toxin Bt maize hybrids have been phased out as companies have shifted to Bt maize hybrids with multiple, pyramided traits (Hurley and Sun, 2019). Furthermore, refuge requirements in the Midwest have changed over time for the different Bt maize hybrids. Initial requirements were for 20% non-Bt maize as structured refuge, but more recently, some Bt maize hybrids with pyramided traits have a 5% or 10% refuge requirement implemented as a seed mix and/or structured refuge (Hurley and Sun, 2019; U.S. Environmental Protection Agency, 2019). Our model does not capture the use of multiple toxins entering the market at different times, overlapping use of hybrids with multiple, pyramided traits at the same time by neighboring farmers, or changes in refuge requirements and methods of implementation. In addition, our model assumes that Bt maize delivers a high dose of the toxin, which is accurate for European corn borer, but not for other lepidopteran pests such as corn earworm (*Helicoverpa zea*) or Bt maize for corn rootworm (Burkness et al., 2010; Gassmann et al., 2014; U.S. Environmental Protection Agency, 2018). Furthermore, the model focuses on a single pest, though farmers simultaneously manage multiple pests with varying levels of control by different Bt maize hybrids (Mason et al., 1996). In addition, the model assumes a single selection by the Bt toxin each year, while many target pests, including the European corn borer, have multiple generations per season with more than one selection event by Bt maize (Mason et al., 1996). Also, economic surplus is not a complete measure of social benefits (Mitchell et al., 2018). For example, as used here, it does not include environmental impacts of insect management, even though a significant benefit of Bt maize is that farmers use it as a substitute for conventional insecticides (Brookes, 2019; Catarino et al., 2019; Perry et al., 2016).

With these caveats, the policy experiments reported here suggest that refuge requirements remain the foundation of resistance mitigation for high-dose technologies, just as they are for resistance management. Based on maximizing social surplus, the optimal policy to mitigate resistance when it emerges was to maintain the current refuge requirement or to modestly increase it for all farmers, rather than to implement localized bans on the sale of Bt maize in areas where resistance develops or to make areawide applications of insecticidal sprays on resistant populations. Based on the economic surplus criterion, the benefits of lower resistance allele frequency for these policies did not adequately compensate for the added costs or loss of the benefits from using Bt maize. Taxes on the sale of Bt maize seed did not cause surplus to differ substantially from the baseline policy, suggesting a possible mechanism to fund various programs to improve Bt maize use, such as development of educational materials and outreach or research activities. However, the results showed that companies bear a large share of these costs, suggesting that it would be more efficient for companies to directly fund these programs based on seed sales rather than creating a seed tax program to fund them. Also, as a caveat, this model did not incorporate the Bt technology market. As a result, for example, the model did not include market competition among companies via differentiated traits, including different regulatory requirements, as, for example, companies would lobby to not have their hybrids included in tax schemes if resistance developed to a competitor's Bt maize.

As a secondary goal of this paper, we demonstrated that social factors can play key roles in the development and management of insect resistance, focusing on the effect that social networks can play on farmer adoption of Bt maize. As modeled here, adoption depended in part on the average adoption of a farmer's neighbors, not just each farmer's expected profitability, as a way to capture the effects of information exchange, integration of multiple farmers' experiences with pests and adoption, shared local institutions and markets, and similar factors. Modeling the mechanisms for this social network and the specific connections among individual farmers is beyond the scope of this analysis. Relative to a model in which farmers responded only to their individual profitability, social networks as modeled here impeded farmer responsiveness to profitability signals, which slowed the initial adoption of Bt maize and its dis-adoption as pest populations declined or resistance developed. Model calibration to observed state-level adoption rates identified model parameters and reduced differences in initial adoption rates with and without social networks. However, this calibrated model implied a relatively slower adjustment in Bt maize use by farmers. As a result, when including social network effects, our model projected that resistance develops about 7 years earlier than without social network effects and the optimal mitigation policy more strongly favored use of moderate increases in refuge for all farmers. Resistance developed earlier because farmers uses Bt maize more intensely since they did not dis-adopt Bt maize as pest populations declined and resistance developed, even though the profitability of Bt maize decreased. With social network effects, the optimal resistance mitigation policy also more strongly favored use of modest increases in refuge because more farmers continued to use Bt maize and obtain it benefits relative to more costly, but effective, policies such as areawide sprays or localized bans. In this example, ignoring social network effects could contribute to making inappropriate policy recommendations for managing pest resistance or mitigating it once it develops.

The intensity and extent of farmer adoption of Bt maize plays a key role in the management and mitigation of pest resistance. This agent-based model incorporated the influence of social factors by having individual farmer adoption respond to expected profitability and the adoption behavior of neighboring farmers. However, many social factors not addressed by this model also affect adoption. For example, expected profitability depends not only on all the market factors driving maize prices, but also technology markets and the pricing behavior of firms selling Bt maize (Shi et al., 2010). Similarly, farmers adopt Bt maize not only for expected profit, but also to manage income risk (Hurley et al., 2004; Shi et al., 2013). Also, social networks for agricultural management rarely have the simple spatial structure assumed here, but typically have varying nodes of importance such as key crop consultants, retailers, and extension agents (Kaup, 2008; Lubell and Fulton, 2007). In addition, social factors affect resistance through more than just adoption of Bt maize, such as through farmer compliance with resistance management and mitigation practices and how Bt maize affects broader cropping systems such as crop rotations (Hurley and Mitchell, 2014; Mitchell and Onstad, 2005). Overall, our results demonstrate that social factors can play an important role in resistance management and mitigation. However, more applied and quantitative social science research would contribute to developing better policy recommendations for resistance management and mitigation, and agent-based models can be a part of this contribution.

Appendix

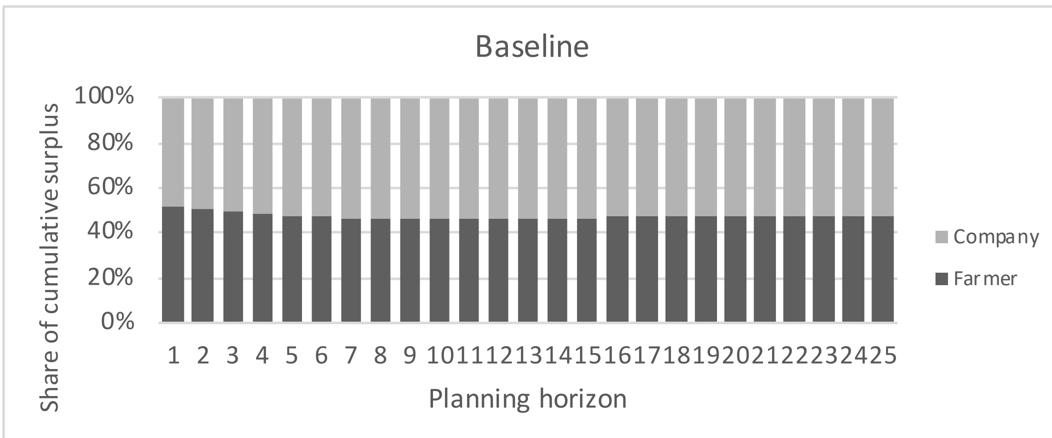

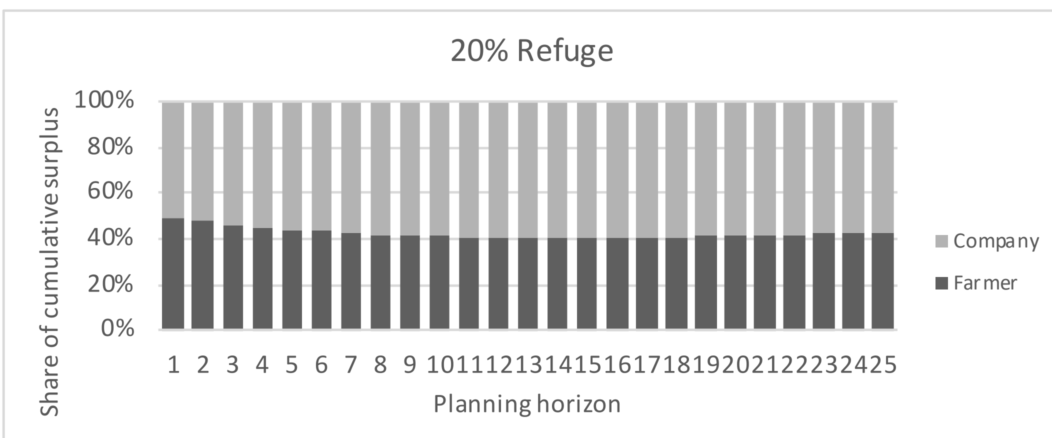

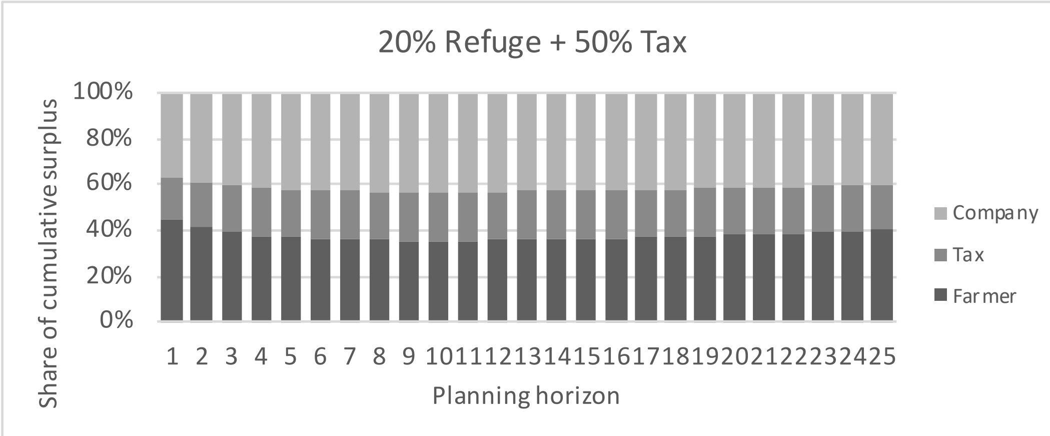

Figure 4.11: Cumulative share of surplus by planning horizon. The top panel is for the baseline, the middle is for the 20% refuge mitigation policy, and the bottom is for the 20% refuge mitigation policy combined with a 50% Bt seed tax.

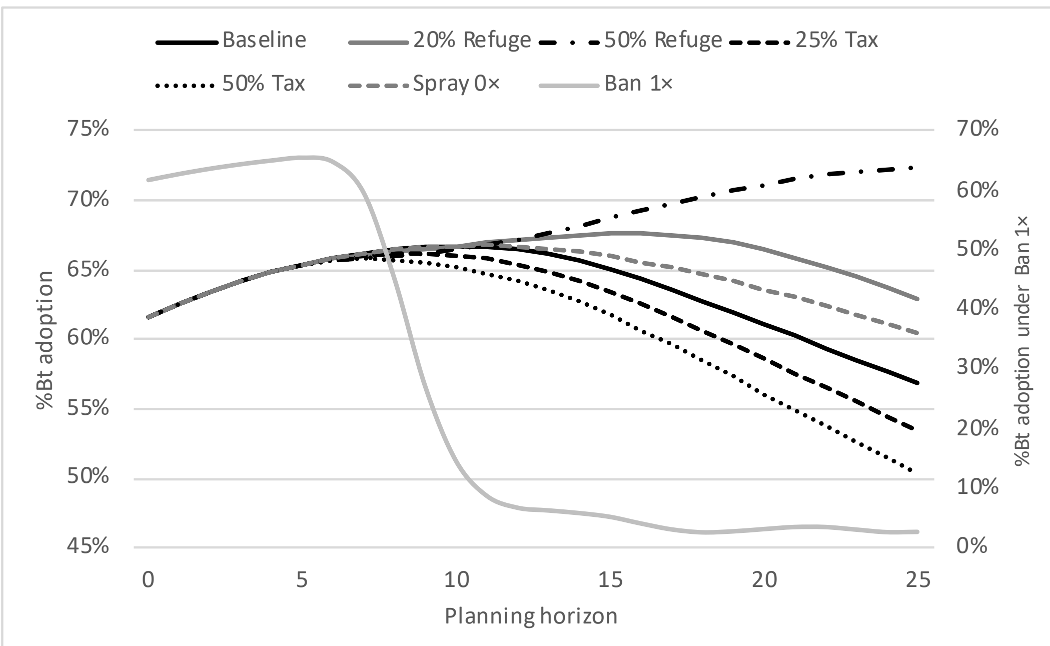

Figure 4.12: Bt adoption rate under simple policies. Each is plotted against the planning horizon without social network effects. The results for each period are averages over 1,000 simulations.

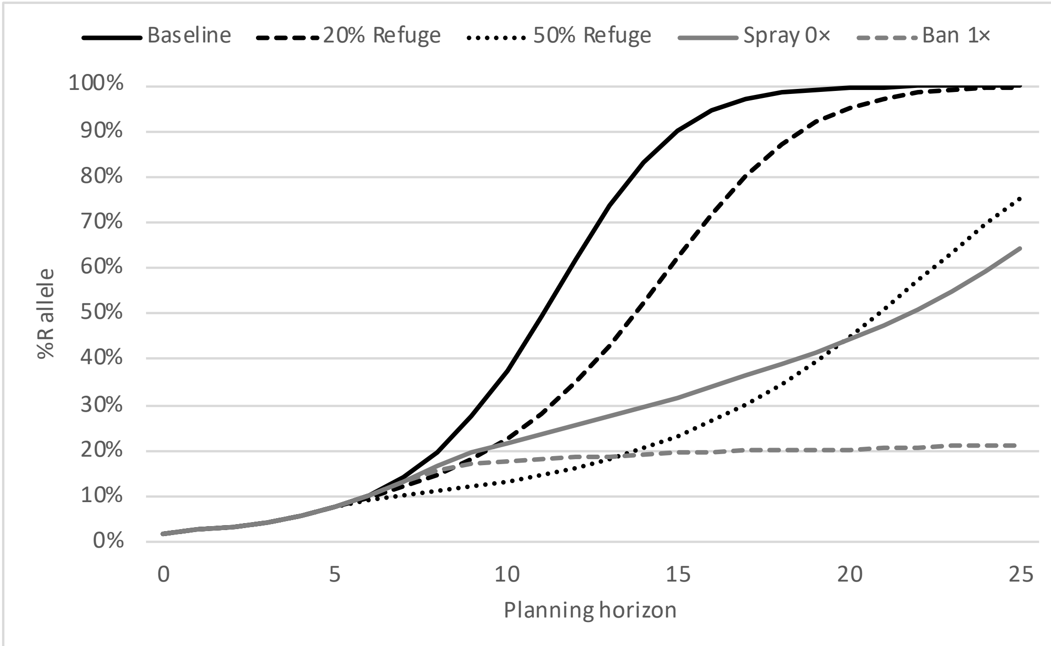

Figure 4.13: R-allele frequency under simple policies. Each is plotted against the planning horizon without social network effects. The results for each period are averages over 1,000 simulations.

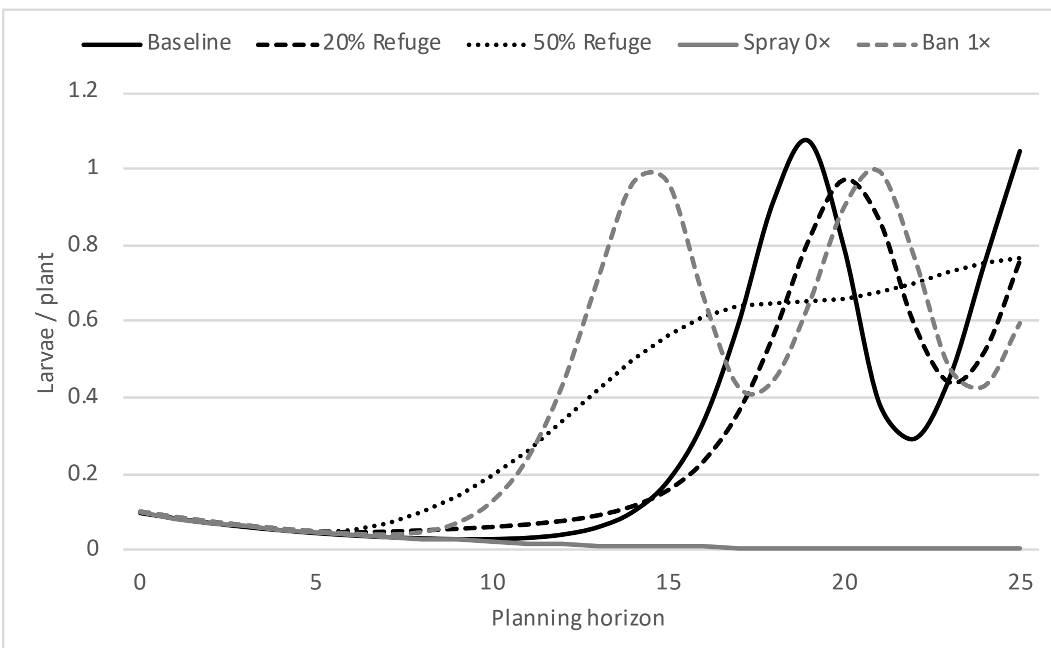

Figure 4.14: Insect population density under simple policies. Each is plotted against the planning horizon without social network effects. The results for each period are averages over 1,000 simulations.

Chapter 5

Conclusions

This dissertation has presented computational models to address complex practical problems in agriculture. The work is motivated by the critical situation surrounding modern agriculture—despite the urgent need for sustainable intensification across the world, the researcher-centricity prevails in academic disciplines yielding research outputs largely irrelevant for real-world complex problems. The models serve two purposes: to demonstrate potential approaches to addressing complex agricultural problems in general and to lay the foundation for selected specific problems.

These models are by no means general solutions; instead, these models are constructed ad hoc for specific problems. This is the point. The current state of knowledge about agricultural systems is inadequate to manage their complexities for practical purposes, while sustainable intensification is an acute problem demanding immediate actions for productivity increase. Although it would be ideal to deductively design solutions based on the first principles, there is no such principle to rely on in agricultural sciences and economics today. Therefore, to address complex practical problems, we must be content with ad hoc approaches for the time being. Fortunately, using the emerging data streams and ever increasing computational capacity, we can build effective empirical models for complex problems. Although the models presented in this dissertation are ad hoc solutions to specific problems, their modeling techniques (Bayesian optimization and agent-based modeling) are very general and applicable to many other practical problems.

Bayesian optimization is a data-driven optimization technique, featuring sample efficiency and flexibility to handle arbitrary objective functions. These characteristics make the technique suitable for optimization in complex systems, including agricultural systems, where objective functions are formulated essentially unknown, and effective trade-off between exploration and exploitation is crucial. While Chapter 2 and 3 have demonstrated applications to static problems, Bayesian optimization can be useful for dynamic problems as well. A common goal in dynamic optimization is to determine an optimal policy, which prescribes an action to each state so that a cumulative reward over time is maximized. Since learning of an optimal policy is guided by maximizing some performance metric, Bayesian optimization can be applied to this internal optimization task. The sample efficiency of Bayesian optimization will be even more beneficial because brute force search is infeasible in realistic dynamic optimization problems.

In contrast to Bayesian optimization and machine learning in general, agent-based modeling plays a unique role in addressing complex agricultural problems. Despite the ever increasing data available in agriculture at large, suitable one for data-driven dynamical systems and control is still lacking in many cases. Agent-based modeling provides a viable option to conduct analyses and gain insights into complex agricultural systems even without much data. Basically, this is achieved by taking advantage of the modular nature of agent-based modeling and imposing assumptions on each component sub-process. Assumptions are typically in the form of parameters and functional relations, which are derived from domain knowledge in relevant disciplines. While the capability of this kind is certainly a strength of agent-based modeling, it is also a danger of undisciplined realism. As mentioned as caveats in Section 4.4, our model is wrong in many aspects of the reality. We could address some or even all of those caveats by making each sub-process and assumption more realistic. However, once again, "*All models are wrong, but some are useful.*" So, whether it is for understanding mechanisms or solving practical problems, a blind pursuit of realism is never a good principle of modeling as it would be wrong anyway but merely inflating model complexities. Instead, every model should serve specific purposes in order to be useful. In agent-based modeling, we should be particularly mindful of such purposes and guard against the temptation of realism. In our case, bringing in multi-trait hybrid seeds and market structure would make the model more realistic but no more useful for illustrating the significance of spatiotemporal dynamics and social factors in designing effective policies for insect resistance management. Ultimately, with purposes firmly kept in mind, it is a trial-and-error process; to discover a useful model, we must build many models and test them to see whether they serve the purposes.

It is my hope that the emphasis on practicality will be shared in the agricultural research community so that we can achieve a better balance between scientific pursuit of understanding mechanisms and practical problem-solving in complex agricultural systems.

Bibliography

Albornoz, Enrique M., Alejandra C Kemerer, Romina Galarza, Nicolás Mastaglia, Ricardo Melchiori, and César E Martínez (2018). "Development and evaluation of an automatic software for management zone delineation". *Precision Agriculture* 19.3, pp. 463–476.

Altieri, Miguel A (2018). *Agroecology: The science of sustainable agriculture*. 2nd. CRC Press.

Andow, David A., Steven G. Pueppke, Arthur W. Schaafsma, Aaron J. Gassmann, Thomas W. Sappington, Lance J. Meinke, Paul D. Mitchell, Terrance M. Hurley, Richard L. Hellmich, and R. Pat Porter (2016). "Early Detection and Mitigation of Resistance to Bt Maize by Western Corn Rootworm (Coleoptera: Chrysomelidae)". *Journal of Economic Entomology* 109.1, pp. 1–12.

Anselin, Luc, Rodolfo Bongiovanni, and Jess Lowenberg-DeBoer (2004). "A spatial econometric approach to the economics of site-specific nitrogen management in corn production". *American Journal of Agricultural Economics* 86.3, pp. 675–687.

Antle, John M (2019). "Data, Economics and Computational Agricultural Science". *American Journal of Agricultural Economics* 101.2, pp. 365–382.

Antle, John M, Bruno Basso, Richard T Conant, H Charles J Godfray, James W Jones, Mario Herrero, Richard E Howitt, Brian A Keating, Rafael Munoz-Carpena, Cynthia Rosenzweig, Pablo Tittonell, and Tim R Wheeler (2017a). "Towards a new generation of agricultural system data, models and knowledge products: Design and improvement". *Agricultural Systems* 155, pp. 255–268.

Antle, John M and Susan M Capalbo (2001). “Econometric-process models for integrated assessment of agricultural production systems”. *American Journal of Agricultural Economics* 83.2, pp. 389–401.

Antle, John M and Susan M Capalbo (2002). “Agriculture as a managed ecosystem: policy implications”. *Journal of Agricultural and Resource Economics*, pp. 1–15.

Antle, John M, James W Jones, and Cynthia E Rosenzweig (2017b). “Next generation agricultural system data, models and knowledge products: Introduction”. *Agricultural Systems* 155, pp. 186–190.

Arthur, Brian W. (2013). “Complexity economics: A different framework for economic thought”. *Complexity and the Economy*. Oxford University Press.

Auernhammer, Hermann (2001). “Precision farming—the environmental challenge”. *Computers and electronics in agriculture* 30.1, pp. 31–43.

Bachmaier, Martin and Markus Gandorfer (2009). “A conceptual framework for judging the precision agriculture hypothesis with regard to site-specific nitrogen application”. *Precision agriculture* 10.2, p. 95.

Bell, James R, Eric C Burkness, Alice E Milne, David W Onstad, Mark Abrahamson, Krista L Hamilton, and William D Hutchison (2012). “Putting the brakes on a cycle: bottom-up effects damp cycle amplitude: Managing epidemic persistence”. *Ecology Letters* 15.4, pp. 310–318.

Bellman, Richard E (2015). *Adaptive control processes: a guided tour*. Vol. 2045. Princeton university press.

Berk, Julian, Vu Nguyen, Sunil Gupta, Santu Rana, and Svetha Venkatesh (2018). “Exploration Enhanced Expected Improvement for Bayesian Optimization”. Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Springer, pp. 621–637.

Bongiovanni, Rodolfo and Jess Lowenberg-DeBoer (2004). “Precision agriculture and sustainability”. *Precision agriculture* 5.4, pp. 359–387.

Box, George E.P. (1976). "Science and statistics". *Journal of the American Statistical Association* 71.356, pp. 791–799.

Box, George E.P. (1979). "Robustness in the strategy of scientific model building". *Robustness in statistics.* Elsevier, pp. 201–236.

Boyer, Christopher N., B. Wade Brorsen, John B. Solie, and William R. Raun (2011). "Profitability of variable rate nitrogen application in wheat production". *Precision Agriculture* 12.4, pp. 473–487.

Bramley, R. G. V. (2009). "Lessons from nearly 20 years of Precision Agriculture research, development, and adoption as a guide to its appropriate application". *Crop and Pasture Science* 60.3, pp. 197–217.

Bramley, R. G. V., KJ Evans, KJ Dunne, and DL Gobbett (2011). "Spatial variation in response to 'reduced input' spray programs for powdery mildew and botrytis identified through whole-of-block experimentation". *Australian Journal of Grape and Wine Research* 17.3, pp. 341–350.

Bramley, R. G. V. and J. Ouzman (2019). "Farmer attitudes to the use of sensors and automation in fertilizer decision-making: nitrogen fertilization in the Australian grains sector". *Precision Agriculture* 20.1, pp. 157–175.

Brochu, Eric, Vlad M Cora, and Nando De Freitas (2010). "A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning". *arXiv preprint arXiv:1012.2599.*

Brookes, Graham (2019). "Twenty-one years of using insect resistant (GM) maize in Spain and Portugal: farm-level economic and environmental contributions". *GM crops & food*, pp. 1–12.

Bullock, David S, Maria Boerngen, Haiying Tao, Bruce Maxwell, Joe D Luck, Luciano Shiratsuchi, Laila Puntel, and Nicolas F Martin (2019). "The data-intensive farm management project: Changing agronomic research through On-farm precision experimentation". *Agronomy Journal.*

Bullock, David S and Donald G Bullock (2000). "From agronomic research to farm management guidelines: A primer on the economics of information and precision technology". *Precision Agriculture* 2.1, pp. 71–101.

Bullock, David S, Jess Lowenberg-DeBoer, and Scott M Swinton (2002). "Adding value to spatially managed inputs by understanding site-specific yield response". *Agricultural Economics* 27.3, pp. 233–245.

Buresh, Roland J., Rowena L. Castillo, Judith Carla Dela Torre, Eufrocino V. Laureles, Marianne I. Samson, Philip Joshua Sinohin, and Marlon Guerra (2019). "Site-specific nutrient management for rice in the Philippines: Calculation of field-specific fertilizer requirements by Rice Crop Manager". *Field Crops Research* 239, pp. 56–70.

Burke, Marshall and David B Lobell (2017). "Satellite-based assessment of yield variation and its determinants in smallholder African systems". *Proceedings of the National Academy of Sciences* 114.9, pp. 2189–2194.

Burkness, Eric C, Galen Dively, Terry Patton, Amy C Morey, and William D Hutchison (2010). "Novel Vip3A Bacillus thuringiensis (Bt) maize approaches high-dose efficacy against Helicoverpa zea (Lepidoptera: Noctuidae) under field conditions: Implications for resistance management". *GM crops* 1.5, pp. 337–343.

Camberato, Jim (2014). *Low soil pH and limestone recommendations for mineral soils: The confusion of 'Buffer pH'*. URL: https://ag.purdue.edu/agry/extension/SiteAssets/soilfertilityassets/Buffer-pH.pdf.

Camberato, Jim and R.L. Nielsen (2017). *Soil Sampling to Assess Current Soil N Availability.* Purdue University. URL: https://www.agry.purdue.edu/ext/corn/news/timeless/assessavailablen.html.

Capalbo, Susan M, John M Antle, and Clark Seavert (2017). "Next generation data systems and knowledge products to support agricultural producers and science-based policy decision making". *Agricultural Systems* 155, pp. 191–199.

Cassman, K. G. (1999). "Ecological intensification of cereal production systems: Yield potential, soil quality, and precision agriculture". *Proceedings of the National Academy of Sciences* 96.11, pp. 5952–5959.

Castle, Michael H, Bradley D Lubben, and Joe D Luck (2016). *Factors Influencing Producer Propensity for Data Sharing & Opinions Regarding Precision Agriculture and Big Farm Data*. UNL Digital Commons.

Castle, Michael H, Bradley D Lubben, Joe D Luck, and Taro Mieno (2017). *Precision Agriculture Adoption and Profitability*. Agricultural Economics, University of Nebraska–Lincoln. URL: https://agecon.unl.edu/cornhusker-economics/2017/precision-agriculture-adoption-profitability.

Catarino, Rui, Francisco Areal, Julian Park, and Nicolas Parisey (2019). "Spatially explicit economic effects of non-susceptible pests' invasion on Bt maize". *Agricultural Systems* 175, pp. 22–33.

Cela, Sebastián, Francisca Santiveri, and Jaime Lloveras (2011). "Optimum nitrogen fertilization rates for second-year corn succeeding alfalfa under irrigation". *Field Crops Research* 123.2, pp. 109–116.

Chavas, Jean-Paul (2004). *Risk analysis in theory and practice*. Academic Press.

Chen, Zhao, Xv Liu, Junpeng Niu, Wennan Zhou, Tian Zhao, Wenbo Jiang, Jian Cui, Robert Kallenbach, and Quanzhen Wang (2019). "Optimizing irrigation and nitrogen fertilization for seed yield in western wheatgrass [Pascopyrum smithii (Rydb.) Á. Löve] using a large multi-factorial field design". *PLOS ONE* 14.6. Ed. by Shahid Farooq, e0218599.

Cid-Garcia, Nestor M., Angel G. Bravo-Lozano, and Yasmin A. Rios-Solis (2014). "A crop planning and real-time irrigation method based on site-specific management zones and linear programming". *Computers and Electronics in Agriculture* 107, pp. 20–28.

Cook, Simon and R. G. V. Bramley (1998). "Precision agriculture—opportunities, benefits and pitfalls of site-specific crop management in Australia". *Australian Journal of Experimental Agriculture* 38.7, pp. 753–763.

Cook, Simon, James Cock, Thomas Oberthür, and Myles Fisher (2013). "On-farm experimentation". *Better Crops* 97.4, pp. 17–20.

Crowder, DW, DW Onstad, ME Gray, PD Mitchell, JL Spencer, and RJ Brazee (2005). "Economic analysis of dynamic management strategies utilizing transgenic corn for control of western corn rootworm (Coleoptera: Chrysomelidae)". *Journal of economic entomology* 98.3, pp. 961–975.

Dai, Xinglong, Xiaohu Zhou, Dianyong Jia, Lili Xiao, Haibo Kong, and Mingrong He (2013). "Managing the seeding rate to improve nitrogen-use efficiency of winter wheat". *Field Crops Research* 154, pp. 100–109.

Donatelli, M., R.D. Magarey, S. Bregaglio, L. Willocquet, J.P.M. Whish, and S. Savary (2017). "Modelling the impacts of pests and diseases on agricultural systems". *Agricultural Systems* 155, pp. 213–224.

Dorhout, David L., Thomas W. Sappington, L. C. Lewis, and Marlin E. Rice (2011). "Flight behaviour of European corn borer infected with Nosema pyrausta". *Journal of Applied Entomology* 135.1, pp. 25–37.

Dorhout, David L., Thomas W. Sappington, and Marlin E. Rice (2008). "Evidence for obligate migratory flight behavior in young European corn borer (Lepidoptera: Crambidae) females". *Environmental entomology* 37.5, pp. 1280–1290.

Duffy, Michael (2013). *Estimated Costs of Crop Production in Iowa - 2013*. Estimated Costs of Crop Production in Iowa — Ag Decision Maker. URL: http://econ2.econ.iastate.edu/faculty/duffy/documents/EstimatedCostsofCropProduction2013.pdf.

Durlauf, Steven N. (2012). "Complexity, economics, and public policy". *Politics, Philosophy & Economics* 11.1, pp. 45–75.

Easley, David and Jon Kleinberg (2010). *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press.

Edwards, Jeffrey T and Larry C Purcell (2005). “Soybean yield and biomass responses to increasing plant population among diverse maturity groups”. *Crop Science* 45.5, pp. 1770–1777.

Elmore, Roger W (2013). *Corn Planting FAQs — Integrated Crop Management*. URL: https://crops.extension.iastate.edu/cropnews/2013/04/corn-planting-faqs.

Epstein, Joshua M. (2006). *Generative social science: Studies in agent-based computational modeling*. Princeton University Press.

Farnham, Dale (2001). *Corn Planting Guide*.

Feola, Giuseppe and Claudia R. Binder (2010). “Towards an improved understanding of farmers' behaviour: The integrative agent-centred (IAC) framework”. *Ecological Economics* 69.12, pp. 2323–2333.

Filatova, Tatiana, Peter H. Verburg, Dawn Cassandra Parker, and Carol Ann Stannard (2013). “Spatial agent-based models for socio-ecological systems: Challenges and prospects”. *Environmental Modelling & Software*. Thematic Issue on Spatial Agent-Based Models for Socio-Ecological Systems 45, pp. 1–7.

Food and Agriculture Organization of the United Nations (2016). *Sustainable development goals*. URL: http://www.fao.org/sustainable-development-goals/en/.

Fraisse, CW, KA Sudduth, and NR Kitchen (2001). “Delineation of site-specific management zones by unsupervised classification of topographic attributes and soil electrical conductivity”. *Transactions of the ASAE* 44.1, p. 155.

Fridgen, Jon J, Newell R Kitchen, Kenneth A Sudduth, Scott T Drummond, William J Wiebold, and Clyde W Fraisse (2004). “Management Zone Analyst (MZA): Software for Subfield Management Zone Delineation”. *Agronomy Journal* 96, pp. 100–108.

Gandorfer, Markus and Andreas Meyer-Aurich (2017). “Economic Potential of Site-Specific Fertiliser Application and Harvest Management”. *Precision Agriculture: Technology and Economic Perspectives*. Ed. by Søren Marcus Pedersen and Kim Martin Lind. Springer International Publishing, pp. 79–92.

Garnett, Tara, Michael C Appleby, Andrew Balmford, Ian J Bateman, Tim G Benton, Phil Bloomer, Barbara Burlingame, Marian Dawkins, Liam Dolan, and David Fraser (2013). “Sustainable intensification in agriculture: premises and policies”. *Science* 341.6141, pp. 33–34.

Gaspar, Adam P, Paul D Mitchell, and Shawn P Conley (2015). “Economic risk and profitability of soybean fungicide and insecticide seed treatments at reduced seeding rates”. *Crop Science* 55.2, pp. 924–933.

Gassmann, Aaron J., Jennifer L. Petzold-Maxwell, Ryan S. Keweshan, and Mike W. Dunbar (2011). “Field-Evolved Resistance to Bt Maize by Western Corn Rootworm”. *PLOS ONE* 6.7, e22629.

Gassmann, Aaron J, Jennifer L Petzold-Maxwell, Eric H Clifton, Mike W Dunbar, Amanda M Hoffmann, David A Ingber, and Ryan S Keweshan (2014). “Field-evolved resistance by western corn rootworm to multiple Bacillus thuringiensis toxins in transgenic maize”. *Proceedings of the National Academy of Sciences* 111.14, pp. 5141–5146.

Gay, Pierre-Emmanuel, Michel Lecoq, and Cyril Piou (2017). “Improving preventive locust management: insights from a multi-agent model”. *Pest Management Science* 74.1, pp. 46–58.

Gebbers, Robin and Viacheslav I Adamchuk (2010). “Precision agriculture and food security”. *Science* 327.5967, pp. 828–831.

Gili, Adriana, Cristian Álvarez, Ramiro Bagnato, and Elke Noellemeyer (2017). “Comparison of three methods for delineating management zones for site-specific crop management”. *Computers and Electronics in Agriculture* 139, pp. 213–223.

Ginsbourger, David, Rodolphe Le Riche, and Laurent Carraro (2010). “Kriging is well-suited to parallelize optimization”. *Computational intelligence in expensive optimization problems*. Springer, pp. 131–162.

Gliessman, Stephen R. (1990). "Agroecology: Researching the Ecological Basis for Sustainable Agriculture". *Agroecology: Researching the Ecological Basis for Sustainable Agriculture*. Ed. by Stephen R. Gliessman. New York, NY: Springer New York, pp. 3–10.

Gomez, Kwanchai A and Arturo A Gomez (1984). *Statistical procedures for agricultural research*. John Wiley & Sons.

Griffin, Terry (2018). *Collating and analysing small data to make big decisions – Can it improve farm productivity and profitability?* Grain Research and Development Corporation. URL: https://grdc.com.au/resources-and-publications/grdc-update-papers/tab-content/grdc-update-papers/2018/02/collating-and-analysing-small-data-to-make-big-decisions.

Haghverdi, Amir, Brian G. Leib, Robert A. Washington-Allen, Michael J. Buschermohle, and Paul D. Ayers (2016). "Studying uniform and variable rate center pivot irrigation strategies with the aid of site-specific water production functions". *Computers and Electronics in Agriculture* 123, pp. 327–340.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2016). *The elements of statistical learning*. 2nd ed. New York: Springer.

He, Xiantao, Youqiang Ding, Dongxing Zhang, Li Yang, Tao Cui, and Xiangjun Zhong (2019a). "Development of a variable-rate seeding control system for corn planters Part I: Design and laboratory experiment". *Computers and Electronics in Agriculture* 162, pp. 318–327.

He, Xiantao, Youqiang Ding, Dongxing Zhang, Li Yang, Tao Cui, and Xiangjun Zhong (2019b). "Development of a variable-rate seeding control system for corn planters Part II: Field performance". *Computers and Electronics in Agriculture* 162, pp. 309–317.

Hodgson, Erin (2010). *Refuge in a Bag is Here: Explaining the Simplified Refuge*. URL: https://crops.extension.iastate.edu/cropnews/2010/06/refuge-bag-here-explaining-simplified-refuge.

Holzworth, Dean P, Neil I Huth, Peter G deVoil, Eric J Zurcher, Neville I Herrmann, Greg McLean, Karine Chenu, Erik J van Oosterom, Val Snow, and Chris Murphy (2014). "APSIM–evolution towards a new generation of agricultural systems simulation". *Environmental Modelling & Software* 62, pp. 327–350.

Huang, Fangneng, David A. Andow, and Lawrent L. Buschman (2011). "Success of the high-dose/refuge resistance management strategy after 15 years of Bt crop use in North America". *Entomologia Experimentalis et Applicata* 140.1, pp. 1–16.

Huang, Jianliang, Fan He, Kehui Cui, Roland J Buresh, Bo Xu, Weihua Gong, and Shaobing Peng (2008). "Determination of optimal nitrogen rate for rice varieties using a chlorophyll meter". *Field Crops Research* 105.1, pp. 70–80.

Hurley, Terrance M. and Paul D. Mitchell (2014). "Insect Resistance Management: Adoption and Compliance". *Insect Resistance Management*. Ed. by David W. Onstad. 2nd ed. San Diego: Academic Press, pp. 421–451.

Hurley, Terrance M., Paul D. Mitchell, and Marlin E. Rice (2004). "Risk and the value of Bt corn". *American Journal of Agricultural Economics* 86.2, pp. 345–358.

Hurley, Terrance M, Bruce A Babcock, and Richard L Hellmich (2001). "Bt corn and insect resistance: an economic assessment of refuges". *Journal of Agricultural and Resource Economics* 26.1, p. 176.

Hurley, Terrance M and Huichun Sun (2019). "Softening Shock and Awe Pest Management in Corn and Soybean Production with IPM Principles". *Journal of Integrated Pest Management* 10.1, p. 7.

Hutchison, W. D., E. C. Burkness, P. D. Mitchell, R. D. Moon, T. W. Leslie, S. J. Fleischer, M. Abrahamson, K. L. Hamilton, K. L. Steffey, M. E. Gray, R. L. Hellmich, L. V. Kaster, T. E. Hunt, R. J. Wright, K. Pecinovsky, T. L. Rabaey, B. R. Flood, and E. S. Raun (2010). "Areawide Suppression of European Corn Borer with Bt Maize Reaps Savings to Non-Bt Maize Growers". *Science* 330.6001, pp. 222–225.

ISAAA (2017). *Global Status of Commercialized Biotech/GM Crops in 2017: Biotech Crop Adoption Surges as Economic Benefits Accumulate in 22 Years*. 53. Ithaca, NY: ISAAA.

Ives, Anthony R., Cate Paull, Andrew Hulthen, Sharon Downes, David A. Andow, Ralph Haygood, Myron P. Zalucki, and Nancy A. Schellhorn (2017). "Spatio-Temporal Variation in Landscape Composition May Speed Resistance Evolution of Pests to Bt Crops". *PLOS ONE* 12.1, e0169167.

Jackson, Matthew O (2010). *Social and economic networks*. Vol. 3. Princeton university press.

Jin, Zewen, Can Chen, Xiaomin Chen, Fei Jiang, Isaac Hopkins, Xiaoling Zhang, Zhaoqiang Han, Grace Billy, and Jhony Benavides (2019a). "Soil acidity, available phosphorus content, and optimal biochar and nitrogen fertilizer application rates: A five-year field trial in upland red soil, China". *Field Crops Research* 232, pp. 77–87.

Jin, Zhenong, Elizabeth A Ainsworth, Andrew D B Leakey, and David B Lobell (2018). "Increasing drought and diminishing benefits of elevated carbon dioxide for soybean yields across the US Midwest". *Global Change Biology* 24.2, e522–e533.

Jin, Zhenong, Sotirios V Archontoulis, and David B Lobell (2019b). "How much will precision nitrogen management pay off? An evaluation based on simulating thousands of corn fields over the US Corn-Belt". *Field Crops Research* 240, pp. 12–22.

Jin, Zhenong, George Azzari, and David B Lobell (2017a). "Improving the accuracy of satellite-based high-resolution yield estimation: A test of multiple scalable approaches". *Agricultural and Forest Meteorology* 247, pp. 207–220.

Jin, Zhenong, Rishi Prasad, John Shriver, and Qianlai Zhuang (2017b). "Crop model- and satellite imagery-based recommendation tool for variable rate N fertilizer application for the US Corn system". *Precision Agriculture* 18.5, pp. 779–800.

Johanns, Ann (2019). *Iowa cash corn and soybean prices*. Cash Corn and Soybean Prices — Ag Decision Maker. URL: https://www.extension.iastate.edu/agdm/crops/html/a2-11.html.

Jones, Donald R, Matthias Schonlau, and William J Welch (1998). “Efficient global optimization of expensive black-box functions”. *Journal of Global optimization* 13.4, pp. 455–492.

Jones, James W, John M Antle, Bruno Basso, Kenneth J Boote, Richard T Conant, Ian Foster, H Charles J Godfray, Mario Herrero, Richard E Howitt, and Sander Janssen (2017). “Toward a new generation of agricultural system data, models, and knowledge products: State of agricultural systems science”. *Agricultural Systems* 155, pp. 269–288.

Karatay, Yusuf Nadi and Andreas Meyer-Aurich (2019). “Profitability and downside risk implications of site-specific nitrogen management with respect to wheat grain quality”. *Precision Agriculture.*

Kaup, Brent Z. (2008). “The reflexive producer: The influence of farmer knowledge upon the use of Bt corn”. *Rural Sociology* 73.1, p. 62.

Khan, Aziz, Leishan Wang, Saif Ali, Shahbaz Atta Tung, Abdul Hafeez, and Guozheng Yang (2017). “Optimal planting density and sowing date can improve cotton yield by maintaining reproductive organ biomass and enhancing potassium uptake”. *Field Crops Research* 214, pp. 164–174.

Leonard, Emma, Rohan Rainbow, A Laurie, David Lamb, R Llewellyn, Ed Perrett, Jay Sanderson, Andrew Skinner, T Stollery, and Leanne Wiseman (2017). *Accelerating precision agriculture to decision agriculture: Enabling digital agriculture in Australia.* Australi: Cotton Research and Development Corporation.

Leroux, Corentin and Bruno Tisseyre (2019). “How to measure and report within-field variability: a review of common indicators and their sensitivity”. *Precision Agriculture* 20.3, pp. 562–590.

Li, Maona, Yunling Wang, Ardeshir Adeli, and Haijun Yan (2018). “Effects of application methods and urea rates on ammonia volatilization, yields and fine root biomass of alfalfa”. *Field Crops Research* 218, pp. 115–125.

Li, Yan, Zhou Shi, Feng Li, and Hong-Yi Li (2007). “Delineation of site-specific management zones using fuzzy clustering analysis in a coastal saline land”. *Computers and Electronics in Agriculture* 56.2, pp. 174–186.

Lindblom, Jessica, Christina Lundström, Magnus Ljung, and Anders Jonsson (2017). “Promoting sustainable intensification in precision agriculture: review of decision support systems development and strategies”. *Precision Agriculture* 18.3, pp. 309–331.

Lo, Tsz Him, Daran R. Rudnick, Brian T. Krienke, Derek M. Heeren, Yufeng Ge, and Tim M. Shaver (2019). “Water effects on optical canopy sensing for late-season site-specific nitrogen management of maize”. *Computers and Electronics in Agriculture* 162, pp. 154–164.

Lobell, David B, Graeme L Hammer, Greg McLean, Carlos Messina, Michael J Roberts, and Wolfram Schlenker (2013). “The critical role of extreme heat for maize production in the United States”. *Nature Climate Change* 3.5, p. 497.

Lobell, David B, Michael J Roberts, Wolfram Schlenker, Noah Braun, Bertis B Little, Roderick M Rejesus, and Graeme L Hammer (2014). “Greater Sensitivity to Drought Accompanies Maize Yield Increase in the U.S. Midwest”. *Science* 344.6183, p. 516.

Lobell, David B, David Thau, Christopher Seifert, Eric Engle, and Bertis Little (2015). “A scalable satellite-based crop yield mapper”. *Remote Sensing of Environment* 164, pp. 324–333.

Lowenberg-DeBoer, Jess and Bruce Erickson (2019). “Setting the record straight on precision agriculture adoption”. *Agronomy Journal.*

Lubell, Mark and Allan Fulton (2007). “Local policy networks and agricultural watershed management”. *Journal of Public Administration Research and Theory* 18.4, pp. 673–696.

Lynch, Joseph P, Elizabeth Glynn, Steven Kildea, and John Spink (2017). “Yield and optimum fungicide dose rates for winter wheat (Triticum aestivum L.) varieties with contrasting ratings for resistance to septoria tritici blotch”. *Field Crops Research* 204, pp. 89–100.

Mason, C.E., Marlin E. Rice, D.D. Calvin, John W. Van Duyn, William B. Showers, W. D. Hutchison, J.F. Witkowski, R. A. Higgins, D. W. Onstad, and G.P. Dively (1996). *European Corn Borer Ecology and Management*. 327. Iowa State University.

McAllister, Ryan, Catherine Robinson, Kirsten Maclean, Angela Guerrero, Kerry Collins, Bruce Taylor, and Paul De Barro (2015). "From local to central: a network analysis of who manages plant pest and disease outbreaks across scales". *Ecology and Society* 20.1.

McCown, RL, PS Carberry, Z Hochman, NP Dalgliesh, and MA Foale (2009). "Re-inventing model-based decision support with Australian dryland farmers. 1. Changing intervention concepts during 17 years of action research". *Crop and Pasture Science* 60.11, pp. 1017–1030.

McCown, Robert L (2002). "Changing systems for supporting farmers' decisions: problems, paradigms, and prospects". *Agricultural Systems* 74.1, pp. 179–220.

Mead, Roger, Robert N Curnow, and Anne M Hasted (2002). *Statistical Methods in Agriculture and Experimental Biology*. 3rd. CRC Press.

Meyer-Aurich, Andreas, Terry W. Griffin, Ruprecht Herbst, Antje Giebel, and Nawaz Muhammad (2010a). "Spatial econometric analysis of a field-scale site-specific nitrogen fertilizer experiment on wheat (Triticum aestuvum L.) yield and quality". *Computers and Electronics in Agriculture* 74.1, pp. 73–79.

Meyer-Aurich, Andreas, Alfons Weersink, Markus Gandorfer, and Peter Wagner (2010b). "Optimal site-specific fertilization and harvesting strategies with respect to crop yield and quality response to nitrogen". *Agricultural Systems* 103.7, pp. 478–485.

Miller, John H. and Scott E. Page (2009). *Complex adaptive systems: An introduction to computational models of social life*. Princeton university press.

Milne, Alice E., James R. Bell, William D. Hutchison, Frank van den Bosch, Paul D. Mitchell, David Crowder, Stephen Parnell, and Andrew P. Whitmore (2015). "The effect of farmers' decisions on pest control with Bt crops: a billion dollar game of strategy". *PLOS Computational Biology* 11.12.

Mitchell, Paul D. (2017). *Methods and assumptions for estimating the impact of pyrethroid insecticides on pest management practices and costs for U.S. crop farmers*. Madison, WI: AgInfomatics.

Mitchell, Paul D. and William D. Hutchison (2008). “Decision making and economic risk in IPM”. *Integrated Pest Management: Concepts, Tactics, Strategies and Case Studies*. Ed. by Edward B. Radcliffe, Rafael E. Cancelado, and William D. Hutchison. Cambridge: Cambridge University Press, pp. 33–50.

Mitchell, Paul D. and David W Onstad (2005). “Effect of extended diapause on evolution of resistance to transgenic Bacillus thuringiensis corn by northern corn rootworm (Coleoptera: Chrysomelidae)”. *Journal of economic entomology* 98.6, pp. 2220–2234.

Mitchell, Paul D, Zachary Brown, and Neil McRoberts (2018). “Economic issues to consider for gene drives”. *Journal of Responsible Innovation* 5 (sup1), S180–S202.

Moćkus, J, V Tiesis, and A Źilinskas (1978). “The Application of Bayesian Methods for Seeking the Extremum.” Dixon, L and G Szego. *Toward Global Optimization*. Vol. 2. Amsterdam, The Netherlands: Elsevier.

Monjardino, M, T McBeath, J Ouzman, Rick Llewellyn, and B Jones (2015). “Farmer risk-aversion limits closure of yield and profit gaps: A study of nitrogen management in the southern Australian wheatbelt”. *Agricultural Systems* 137. Publisher: Elsevier, pp. 108–118.

Montgomery, Douglas C (2017). *Design and analysis of experiments*. John Wiley & Sons.

Onstad, David W. and Lance J. Meinke (2010). “Modeling Evolution of Diabrotica Virgifera Virgifera (Coleoptera: Chrysomelidae) to Transgenic Corn with Two Insecticidal Traits”. *Journal of Economic Entomology* 103.3, pp. 849–860.

Onstad, David W and Joseph V Maddox (1989). “Modeling the effects of the microsporidium, Nosema pyrausta, on the population dynamics of the insect, Ostrinia nubilalis)”. *Journal of Invertebrate Pathology* 53.3, pp. 410–421.

Orlowski, John M, Bryson J Haverkamp, Randall G Laurenz, David Marburger, Eric W Wilson, Shaun N Casteel, Shawn P Conley, Seth L Naeve, Emerson D Nafziger, and Kraig L Roozeboom (2016). "High-input management systems effect on soybean seed yield, yield components, and economic break-even probabilities". *Crop Science* 56.4, pp. 1988–2004.

Ouedraogo, Frederic and B Wade Brorsen (2018). "Hierarchical Bayesian Estimation of a Stochastic Plateau Response Function: Determining Optimal Levels of Nitrogen Fertilization". *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* 66.1, pp. 87–102.

Pahlmann, Ingo, Ulf Böttcher, and Henning Kage (2017). "Developing and testing an algorithm for site-specific N fertilization of winter oilseed rape". *Computers and Electronics in Agriculture* 136, pp. 228–237.

Pannell, Dave, Markus Gandorfer, and Alfons Weersink (2019). "How flat is flat? Measuring payoff functions and the implications for site-specific crop management". *Computers and Electronics in Agriculture* 162, pp. 459–465.

Panten, K., R. G. V. Bramley, R. M. Lark, and T. F. A. Bishop (2010). "Enhancing the value of field experimentation through whole-of-block designs". *Precision Agriculture* 11.2, pp. 198–213.

Pardo-López, Liliana, Mario Soberón, and Alejandra Bravo (2013). "Bacillus thuringiensis insecticidal three-domain Cry toxins: mode of action, insect resistance and consequences for crop protection". *FEMS Microbiology Reviews* 37.1, pp. 3–22.

Park, Eunchun, Wade Brorsen, and Xiaofei Li (2018). *How to Use Yield Monitor Data to Determine Nitrogen Recommendations: Bayesian Kriging for Location Specific Parameter Estimates*. Agricultural and Applied Economics Association.

Pathak, Hari Sharan, Philip Brown, and Talitha Best (2019). "A systematic literature review of the factors affecting the precision agriculture adoption process". *Precision Agriculture*.

Peck, Steven L. (2004). "Simulation as experiment: a philosophical reassessment for biological modeling". *Trends in Ecology & Evolution* 19.10, pp. 530–534.

Perry, Edward D, Federico Ciliberto, David A Hennessy, and GianCarlo Moschini (2016). "Genetically engineered crops and pesticide use in US maize and soybeans". *Science advances* 2.8, e1600850.

Petersen, Roger G (1994). *Agricultural field experiments: design and analysis*. CRC Press.

Piepho, Hans-Peter, Christel Richter, Joachim Spilke, Karin Hartung, Arndt Kunick, and Heinrich Thöle (2011). "Statistical aspects of on-farm experimentation". *Crop and Pasture Science* 62.9, pp. 721–735.

Pringle, M. J., Simon Cook, and A. B. McBratney (2004a). "Field-Scale Experiments for Site-Specific Crop Management. Part I: Design Considerations". *Precision Agriculture* 5.6, pp. 617–624.

Pringle, M. J., A. B. McBratney, and Simon Cook (2004b). "Field-Scale Experiments for Site-Specific Crop Management. Part II: A Geostatistical Analysis". *Precision Agriculture* 5.6, pp. 625–645.

Rasmussen, Carl Edward and Christopher K Williams (2006). *Gaussian Processes for Machine Learning*. MIT Press.

Rebaudo, François and Olivier Dangles (2013). "An agent-based modeling framework for integrated pest management dissemination programs". *Environmental Modelling & Software*. Thematic Issue on Spatial Agent-Based Models for Socio-Ecological Systems 45, pp. 141–149.

Ren, Tao, Bo Liu, Jianwei Lu, Zhonghua Deng, Xiaokun Li, and Rihuan Cong (2017). "Optimal plant density and N fertilization to achieve higher seed yield and lower N surplus for winter oilseed rape (Brassica napus L.)" *Field Crops Research* 204, pp. 199–207.

Rens, Libby R, Lincoln Zotarelli, Diane L Rowland, and Kelly T Morgan (2018). "Optimizing nitrogen fertilizer rates and time of application for potatoes under seepage irrigation". *Field Crops Research* 215, pp. 49–58.

Renton, Michael (2013). “Shifting focus from the population to the individual as a way forward in understanding, predicting and managing the complexities of evolution of resistance to pesticides”. *Pest management science* 69.2, pp. 171–175.

Renton, Michael, Roberto Busi, Paul Neve, David Thornby, and Martin Vila-Aiub (2014). “Herbicide resistance modelling: past, present and future”. *Pest Management Science* 70.9, pp. 1394–1404.

Rose, David C, Caroline Parker, Joe Fodey, Caroline Park, William J Sutherland, and Lynn V Dicks (2018). “Involving stakeholders in agricultural decision support systems: Improving user-centred design”. *International Journal of Agricultural Management* 6.3, pp. 80–89.

Ruffo, Matías L, Germán A Bollero, David S Bullock, and Donald G Bullock (2006). “Site-specific production functions for variable rate corn nitrogen fertilization”. *Precision Agriculture* 7.5, pp. 327–342.

Ryan, Bryce and Neal C. Gross (1943). “The diffusion of hybrid seed corn in two Iowa communities.” *Rural sociology* 8.1, p. 15.

Saikai, Yuji, Vivak Patel, Shawn P Conley, and Paul D Mitchell (2020). *Adaptive experimental design using Bayesian optimization to improve the cost efficiency of field trials*. URL: https://github.com/ysaikai/AEDBO.

Sawyer, John (2008). *Measuring the Nitrogen Status — Integrated Crop Management*. Iowa State University. URL: https://crops.extension.iastate.edu/cropnews/2008/06/measuring-nitrogen-status.

Sawyer, John (2019). *The Corn Nitrogen Rate Calculator*. URL: http://cnrc.agron.iastate.edu/nRate.aspx.

Schimmelpfennig, David (2016). *Farm profits and adoption of precision agriculture*. United States Department of Agriculture, Economic Research Service.

Scudiero, Elia, Pietro Teatini, Dennis L. Corwin, Rita Deiana, Antonio Berti, and Francesco Morari (2013). “Delineation of site-specific management units in a saline region at the

Venice Lagoon margin, Italy, using soil reflectance and apparent electrical conductivity". *Computers and Electronics in Agriculture* 99, pp. 54–64.

Shahriari, Bobak, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas (2016). "Taking the human out of the loop: A review of bayesian optimization". *Proceedings of the IEEE* 104.1, pp. 148–175.

Shi, Guanming, Jean-Paul Chavas, and Joseph Lauer (2013). "Commercialized transgenic traits, maize productivity and yield risk". *Nature biotechnology* 31.2, p. 111.

Shi, Guanming, Jean-paul Chavas, and Kyle Stiegert (2010). "An Analysis of the Pricing of Traits in the U.S. Corn Seed Market". *American Journal of Agricultural Economics* 92.5, pp. 1324–1338.

Showers, William B., Richard L. Hellmich, M. Ellison Derrick-Robinson, and William H. Hendrix (2001). "Aggregation and dispersal behavior of marked and released European corn borer (Lepidoptera: Crambidae) adults". *Environmental Entomology* 30.4, pp. 700–710.

Storer, K E, P M Berry, D R Kindred, and R Sylvester-Bradley (2018). "Identifying oilseed rape varieties with high yield and low nitrogen fertiliser requirement". *Field Crops Research* 225, pp. 104–116.

Storer, Nicholas P. (2003). "A spatially explicit model simulating western corn rootworm (Coleoptera: Chrysomelidae) adaptation to insect-resistant maize". *Journal of economic entomology* 96.5, pp. 1530–1547.

Swinton, Scott M., Frank Lupi, G. Philip Robertson, and Stephen K. Hamilton (2007). "Ecosystem services and agriculture: Cultivating agricultural ecosystems for diverse benefits". *Ecological Economics* 64, pp. 245–252.

Tabashnik, Bruce E. (2008). "Delaying insect resistance to transgenic crops". *Proceedings of the National Academy of Sciences* 105.49, pp. 19029–19030.

Tabashnik, Bruce E. (2015). "ABCs of Insect Resistance to Bt". *PLOS Genetics* 11.11, e1005646.

Tabashnik, Bruce E., Thierry Brévault, and Yves Carrière (2013). "Insect resistance to Bt crops: lessons from the first billion acres". *Nature Biotechnology* 31.6, pp. 510–521.

Thöle, Heinrich, Christel Richter, and Detlef Ehlert (2013). "Strategy of statistical model selection for precision farming on-farm experiments". *Precision Agriculture* 14.4, pp. 434–449.

Tilman, David, Christian Balzer, Jason Hill, and Belinda L Befort (2011). "Global food demand and the sustainable intensification of agriculture". *Proceedings of the National Academy of Sciences* 108.50, pp. 20260–20264.

U.S. Environmental Protection Agency (2018). *White Paper on Resistance in Lepidopteran Pests of Bacillus thuringiensis (Bt) Plant-Incorporated Protectants in the United States.* URL: https://www.epa.gov/sites/production/files/2018-07/documents/position_paper_07132018.pdf.

U.S. Environmental Protection Agency (2019). *Insect Resistance Management for Bt Plant-Incorporated Protectants.* URL: https://www.epa.gov/regulation-biotechnology-under-tsca-and-fifra/insect-resistance-management-bt-plant-incorporated.

USDA (2010). *Biopesticides Regrisrtaiotn Action Dcouments: Cry1Ab and Cry1F Bacillus thuringiensis (Bt) Corn Plant-Incorporated Protectants.* Washington, DC.

USDA (2012). *Census of agriculture.* URL: https://www.agcensus.usda.gov/Publications/2012/.

USDA (2016). *Commodity Costs and Returns.* URL: http://www.ers.usda.gov/data-products/commodity-costs-and-returns/commodity-costs-and-returns/#Recent%20Costs%20and%20Returns:%20Corn.

USDA (2018a). *Agricultural Prices.* URL: https://usda.library.cornell.edu/concern/publications/c821gj76b.

USDA (2018b). *World Agricultural Supply and Demand Estimates.* 581. Washington, DC.

USDA (2019). *Quick Stats Lite.* URL: https://www.nass.usda.gov/Quick_Stats/Lite/.

Useche, Pilar, Bradford L. Barham, and Jeremy D. Foltz (2009). "Integrating Technology Traits and Producer Heterogeneity: A Mixed-Multinomial Model of Genetically Modified Corn Adoption". *American Journal of Agricultural Economics* 91.2, pp. 444–461.

Wang, Hongyuan, Yitao Zhang, Anqiang Chen, Hongbin Liu, Limei Zhai, Baokun Lei, and Tianzhi Ren (2017). "An optimal regional nitrogen application threshold for wheat in the North China Plain considering yield and environmental effects". *Field Crops Research* 207, pp. 52–61.

Wang, Weini, Jianwei Lu, Tao Ren, Xiaokun Li, Wei Su, and Mingxing Lu (2012). "Evaluating regional mean optimal nitrogen rates in combination with indigenous nitrogen supply for rice production". *Field Crops Research* 137, pp. 37–48.

Whelan, B M, J A Taylor, and A B McBratney (2012). "A small strip approach to empirically determining management class yield response functions and calculating the potential financial net wastage associated with whole-field uniform-rate fertiliser application". *Field Crops Research* 139, pp. 47–56.