Mitigating floods and pestilence: Examining the provision of public goods under uncertainty

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Stephen Seyler Aultman

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Adviser: Dr. Terrance M. Hurley

January, 2012

Acknowledgements

I would like to thank my dissertation committee for their support, suggestions and guidance. I would also like to thank my friends, family and especially my wife for their support and encouragement.

Dedication

To my wife, Sara.

Abstract

This dissertation contains two essays on the provision of a specific type of public good, specifically public goods that affect the probability with which different states of the world occur. Two potential examples of this type of public good are levees and monitoring networks for disease. To distinguish this type of public good from those that have no relationship with risk we adopt the term risk-modifying public good. The first essay focuses on the provision of risk-modifying public goods within a mechanism design framework. The main result of the first essay is proof that there exists a family of social choice functions that provide a Pareto efficient level of a risk-modifying public good and that these social choice functions are implementable in dominant strategies. The family of social choice functions we identify provide a unique level of the public good and vary only in how the gains from providing the public good are distributed amongst the agents in the economy. The second essay is empirical in nature and focuses on measuring the returns to information gathered by a monitoring network for an invasive wind borne crop disease (soybean rust). What makes this information a risk-modifying public good is that it is used by agricultural producers to inform themselves about their risk of being affected by this invasive species and to make fungicide application decisions. In this type of situation, applying fungicides when the disease is actually present can be thought of as one state-of-the-world and applying fungicides when one shouldn’t another. Thus, a state-of-the-world can be thought of as a combination of agent’s beliefs about the environment relative to the reality that they face. A dynamic model of farmer decision making is presented where the information from the monitoring

network can help farmers better learn about their risk of soybean rust infection and can potentially be used as an alternative to a farmer scouting soybean fields if the information is sufficiently accurate. Using this model, the marginal returns produced by additional years of monitoring and the question of the optimal spatial arrangement of sentinel plots are considered.

Table of Contents

List of Tables vi
List of Figures vii
Introduction 1
Insurance, Mechanism Design and the Provision of Risk-Modifying Public Goods 3
  Introduction 3
  The Economic Environment 10
  Efficient Social Choice Functions Implementable in Dominant Strategies 12
  Numerical Example 20
  Discussion 24
Valuing Monitoring Networks for New Pathogens: The Case of Soybean Rust 28
  Introduction 28
  Model 30
  The Producer's Problem when no Monitoring Network Exists 36
  The Producer's Problem when a Monitoring Network Exists 37
  Results 46
  Conclusion 50
Conclusion 53
References 54

List of Tables

Table 1.1: Numerical Example Parameters 21
Table 1.2: Insurance Premiums for Agents by Social Choice Function 22
Table 2.1: Cost of Following Management Strategies (dollars/acre) 31
Table 2.2: Present Value of an Additional Year of Soybean Rust Monitoring 47

List of Figures

Figure 1.1: Utility Possibility Frontier 24
Figure 2.1: Optimal Management Decisions for a Farmer without a Monitoring Network 41
Figure 2.2: Optimal Management Decisions for a Farmer with a 4-Year Monitoring Network with Signal Quality of 0.75 42
Figure 2.3: Signal Quality as a Function of Distance 43
Figure 2.4: High Signal Quality 200 Sentinel Plot Locations 48
Figure 2.5: Sentinel Plot Locations for 2010 Growing Season 48
Figure 2.6: Marginal Value of Extending Soybean Rust Sentinel Plot System........ 50

Introduction

This dissertation contains two essays on the provision of public goods. While public goods, those goods that are both non-rival and non-excludable, have been studied extensively in the economics literature, these essays focus on a specific type of public good that has received comparatively little attention. Specifically, these essays focus on public goods that affect the probability with which different states of the world occur. Two potential examples of this type of public good are levees and monitoring networks for disease. To distinguish this type of public good from those that have no relationship with risk we adopt the term risk-modifying public good.

The first essay focuses on the provision of risk-modifying public goods within a mechanism design framework. While the literature on providing an efficient level of traditional public goods is fraught with road blocks and impossibility results, the prospects for risk-modifying public goods are much brighter. The main result of the first essay is a proof showing the existence of a family of social choice functions that provide a Pareto efficient level of a risk-modifying public good and that these social choice functions are implementable in dominant strategies. All social welfare functions within the family provide the same level of the public good. The functions vary only in how the gains from providing the public good are distributed amongst the agents in the economy.

The second essay is more empirical in nature and focuses on measuring the returns to a monitoring network for an invasive wind borne crop disease (soybean rust).

The information provided by this monitoring is a public good by its non-rival and non-excludable nature. This information helps agricultural producers to inform themselves about their risk of being affected by soybean rust about which they have no previous experience. Without this information, it is possible that farmers could end up holding incorrect beliefs about the probability that soybean rust will infect their farm. Thus, a state-of-the-world in this case is a combination of a farmer’s beliefs as learned from Bayesian updating relative to the reality of the situation. In this type of situation, applying fungicides when the disease is actually present can be thought of as one state-of-the-world and applying fungicides when the disease is not present another.

An interesting feature of the second essay is that it explicitly incorporates a statistical model where information produced by the monitoring network is incorporated into a farmer’s knowledge of their likelihood of being affected by soybean rust. After presenting a conceptual model of the farmer’s problem, a spatially explicit county level model of soybean production in the US is created using the first several years of data from the monitoring network and USDA production data. The data from the monitoring network is made publicly available by the Integrated Pest Management Pest Information Platform for Extension and Education (IPM PIPE) website. The benefits of the network are calculated by comparing the costs farmers incur (both in fungicide costs and yield loss) with and without the network present. Lastly, the problem is reformulated in terms of spatial optimization and we identify the optimal placement of a limited number of monitoring stations given the history of infections that have been observed to date.

Insurance, Mechanism Design and the Provision of Risk-Modifying Public Goods

Introduction

A town lies in quiet repose along the banks of a flood-prone river. Each spring the high water threatens and each summer the decision to build a protective levee is deferred yet again. Sometimes the floods come; sometimes they do not. Some residents purchase flood insurance to protect against potential losses; some do not. The levee under consideration is among the class of canonical public goods. If built, it would confer its benefits upon all those living behind it, reducing or eliminating the risk that this year's flood might overwhelm the town. Suppose that the town council has decided at last to make a decision. How might the council, as social planner, decide whether to build the levee and, if it is built, how high? One approach would be to take the advice of Samuelson (1954), just as canonical as the levee itself. Though his approach means ignoring the uncertainty inherent in the problem, it is a sensible guide to behavior for a well-informed planner. It says to build a levee to such a height as to equate the vertical sum of marginal rates of substitution between the levee and a private good on one hand, and the marginal cost of the levee on the other. Quite apart from the fact that it ignores uncertainty, Samuelson's method has an important, even fatal, practical shortcoming: it requires that the planner know each citizen's preferences for the levee. Asking citizens to report their preferences is generally futile. At least some people will gain by lying, a problem that did not escape Samuelson's notice[1].

So what should the council do? If, like Samuelson, it is willing to ignore uncertainty, and if it feels justified in supposing that citizens' preferences are quasi-linear, an alternative and powerful tool becomes available. With quasi-linearity, the Vickrey-Clarke-Groves (VCG) mechanism can be brought to bear in selecting a levee height and deciding whether to build the levee (Krishna and Perry 2000). The VCG mechanism provides attractive incentives for respondents to the planner's poll. It is a dominant-strategy mechanism: all will report their true net valuation of the levee. The trick that achieves this result is that a specialized tax scheme forces all pivotal respondents to pay for the externality that their presence in the mechanism imposes on others. Krishna and Perry show that the VCG mechanism, extended to incorporate a “basis" that accounts for each respondent's worst outcome given the set of people, maximizes the revenue collected among the class of efficient mechanisms. The VCG mechanism, though appealing, has its own drawbacks. Chief among them is that it does not guarantee that the budget constraint will be satisfied. In general, the VCG mechanism raises more revenue than is needed to fund the public project.

An essential feature of the flooding problem, and by extension of the levee problem, is that it is fundamentally uncertain. It is uncertain, in the sense that an exogenous random variable influences agents' outcomes and a decision is to be made *ex ante*, before the realization is known. Accounting for this risk would appear to be

---

[1] Samuelson writes, ``*However no decentralized pricing system can serve to determine optimally these levels of collective consumption*. Other kinds of ``voting" or ``signaling" would have to be tried. But, and

fundamentally important. Our enterprise is to examine carefully public-goods problems in this class, which can be seen to include public-health threats, community and national security, as well as many environmental hazards.

We note that uncertainty of the sort considered here (whether or not a river will flood, or whether or not an exotic flu will reach a given city) is in the sense of Arrow (1964) and is different from the uncertainty encountered elsewhere in the mechanism-design literature. It is different from uncertainty, on the part of one agent, regarding the preferences of others, where Bayesian implementation comes into play (Jackson 1991). It is different from uncertainty regarding whether the public good will be provided at all (Neeman 2004). And it is different from stochastic social-choice rules (Pattanaik and Peleg 1986) where an outcome is selected from a finite set of alternatives at random. In Pattanaik and Peleg and related literature randomness arises from the social-choice function itself. To distinguish the type of public good we are referring to from those without a direct relationship to uncertainty, we adopt the term risk-modifying public good.

Even apart from concerns over mechanism design, the presence of state-contingent uncertainty adds a significant complication to the study of public goods. Indeed, identifying Pareto-efficient outcomes in this case is itself far from straightforward. In a remarkable paper that shares many features with ours, Graham (1992) devised a general-purpose criterion for identifying Pareto-efficient public projects under uncertainty. Graham's net-benefit criterion is capable of determining the set of efficient vectors of public goods and, at the same time, the set of efficient accompanying state-contingent payments between agents across states.

The combination of, and distinction between, the public good and the state-contingent payments that support it is the essential feature of Graham. Upon it, we build our efficient, dominant-strategy, and budget-balancing mechanism. In Graham the contingent payments serve two distinct but interrelated functions. They fund the vector of public goods and they also allow for the optimal sharing of risk between agents with heterogeneous risk preferences. Graham shows how the problem of public-goods provision presents a valuable opportunity: the public good, together with the funding scheme, provides optimal risk sharing that would otherwise be impossible given incomplete markets for specific risks. Graham's criterion has the added virtue of distinguishing the welfare gains due to the public goods from those due to the risk sharing achieved by the contingent payments.

Graham's framework is a powerful tool for identifying efficient public projects, but in practice it has three important shortcomings. One is that, though it takes us to the efficient frontier, it abandons us there. The second is that Graham's criterion is not helpful in achieving the desirable outcomes he describes. Eschewing implementation, he leaves unanswered the question of how any efficient outcome might be achieved in a decentralized fashion.[2]

[2] Chavas and Coggins (2003) refine the Graham criterion so as to identify a subset of the Pareto-efficient outcomes having attractive equity properties. They adapt the notion of egalitarian equivalence (Pazner and Schmeidler 1978) to the case of uncertainty in order to rule out most of the alternatives that pass Graham's net-benefit criterion. Chavas and Coggins call the remaining efficient alternatives "fair equivalent," and argue that decisions regarding public goods should be limited to this subset of the efficient alternatives.

The third shortcoming is that in some cases the computational burden of using Graham's criterion in applied problems quickly becomes insurmountable. Graham assumes that all agents experience the same state of the world. If this is not true, if risk is idiosyncratic, then the number of possible states grows exponentially large. Because his scheme requires computing a hypothetical price for each possible state, this leads to difficulties. If, for example, s independent random outcomes are possible for each of n agents, then the number of possible parameters that need to be calculated is given by ns^n. In the case of the flu, if 32 residents in a village each face an independent binary outcome of infection or no infection, the number of states exceeds 137 billion.[3]

We contribute to the literature by describing a mechanism implementable in dominant strategies that achieves Pareto efficient outcomes for a particular class of public-goods problems. It appears that our treatment is the first to address implementation of a public-good decision under uncertainty. The class to which our analysis applies is characterized by the following properties. There are two states of the world that an agent can realize. In the good state the agent experiences no harm; in the bad state each suffers a monetary loss that is public information. Importantly, we do not require that all agents experience the same state realization, that is to say the model is flexible enough to allow idiosyncratic risk. Damages can differ across people, but are assumed to fall short of an endowed level of wealth. The public good in question has the

[3] The entirety of the world's computing capacity would be unable to solve this problem for the city of St. Paul, Minnesota and its 115,435 households.

effect of reducing the probability that agents experience the bad state of the world according to a well-behaved cost function.[4]

At an efficient outcome, which our mechanism implements, the planner collects state-dependent payments from each agent. A portion of the aggregate payments is used to finance the public good. The remainder is used to buy reinsurance from an exogenous competitive financial market, with the returns used, in the bad state, to indemnify policyholders. By offering this package, the planner can induce agents to make individual decisions that, when aggregated, lead to adoption of the efficient public outcome. Indeed, the choice to accept the offered package is a dominant strategy for each agent.[5] An optimal insurance package in this setup, it turns out, renders agents' risk preferences irrelevant. All decisions, the level of the public good and the state-contingent payment scheme, are ex ante and thus require that budgets balance on average. When offered fair insurance all will fully insure and so become indifferent to which outcome is realized. Thus, the planner can select the optimal level of the public good, and also the optimal insurance policy, without holding any information about agents' preferences. Unlike with the VCG mechanism, our scheme also balances the budget. And, as in Graham, our scheme selects the optimal level of the public good while at the same time selecting an optimal sharing of risk.

There are two assumptions that are essential to our scheme. First is the existence of an exogenous competitive financial market that stands ready to be the counter-party to

---

[4] In the parlance of Ehrlich and Becker (1972), this sort of public good provides "self-protection,'' albeit at the public level.

[5] We show that if all agents are risk averse, approving the package is a strictly dominant strategy. For a risk neutral agent, approval is weakly dominant.

any actuarially fair bet. We argue that this departure from Graham's setup is justifiable both from a technical and a practical standpoint. The Graham model is self-contained in the sense that funding for his public goods comes exclusively from the state-contingent payments of the agents. That is to say the state with the lowest economy-wide realization of wealth is a binding constraint because the public good must be funded by the agents in every possible state of the world.

Our assumption of an exogenous financial market is also practical. Reinsurance is a major industry, and so it seems natural to mimic the real world in this regard. An insurer whose only business is to provide flood insurance to residents of Fargo, North Dakota, would face a severe risk of default in years when flooding is severe. Insurers do, and must, diversify their portfolios across many policy holders in many different areas so that their indemnity obligations can be met in Fargo out of premium payments from communities that escaped catastrophe that year. Requiring all of the funding for a local public good to come from local state-contingent payments is unattractive for both technical reasons (default risk) and practical reasons (it does not happen in practice).

The second assumption is that the planner knows individual damages in the bad state. Local governments, raising a significant portion of revenue through property taxes, have a keen interest in knowing how much a home is worth. Private insurers have an equally keen interest in understanding their financial exposure should a home be damaged by flood or fire. Insurance companies spend significant time and money assessing damages that must sometimes withstand court scrutiny.

The class of problems to which our scheme applies overlaps with, but is different from, that addressed by Graham. The levee described above is one example. Another is a public-health threat such as influenza. Another is the infestation of agricultural or forest lands by invasive pests. Another is an environmental threat such as the London fog or a local air inversion that can have significant health effects. In each of these cases the bad state is experienced widely, the harm to individuals can be closely approximated without relying upon signals from agents (think of the city assessor's data on property values), and a public good can be designed to alter the probability of the bad state.

Our approach and our results have important policy implications. Chief among them, perhaps, is the suggestion that public insurance should be provided in concert with public goods. The US Federal Emergency Management Agency, for example, provides flood insurance. Much of the flood-protection work in the country is undertaken by the U.S. Army Corps of Engineers. Our work suggests these agencies should work together more than they currently do, or should perhaps even be merged for the purposes of flood protection and flood insurance.

The Economic Environment

Consider an uncertain economy with N agents indexed by i with $i = 1,\ldots, N$, and two possible states of the world, indexed by j with $j = 1$ (the bad state) and $j = 2$ (the good state) such that $j \in J = \{1,2\}$. The probability of the bad state occurring depends on the level of the public good $Z \in \mathbb{R}_+$. The cost of providing Z is given by the C^2 cost function $c : \mathbb{R}_+ \to \mathbb{R}_+$ with the following properties: $c(0) = 0$, $c'(Z) > 0$, and $c''(Z) > 0$. The probability with which the bad state of the world occurs is given by a C^2 function $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ with the following properties: $p(0) = \bar{p}$, $p'(Z) < 0$, and $p''(Z) > 0$.

The *N* agents are characterized by types. The *i*th agent's type depends on initial wealth in the good state W_i, losses to this wealth in the bad state D_i, and a strictly concave and twice differentiable indirect utility function U_i: $\mathbb{R} \rightarrow \mathbb{R}$. Let $\theta_i = (W_i, D_i, U_i)$ be the *i*th agents type such that

$$\theta_i \in \Theta = \{(W, D, U): W \in \mathbb{R}_{++}, D \in \mathbb{R}_{++}, \text{and } U \in \text{concave } C^2\}.$$

Let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_N)$ be the set of agent types. *Nature* randomly chooses a type from the probability distribution function $F(\theta)$. Nature reveals θ_i to agent *i* for all *i*. Furthermore, D_i for all *i* is public information and revealed to all agents. All agents also know $F(\theta)$, *p*(*Z*), and *c*(Z).

An individual's preference relation is defined over allocations

$$t \in T = \{(t^1, t^2, Z): t^1 \in \mathbb{R}^N, t^2 \in \mathbb{R}^N, Z \in \mathbb{R}_+\}$$

and characterized by the expected utility function

$$u_i(t, \theta_i) = p(Z)U_i(W_i - t_i^1 - D_i) + \big(1 - p(Z)\big)U_i(W_i - t_i^2)$$

such that for *t*, *t'* $\in T$, *t* is strictly preferred to *t'* for individual *i* if and only if $u_i(t, \theta_i) > u_i(t', \theta_i)$.

In this framework t_i^1 and t_i^2 represent state-contingent transfers between agent *i* and the public agency or planner (equivalently, a government licensed firm) that has the exclusive authority to make investments in the public good *Z*. The planner is restricted to raising funds by offering insurance products and has no authority to levy compulsory taxes. The planner is required, however, to earn zero profits. A critical assumption of this work is that the planner has access to a competitive financial market that stands ready to serve as counterparty to any actuarially fair bet. Thus, the risk facing agents in the model is fully diversifiable. This assumption has important implications for the budget constraint faced by the planner. Expressed mathematically, the planner's feasible set of allocations is:

$$T^f = \left\{ t \in T : p(Z) \sum_i t_i^1 + \left(1 - p(Z)\right) \sum_i t_i^2 = c(Z) \right\}$$

An important property of this budget constraint is that *i* holds only in expectation, which is why $p(Z)$ is present, and not in each state of the world. This can be thought of as the planner striking a deal with a reinsurer to diversify their risk. To summarize, our *economic environment* is given by $\varepsilon = (\boldsymbol{\theta}, p(Z), c(Z), T^f)$.

Having addressed these preliminaries, we now turn to our contribution to the literature. Specifically, we will show that there exists a family of social choice functions for this kind of economic environment which are both Pareto efficient and implementable in dominant strategies. We will present the family of social choice functions first, followed by proofs of its efficiency and implementability.

Efficient Social Choice Functions Implementable in Dominant Strategies

Here we present a family of social choice functions that are Pareto efficient and implementable in dominant strategies. The class or family of social choice functions we present always yield a unique level of the public good and quantity of insurance for each agent. What varies across the social choice functions is how the gains from the public good are to be divided amongst the agents. The class of social choice functions is described by the following equations:

$$p'(Z)\sum_i D_i + c'(Z) = 0 \tag{1}$$

$$t_i^1 = a_i - D_i \;\forall\, i \tag{2}$$

$$t_i^2 = a_i \;\forall\, i \tag{3}$$

$$a_i < p(0)D_i \;\forall\, i \tag{4}$$

$$\sum_i a_i = p(Z)\sum_i D_i + c(Z) \tag{5}$$

Proposition 1. The allocation $t^* = \left(t^{1^*}, t^{2^*}, Z^*\right)$ described by equations (1) - (5) is Pareto-efficient.

Proof. To show that this family of social choice functions is Pareto efficient we will setup the social planner's problem for the economy and show that the family of social choice functions satisfies all of the first-order necessary conditions. Because the functions *c*(*Z*) and *p*(*Z*) are well behaved the first-order necessary conditions are also sufficient. The social planner's problem is given by equation (6) with weights on each agent's utility given by $\mu_i > 0$:

$$\max_{t_i^1, t_i^2, Z} \sum_i \mu_i u_i(t, \theta_i) \tag{6}$$

$$s.t. \quad p(Z)\sum_i t_i^1 + \left(1 - p(Z)\right)\sum_i t_i^2 - c(Z) = 0.$$

Setting up the maximization problem in Lagrangian form we get

$$L = \sum_i \mu_i u_i(t, \theta_i) + \lambda(p(Z)\sum_i t_i^1 + \left(1 - p(Z)\right)\sum_i t_i^2 - c(Z)). \tag{7}$$

Taking the first-order conditions yields

$$\frac{\partial L}{\partial t_i^1}: \quad -\mu_i * p(Z)U_i{}'(W_i - t_i^1 - D_i) + \lambda p(Z) \quad = 0 \; \forall \, i, \tag{8}$$

$$\frac{\partial L}{\partial t_i^2}: \quad -\mu_i * (1 - p(Z))U_i{}'(W_i - t_i^2) + \lambda(1 - p(Z)) \quad = 0 \; \forall \, i, \tag{9}$$

$$\frac{\partial L}{\partial Z}: \; - p'(Z)U_i(W_i - t_i^2) + p'(Z)U_i(W_i - t_i^1 - D_i)$$

$$-\lambda \left(p'(Z) \textstyle\sum_i t_i^1 - p'^{(Z)} \sum_i t_i^2 - c'(Z)\right) = 0, \text{ and} \tag{10}$$

$$\frac{\partial L}{\partial \lambda}: \; p(Z) \textstyle\sum_i t_i^1 + \left(1 - p(Z)\right) \sum_i t_i^2 - c(Z) = 0. \tag{11}$$

We begin by focusing on the first-order condition for *Z*. Substituting equations (2) and (3) into (10) we get

$$-p'(Z)U_i(W_i - a_i) + p'(Z)U_i(W_i + D_i - a_i - D_i)$$

$$-\lambda \left(p'(Z) \textstyle\sum_i (-D_i + a_i) - p'(Z) \sum_i a_i - c'(Z)\right) = 0. \tag{12}$$

Simplifying (12) we get

$$-\lambda \left(p'(Z) \textstyle\sum_i (-D_i + a_i) - p'(Z) \sum_i a_i - c'(Z)\right) = 0. \tag{13}$$

Simplifying (13) we arrive at $\lambda \left(p'(Z) \sum_i D_i + c'(Z)\right) = 0$, which is satisfied when (1) is satisfied.

Next we turn our attention to the first-order condition with respect to λ. Substituting in (2) and (3) into (11) we get

$$p(Z) \textstyle\sum_i a_i - D_i + \left(1 - p(Z)\right) \sum_i a_i - c(Z) = 0. \tag{14}$$

Simplifying (14) we get

$$\textstyle\sum_i a_i = p(Z) \sum_i D_i + c(Z). \tag{15}$$

Equation (15) is identical to (5) for our family of social choice functions. This equation states that the revenue collected must be sufficient to pay for the expected damages and the cost of levee, but how that burden is distributed over the agents (note that it is only the sum of a_i that is important) is flexible.

Lastly, we turn our attention to (8) and (9). This part of the proof is slightly different from the others because we are working in the opposite direction than is typically used for a social planner's problem. In a traditional social planner's problem the goal is to find an allocation that maximizes the weighted sum of utility. Here we are doing the reverse, given an allocation we are proving that there exists a set of weights on each agent's utility such that the allocation satisfies the social planner's problem. Substituting (2) and (3) into (8) and (9) we get

$$-\mu_i p(Z) U_i{}'(W_i + D_i - a_i - D_i) + \lambda(p(Z) \ = 0 \ \forall \, i \ \text{ and} \tag{16}$$

$$-\mu_i (1 - p(Z)) U_i{}'(W_i - a_i) + \lambda((1 - p(Z)) \ = 0 \ \forall \, i. \tag{17}$$

Simplifying either (16) or (17) yields the same result, namely

$$\mu_i \ = \frac{\lambda}{U_i{}'(W_i - a_i)} \ \forall \, i. \tag{18}$$

Thus, for any allocation satisfying (1) - (5) we can show that it is Pareto efficient because it satisfies first-order necessary conditions of the social planner's problem. However, while we can know that the family of social choice functions is efficient without knowledge of agent's risk aversion or wealth, it is not possible to know the values of the weights. Put another way, while we can be sure that an allocation satisfying (1) - (5) is on the Pareto-efficient frontier, without knowledge of each agent's utility function, it is not possible to know *where* on the frontier a particular allocation lies. This concludes the proof for Proposition 1.

Next, we turn our attention to the main result of the paper: the development and analysis of a mechanism that implements the Pareto-efficient allocation t^* in dominant strategies. By showing one mechanism that implements the family of social choice functions described in (1) - (5), we prove that the family is implementable in dominant strategies. The mechanism-design environment consists of the following elements. Let $\boldsymbol{D} = \{D_1, \dots D_N\}$ and let us define a profile as $(\boldsymbol{D}, c(\cdot), p(\cdot))$. An outcome is defined as $H = (t^1, t^2, Z) \in \mathbb{R}^{2N+1}$. The *social choice function* f maps a profile into an outcome. Note that our profile only contains public information and it is this fact that makes it possible to implement a Pareto efficient outcome in dominant strategies while balancing the budget.

Given this economic environment, a mechanism denoted by g, endows agents with a strategy set M_i, with $\boldsymbol{M} = M_1 \times \dots \times M_N$, and maps a vector of strategies $\boldsymbol{m} \in \boldsymbol{M}$ into an outcome. Let $M_i = \{0,1\}$, where $m_i = 0$ means that agent *i* does not purchase the insurance package that is offered and $m_i = 1$ means that the agent *i* does purchase the offered insurance package. The mechanism may be thought of as a two-stage game, though the first "stage" is an announcement by the planner (who is not a player) of a comprehensive schedule of insurance policies, state-contingent payment vectors, and a level of public goods provision, *Z*. The schedule shows the amount of insurance coverage agent *i* will receive if she accepts the package. It also gives the premium she will pay and the level of public good that will be provided.

Expressed mathematically, the mechanism we propose works as follows:

$$g_i(\boldsymbol{m}) = (t_i^1, t_i^2, Z) = \begin{cases} (0,0,0) & \text{for} \quad m_i = 0 \\ ({a_i}^* - D_i, {a_i}^*, Z^*) & \text{for} \ \sum_i m_i = N \\ (p(0)D_i - D_i, p(0)\,D_i\,, 0) & \text{otherwise} \end{cases}$$

where $\{{a_i}^*, Z^*\}$ are values that satisfy the family of social choice functions described in (1) - (5).

Explaining the mechanism in words, if an agent submits a message of 0, she pays no insurance premium, receive no insurance indemnity and no public good is provided. In this way we include an implicit participation constraint in our mechanism because this is the *status quo* at the start of the model. If all agents submit a message of 1, then they each receive full insurance and pay a premium ${a_i}^*$ and a quantity of public good Z^* is provided where ${a_i}^*$ and Z^* are a solution from the family of Pareto efficient social choice functions described by (1) - (5) which are calculable with only publically available information. If an agent submits a message of 1 but not all agents do so, then the agents submitting a message of 1 receive full insurance and pay an actuarially fair premium, but the public good is not provided.

Definition 1. We say that the mechanism $g(\cdot)$ **implements** the social choice function $f(\cdot)$ in dominant strategies if for all profiles $(\boldsymbol{D}, c(\cdot), p(\cdot)\,)$, each agent *i* accepts the insurance package $(m_i = 1)$ regardless of the strategy vector $\boldsymbol{m}_{-i}$ and the resulting dominant-strategy equilibrium outcome $g(\boldsymbol{D}, c(\cdot), p(\cdot)\,)$ coincides with $f(\boldsymbol{D}, c(\cdot), p(\cdot)\,)$.

Proposition 2. Given an economic environment ε as specified above, the mechanism *g* implements the Pareto-efficient outcome in dominant strategies. Moreover, the planner's budget is balanced.

Our proof of Proposition 2 requires establishing two separate results. First, we show that agents will always prefer to receive full insurance and no public good relative to no insurance and no public good. That is they will prefer to accept the insurance package if at least one agent has not accepted the package. Second, we must show that each agent prefers to accept the insurance package if all others have done so. In this case i's decision is pivotal in the sense that it causes the public good to be provided in a positive quantity. Here the strategy of proof is to show that the utility of accepting the insurance package when all the other agents have is greater than accepting the insurance package when at least one agent hasn't. Then by transitivity, each will maximize her utility by always accepting the insurance package regardless of the actions of the other agents, hence a dominant strategy.

Lemma 1. Given an economic environment ε as specified above, a risk-averse agent will fully insure against a potential loss if the insurance premium is actuarially fair. Equivalently, $u_i\big((0,0,0), \theta_i\big) < u_i\big((p(0)\ D_i - D_i, p(0)\ D_i, 0), \theta_i\big)\ \forall\ \theta_i \in \Theta$ and $\forall\ i$.

Proof. Consider an agent *i* from the economic environment described above. Suppose that the agent is offered insurance against loss D_i, with premium rate *a*. The agent chooses a level of insurance coverage I_i to solve

$$\max_{\mathrm{I_i} \in [0, D_i]} \ \big(1 - p(0)\big)U_i(W_i - aI_i) + p(0)U_i(W_i - aI_i - D_i + I_i)$$

where a represents the premium they pay per unit of insurance coverage. Differentiate with respect to I_i to obtain the first-order necessary condition

$$-a\big(1-p(0)\big)U_i^{'}(W_i - aI_i) + p(0)(1-a)U_i^{'}(W_i - D_i + (1-a)I_i). \tag{19}$$

If the insurance premium is actuarially fair, so that $a = p(0)$, then (19) becomes

$$U_i'(W_i - p(0)I_i) = U_i'(W_i - D_i + (1-p(0))I_i). \tag{20}$$

Because $U_i'' < 0$, equation (20) can be true only if

$$W_i - p(0)I_i = W_i - D_i + (1-p(0))I_i,$$

which in turn implies that at the optimal indemnity level we must have $I_i^* = D_i$ as was to be shown. Furthermore, when fully insured, a risk-averse agent is indifferent over the value of p and so is strictly better off than without insurance:

$$U_i'(W_i - p(0)I_i^*) > \big(1-p(0)\big)U_i(W_i) + p(0)U_i(W_i - I_i^*).$$

This completes the proof of Lemma 1.

Lemma 2. Given an economic environment ε as specified above, all agents will prefer the outcome where all agents accept the insurance package to the outcome where they receive full insurance at actuarially fair prices but with no public good. Mathematically, $u_i\big((a_i^* - D_i, a_i^*, Z^*), \theta_i\big) > u_i\big((p(0)\,D_i - D_i, p(0)\,D_i, 0), \theta_i\big)\ \forall\ \theta_i \in \Theta$ and $\forall\ i$.

Proof. Expanding the mathematical statement of Lemma 2 we get

$$\big(1-p(0)\big)U_i\big(W_i - D_i p(0)\big) + p(0)U_i(W_i - D_i p(0) - D_i + D_i) <$$
$$\big(1-p(Z^*)\big)U_i(W_i - a_i^*) + p(Z^*)U_i(W_i - a_i^* - D_i + D_i)\ \forall\ i. \tag{21}$$

Simplifying (21) we get

$$\big(1-p(0)\big)U_i\big(W_i - D_i p(0)\big) + p(0)U_i\big(W_i - D_i p(0)\big) <$$

$$\left(1 - p(Z^*)\right)U_i(W_i - a_i{}^*) + p(Z^*)U_i(W_i - a_i{}^*)\ \forall\ i. \tag{22}$$

Simplifying (22) we get

$$U_i\left(W_i - D_i p(0)\right) < U_i(W_i - a_i{}^*)\ \forall\ i. \tag{23}$$

Because $U_i{}' > 0$, (23) can be true only if $D_i p(0) > a_i{}^*\ \forall\ i$. This is precisely condition (4) that describes the family of social choice functions. This concludes the proof of Lemma 2.

Using Lemmas 1 and Lemma 2, we can now prove Proposition 2. The mechanism $g(\cdot)$ is defined such that when all agents accept the insurance package, submit a message of 1, that the mechanism yields the same outcome as the social choice function. Thus, to prove that $g(\cdot)$ implements $f(\cdot)$ it is necessary to show that it is a dominant strategy for all agents to submit a message of 1. This follows from Lemma 1, Lemma 2 and the transitivity of real numbers. By Lemma 1, we have shown $u_i\left((0,0,0), \theta_i\right) < u_i\left((p(0)\ D_i - D_i, p(0)\ D_i, 0), \theta_i\right)\ \forall\ \theta_i \in \Theta$ and $\forall\ i$. By Lemma 2 we have shown $u_i\left((a_i{}^* - D_i, a_i{}^*, Z^*), \theta_i\right) > u_i\left((p(0)\ D_i - D_i, p(0)\ D_i, 0), \theta_i\right)\ \forall\ \theta_i \in \Theta$ and $\forall\ i$. Thus, by transitivity, $u_i\left((0,0,0), \theta_i\right) < u_i\left((p(0)\ D_i - D_i, p(0)\ D_i, 0), \theta_i\right)\ \forall\ \theta_i \in \Theta$ and $\forall\ i$, which is the last item that needs to be shown for submitting a message of 1 to be a dominant strategy for each agent. This concludes the proof of Proposition 2.

Numerical Example

In this section, a numerical example involving two possible social choice rules that satisfy (1) - (5) will be presented and their effect on inequality considered. We consider an economy with two agents. The description of the economy and the agents is presented in Table 1. For this example, we assume that agents have constant relative risk aversion utility functions and that there is a significant disparity in wealth between the two agents.

Table 1.1: Numerical Example Parameters

	Agent 1	Agent 2
W_i	10	20
D_i	5	10
U(x)	$\frac{x^{0.5}}{0.5}$	
$c(Z)$	$0.5Z^2$	
$p(Z)$	$0.1 + \frac{0.1}{1+Z}$	

Using the condition specified in (1) to choose the optimal level of public good we choose $Z^* = 0.592$. This level of the public good will be the same for the entire family of social choice functions that satisfy (1) - (5). What can vary across potential social choice functions is how the gains created by the public good are distributed amongst the agents in the economy. For this example, the gains provided by the public good can be calculated as follows:

$$gains = \big(P(0) - P(0.592)\big)15 - C(0.592) = 0.383. \tag{24}$$

To uniquely specify a social choice function of the kind that satisfies (1) - (5), it is necessary to include a function that gives the value of a_i for each agent in the economy. We present two possible options given in (25) and (26). In the first option, where N signifies the number of agents in the economy, each agent in the economy shares in the gains from the public good equally. In the second option, the agents in the economy split the gains in proportion to the amount of damage that they will suffer in the bad state.

$$Option\ 1\text{:}\quad a_i = p(0)D_i - \frac{gains}{N} \quad \forall\, i \tag{25}$$

$$Option\ 2\text{:}\quad a_i = p(0)D_i - gains\frac{D_i}{\sum_i D_i} \quad \forall\, i \tag{26}$$

The insurance premium, a_i, that each agent would pay under the two rules are presented in Table 2. In option 1, the gains from the public good are distributed uniformly back to all of the agents in the economy. In option 2, the gains from the public good are distributed in proportion to the amount that agents suffer in the bad state of the world. As these agents will also be the ones paying the highest premium, this type of rule is very similar to the way that consumer cooperatives distribute profits.

Table 1.2: Insurance Premiums for Agents by Social Choice Function

	Agent	
	1	2
Option 1	0.809	1.809
Option 2	0.872	1.745

In Figure 1, we present the utility possibility frontier for the agents in the economy denoted by a dashed line. In addition to the utility possibility frontier, the utilities achieved at the initial endowment, option 1 and option 2 are designated with points in the figure. From this figure we can see that option 1 benefits agent 1 (the poorer of the two agents) more than option 2. In general, however, the effect of this type of mechanism will depend critically on whether the damages agents suffer in the bad state of the world are increasing or decreasing in wealth. Finally, taking into account the implementability constraint embodied in (4), it is possible to plot the set of Pareto efficient outcomes that can be achieved by the mechanism. This is shown in Figure 1 as the thick section of the utility possibility frontier. For this particular example, the mechanism is capable of implementing a significant fraction of the possible Pareto improving outcomes. The outcomes that are not achievable stem from the fact that the planner does not know the curvature of each agent's utility function and thus must ensure (4) holds to be certain that each agent will agree to the proposed plan. If all of the agents in the economy were risk neutral the mechanism we have presented could achieve any of the Pareto improving outcomes. Agreeing to the proposed plan would be only a weakly dominant strategy for each agent.

Figure 1.1: Utility Possibility Frontier

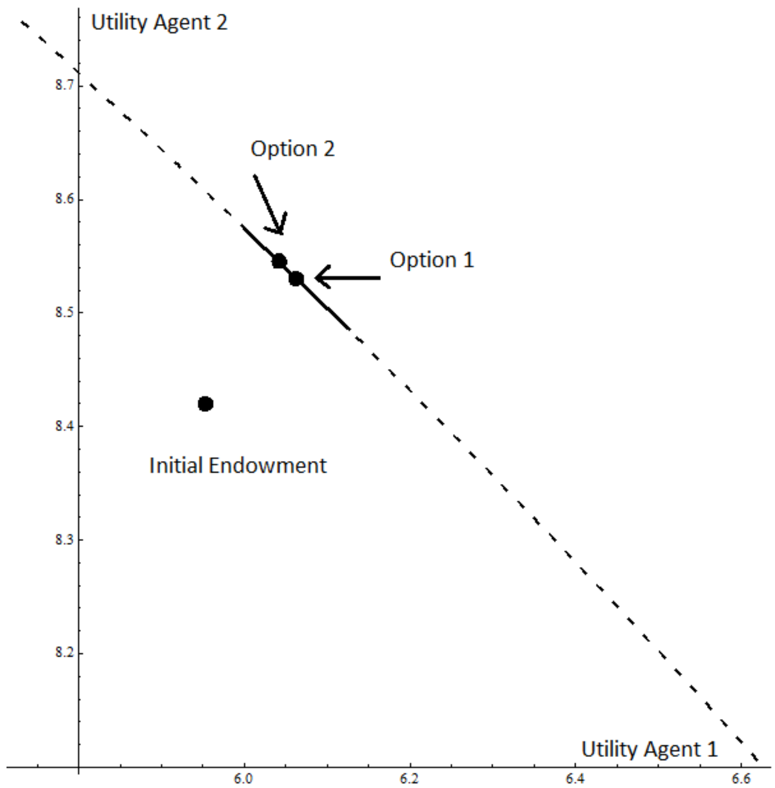

Discussion

The key insight in this work is that, when there is a competitive financial market present, it is possible to offer all agents full insurance at actuarially fair prices. When agents are fully insured (they have the same realization of wealth in all states, they are indifferent to which state of the world actually occurs and are only concerned with the cost of insurance. This makes it possible to implement a Pareto efficient outcome without knowing anything about the utility functions of the agents, other than that they are increasing in wealth. However, knowing nothing about preferences restricts the set of implementable outcomes. For instance, the reason that our mechanism charges everyone at most the actuarially fair rate for insurance when no public good is provided is because if one were to charge more, risk neutral individuals wouldn't find it in their interests to accept the plan and thus the mechanism is no longer dominant strategy implementable.

A potential complaint of this work is that by having the potential for full insurance, the optimal level of the public good doesn't depend on a decentralized pool of knowledge that needs to be aggregated. If a policy maker knows the damage that each agent will suffer in the event of the bad state occurring, the cost of providing the public good and how the public good affects the probability with which different states of the world occur, they can calculate the efficient level of public goods provision. While this is true, it neglects the very important question of how to pay for the public good.

In the United States, the Army Corps of Engineers is responsible for selecting the appropriate size of flood control projects by conducting a cost-benefit analysis and then selecting the project which most helps “national economic development.” To pay for the flood control projects the government levies a variety of taxes such as on labor, consumption and capital. These taxes distort prices in the economy and create dead weight loss. So, while it is possible to determine and construct the efficient level of the public good separate from flood insurance, such a scheme isn't Pareto efficient. If the Army Corps of Engineers were to be combined with the Federal Emergency Management Agency (FEMA), which is responsible for the national flood insurance program, then the joint agency could fund the ideal sized project out of insurance revenue as in the mechanism we propose and avoid any dead weight loss associated with traditional taxes.

There are also several interesting institutional considerations that arise from the possible combination of the Army Corps of Engineers, or at a minimum its flood control activities, with FEMA. First, the requirement that projects be self-funding through insurance revenue places a check on projects that don't make economic sense. Currently, the Federal government pays for approximately 70 percent of the cost of projects. While the Corps of Engineers is supposed to select the project that maximizes "national economic development," this requirement can be waived by the head of the Army Corps of Engineers. Second, the assistance provided by FEMA can be viewed as an implicit insurance policy for natural disasters. It is implicit because it is funded out of general tax revenues instead of entering into explicit agreements with residents. While there are situations where such an arrangement can be welfare improving for certain kinds of risks, this implicit style of insurance is especially vulnerable to the problem of adverse selection. For instance, individuals building on property in a flood plain have the expectation that they will be bailed out by FEMA in the event of a disaster. By combining FEMA with the Army Corps of engineers, FEMA might move from implicit to explicit insurance with respect to flooding and thus have an incentive to price risk appropriately and send correct price signals to investors about the returns to building in risky areas.

Lastly, this work suggests a way forward to reform the Corps' unusual and complex cost-benefit rules. In this paper, we show that the optimal levee height when agents are fully insured is the levee that minimizes the sum of expected indemnity payments plus the cost of the levee. The Corps' cost benefit rules, however take a different approach and instead view the situation through their concept of "national economic development." For instance, a grocery store owner will have lost income in the event of a flood that they would like to insure against. While in our mechanism these losses would be considered, the Corps disregards these type of costs on the basis that residents will buy groceries somewhere else and thus even though it has flooded, the level of economic activity at the national level with respect to groceries has remained the same. While it is easy to see why an agency funded with money from the Federal government would want to take a national perspective on the consequences of a disaster, this is <u>not</u> sound economic thinking. Especially when viewed as a large aggregate, the quantity of economic activity can be a very poor measure of welfare. In the event of a major flood, displaced residents might spend *more* on groceries and food than they would have without the flood, however, this should not be confused as an increase in the national welfare. By definition, if an agent is willing to pay a reasonable insurance premium to cover a potential loss, then that loss is worth considering when conducting an accounting to determine the efficient level of the public good. By utilizing a mechanism similar to the one we have described in our paper, the Corp's cost benefit rules can be simplified while still protecting against the possibility of over-sized flood control projects.

Valuing Monitoring Networks for New Pathogens: The Case of Soybean Rust

Introduction

In 2004, soybean rust (*Phakopsora pachyrhizi*) arrived in the United States, likely carried from South America by the winds of Hurricane Ivan (Isard *et al.*, 2005). Soybean rust is a fungal plant pathogen that cannot overwinter in temperate climates found in most of the United States, but thrives throughout the year in the heat and humidity of the Gulf Coast. Rust spores originating in the Gulf Coast are capable of traveling hundreds of miles on wind currents during the soybean growing season. When weather conditions are suitable, rust spores deposited on soybean fields can take hold and cause substantial yield loss. When soybean rust first arrived in the United States, there were concerns that the economic costs of its presence could be large. A study by the United States Department of Agriculture's Economic Research Service predicted that the losses in the first year of infestation could be between $640 million and $1.3 billion (Livingston *et al.* 2004). In 1984, a different study estimated that the economic consequences of soybean rust, once fully established, could be as high $7.1 billion per year (Kuchler *et al.* 1984). In response to the arrival of this pathogen, the USDA and the United Central Soybean Board created a national monitoring system for the disease.

The soybean rust monitoring network consists of several components. A key component is the sentinel plot. Sentinel plots, managed through each state's extension service, are areas of early-maturing soybeans grown specifically to detect rust. Leaf samples from the sentinel plots are sent to processing labs at each state's land grant university. A microscope or magnifying hand lens is used to diagnose the pathogen while it is in a treatable phase of its lifecycle. In 2007, there were more than 700 sentinel plots in operation across the US. The sentinel plot data is made publicly available via the integrated pest management pest information platform for extension and education's website (IPM PIPE). This website displays a map of all the confirmed cases of rust, and farmers can sign up to receive email alerts if a case of rust is detected in their region.

While initial estimates of the magnitude of potential losses due to soybean rust were large, the problem so far has been less severe than anticipated. There have not been any infections of economic consequence in the corn-belt region of the country (which accounts for the majority of US soybean production) since soybean rust came to the US. Because of this, the USDA and United Central Soybean Board have begun to consider the costs and benefits of changing the extent and intensity of the monitoring effort.

Roberts *et al.* (2007) conducted early work on this topic using a farm-level model of decision-making. In their model, the value of the monitoring network stemmed from improved pre-planting decisions based on more accurate soybean rust forecasts. We build upon this work in several ways. First, we develop a dynamic model so we can consider how the value of the information produced by the monitoring network changes over time as producers gain more knowledge about their risk of infection. Second, we add to the management options modeled in Roberts *et al.*: in our model, we add the option to follow a within-season signal provided by the monitoring network. Third, we allow for signal quality to decline with distance from the nearest sentinel plot and use this in our development of a spatially explicit optimization model for the arrangement of a limited number of sentinel plots in the US. Our primary goal is to provide useful information for guiding future investments in the current soybean rust network both in terms of the best distribution of sentinel plots across the landscape and the value of additional years of the monitoring program.

Model

First, we review the strategies available to farmers to manage soybean rust. Roberts *et al.* (2006) considered three strategies: applying a preventative fungicide at the beginning of the susceptible period, scouting for rust and applying a curative fungicide only if rust is found and, lastly, taking no action. We consider these three strategies, and add the possibility of farmers conditioning their strategy choice on the infection status of the sentinel plot nearest to them (within-season signal strategy). Specifically, we assume that farmers will apply a preventative fungicide treatment to their fields only if they observe an infection at the sentinel plot closest to them. If they do not observe an infection, they take no management action. In our model, farmers choose which strategy to follow at the beginning of the growing season. We use the cost estimates from Roberts *et al.* for following the preventative, curative and no action strategy which are shown in Table 2.1. Table 2.1 has two columns, the left shows the cost of the strategy if rust does not come and the right column denotes the cost if it does. The payoff from choosing the within-season signal strategy is a composite of the preventative and no action strategy and will depend on the degree of correlation between the infection status of the sentinel plot and the farmer's field. We choose to represent this with a single signal accuracy parameter, *S*, that varies between 0.5 and 1.

Table 2.1: Cost of Following Management Strategies (dollars/acre)

Management Strategy	No Infection	Infection
Preventative	$25.63	$25.63+1% yield loss
Curative	$6.71	$20.52+7% yield loss
No Action	$0.00	25% yield loss
Follow the within season signal with accuracy, *S*	(S)$0+(1-S)$25.63	(S)*($25.63+1% yield loss)+ (1-S)(25% yield loss)

Choosing the expected cost minimizing strategy, however, requires knowledge of the probability of infection. Because soybean rust is an invasive species, the probability of infection was unknown when it first arrived in the US. Compounding the problem, producers face a dilemma in choosing a management strategy because the strategy they choose affects how much they learn about the probability of infection. Specifically, if a farmer applies a preventative fungicide at the beginning of the susceptible period then there is a significant risk that they might observe a false negative; inoculum that lands on the fields may not develop because of the presence of the fungicide. To capture this tradeoff, we assume that, if no monitoring network is present, farmers do not get information to update their prior beliefs about the infection probability if they apply a preventative fungicide. Due to the nature of soybean rust, false positives are less of a problem because infected leaves have a distinct appearance. Alternatively, using a curative fungicide controls an infection after it occurs and provides an opportunity for the producer to learn how often an infection is likely to occur. The downside to this strategy is that the farmer incurs scouting costs and still experiences nontrivial yield losses.

We conceptualize the soybean rust sentinel plot system as providing two major benefits. The first is that it provides a within-season signal that farmers can condition their management choice upon. Because soybean rust can be difficult to detect early in its life cycle we associate an accuracy parameter, *S*, with the within season signal. This accuracy parameter can be thought of as a function of how spatially correlated infections are and the probability of identifying the disease in time to treat with fungicides. Applying fungicides after the infection is sufficiently established has no economic benefit, so timing is very important. The second benefit that the sentinel plot system provides is that it allows producers to learn about their risk of rust infection even if they apply a preventative fungicide.

Because of the tradeoffs between current returns and future knowledge, we frame the producer's problem in terms of the two-armed bandit problem first discussed in the economics literature by Rothschild (1974). A two-armed bandit problem is a situation in which an agent is in a casino and forced to repeatedly play one of two slot machines (one-armed bandits) each time period for eternity. Each machine yields a payoff, *R*, with probability, *P*, and no payoff with probability (1 - *P*). The probability of a payoff can be different between the two machines. The agent is assumed to have complete information about one machine but no information about the other. The agent's problem is to choose which machine to play each time period to maximize his returns or, more likely, minimize his losses. Rothschild solved this problem using dynamic programming and Bayesian updating and found, when there is a positive discount rate, that the optimal solution takes the form of a stopping rule; a series of sufficiently disappointing outcomes from the unknown slot machine will lead the agent to play the known slot machine for eternity. Rothschild also showed that there is a positive probability that the agent observes such a series of disappointing outcomes from the unknown slot machine and stops playing it even if it yields an objectively higher return.

The basic insights provided by this Rothschild paper are the foundation for the adaptive management literature. A distinguishing characteristic of adaptive management is uncertainty about the properties of the system that is being managed and the ability of different management choices to provide varying levels of information about the nature of the system. As the manager gains information about the system, they use this information to update their prior beliefs (often using Bayesian updating) about the properties of the system. In problems of adaptive management, variables that fully describe a manager's prior beliefs about the system necessarily become state variables in the dynamic optimization problem.

The two armed bandit problem is a useful metaphor for the soybean rust problem. On the one hand, producers can apply a preventative fungicide and get a certain return, but learn nothing about the risk of soybean rust. On the other hand, they can apply a curative fungicide only if they observe a rust infection and learn about their risk. The expected cost minimizing choice depends on the probability that rust will infect their farm in a given year. In this problem, the preventative fungicide strategy is equivalent to the agent playing the slot machine with the known return and the curative and no action strategies are like playing the slot machine with an unknown return. In this context, one benefit of the soybean rust monitoring network is that it allows an agent to apply the preventative fungicide but still receive information about what his or her return would have been had they chosen to apply a curative fungicide if soybean rust were observed. The monitoring network reduces the probability that agents will make poor management decisions because it allows for greater learning and, while it is being provided, eliminates the tradeoff between current returns and future learning.

To better join the problem of soybean rust and the two-armed bandit model, we assign a constant payoff to the preventative strategy and allow for an explicit strategy of following a within season signal. If the payoffs to the preventative strategy were different when rust came and when it didn't, the farmer could deduce the outcome of that year and update his or her prior beliefs. Field trials cited by Roberts *et al.* (2007) found an average of 1% yield loss if rust spores infect a field treated with a preventative fungicide. Because there are many phenomena that can cause a 1% yield loss, it would be difficult for a farmer to know whether rust spores arrived at their farm in years when they apply a preventative fungicide just from their yield. Therefore, it is safe to assume that the payoff from a preventative strategy is invariant whether or not soybean rust arrives.

The farmer's problem can be formulated as a dynamic programming problem where the state variable is the farmer's belief about the annual probability of a rust infection p. To make this problem tractable, it is important to find a parsimonious way of representing the farmer's belief. To do this, we assume that the probability that a farmer will observe Y infections in Z years will follow a binomial distribution with parameter p. When we make this assumption, the beta distribution is an ideal way to describe a famer's beliefs about the value of p for two reasons. First, the beta distribution is defined on the interval (0, 1) and thus is constrained to be valid values for p. Second, when a beta distribution is used to describe farmer's prior belief about the value of p, then the posterior distribution that the farmer has after updating will also be a beta distribution; the beta distribution is a conjugate prior with respect to the Bernoulli and binomial distributions. This is an important property because it allows a farmer's belief about the value p to be completely described by two parameters, M and N, in *all* time periods of the model.

The probability density function of the beta distribution

$$pdf(p) = beta(M, N) = \frac{(1-\mathrm{p})^{\mathrm{N}-1}\mathrm{p}^{\mathrm{M}-1}}{\mathrm{Beta[M,N]}} \tag{1}$$

where beta(M,N) refers to a beta distribution and Beta[M,N] is the Euler beta function. Also, the Bayesian updating that occurs as farmers observe new information is simple to express. Specifically, given a prior belief about p described by a beta distribution with parameters $\{M, N\}$, if a famer observes an infection in the current period then their posterior belief from applying Bayes rule is given by beta distribution with parameters $\{M + 1, N\}$. Conversely, if the farmer observes no infection, the posterior belief is described by the parameters $\{M, N + 1\}$. A final statistical preliminary is to be able to express farmer's expectation over p when their prior beliefs are described by a beta distribution:

$$E[p] = \frac{M}{M+N} \tag{2}$$

The Producer's Problem when no Monitoring Network Exists

Now we can formulate the farmer's decision problem when there is no monitoring network as a dynamic program with three strategies. Recall that there is no within-season strategy because there is no information provided by a monitoring network. The optimal value function, V, is a function of the parameters *M* and *N*:

$$V[M, N] =$$

$$Max[\pi_{N1}\frac{M}{M+N} + \pi_{N0}\left(1 - \frac{M}{M+N}\right) + \delta\left(V[M+1, N]\frac{M}{M+N} + V[M, N+1]\left(1 - \frac{M}{M+N}\right)\right),$$

$$\pi_{C1} * \frac{M}{M+N} + \pi_{C0}\left(1 - \frac{M}{M+N}\right) + \delta\left(V[M+1, N]\frac{M}{M+N} + V[M, N+1]\left(1 - \frac{M}{M+N}\right)\right),$$

$$\pi_P + \delta\, V[M, N]] \tag{3}$$

where π_{N1} and π_{N0} represent the one period return from following the no action strategy with and without a rust infection; π_{C1} and π_{C0} represents the one period returns from following the curative strategy with and without a rust infection; π_P represents the return from following the preventative strategy; and δ is the discount rate. The important part of this formulation is to note that choosing the preventative strategy prohibits the farmer from gaining additional information about their risk of infection. A reasonable initial prior for a farmer with no information is to assume that *p* is uniformly distributed over the interval [0, 1]. This is equivalent to a beta distribution with $M = N = 1$. As Rothschild showed in his paper on bandit problems, once the preventative strategy is optimal, it will be chosen for all remaining time periods.

The Producer's Problem when a Monitoring Network Exists

Next, we turn our attention to how the existence of the monitoring network would change the farmer's problem. From an informational standpoint, the soybean rust monitoring network allows farmers to gather evidence about their risk of infection regardless of the management strategy they choose. The monitoring network also provides a within season signal upon which farmers can condition their actions. In the context of the two armed bandit problem, the monitoring network allows a producer to see what return would have been earned had he played the unknown arm. To place a value on the monitoring network, farmers need to know when and how long the monitoring network will be in existence. For instance, a monitoring network that exists during years one to three will be more valuable than a network that only exists during years one and two. Likewise, because of discounting, a monitoring network that exists during years one to three will be more valuable than one that exists during years four to six. Because of this, the dynamic program representing the farmer's problem with information will require more notation than the farmer's problem without information. To simplify this notation, we assume that the monitoring network always begins in year one and lasts a finite number of periods given by the parameter L. To measure the value of a monitoring network that exists for a specified number of periods, we have to re-solve the dynamic programming problem allowing farmers to observe evidence when choosing the preventative strategy only during periods when the monitoring network exists. After solving this modified problem, the value of the information can be found by comparing the two different value functions. We represent the farmer's value function without information as $V_t[M, N]$ to mean the value of arriving in year t with the prior described by M and N. We represent the farmer's value function with information as $V1_t[M, N, L, S]$ to mean the value of arriving in year t with priors M and N, a monitoring network that exists during years 1 through L and a within season signal quality of S.

An efficient way of calculating the value function with the monitoring network is to make use of the previously calculated value functions from the farmer's original decision problem. The key insight is that the value function for the decision problem with information becomes the same as the original problem once the monitoring network is discontinued. Mathematically, the functional equation can be written:

$$V1_t[\mathrm{M}, \mathrm{N}, L, \mathrm{S}] = \begin{cases} Max[\Pi_N(M,N), \Pi_C(M,N), \Pi_P(M,N), \Pi_{NP}(M,N,S)], & \text{for } t < L \\ V_t[\mathrm{M}, \mathrm{N}], & \text{for } t \geq L \end{cases}$$

where

$$\Pi_N(M,N) = \pi_{N1}\frac{M}{M+N} + \pi_{N0}\left(1 - \frac{M}{M+N}\right) + \delta\left(V1_{t+1}[\mathrm{M}+1, \mathrm{N}, L]\frac{M}{M+N} + V1_{t+1}[\mathrm{M}, \mathrm{N}+1, L]\left(1 - \frac{M}{M+N}\right)\right),$$

$$\Pi_C(M,N) = \pi_{C1}\frac{M}{M+N} + \pi_{C0}\left(1 - \frac{M}{M+N}\right) + \delta\left(V1_{t+1}[\mathrm{M}+1, \mathrm{N}, L]\frac{M}{M+N} + V1_{t+1}[\mathrm{M}, \mathrm{N}+1, L]\left(1 - \frac{M}{M+N}\right)\right),$$

$$\Pi_P(M,N) = \pi_P + \delta\left(V1_{t+1}[\mathrm{M}+1, \mathrm{N}, L]\frac{m}{m+n} + V1_{t+1}[\mathrm{M}, \mathrm{N}+1, L]\left(1 - \frac{M}{M+N}\right)\right),$$

and

$$\Pi_{NP}(M,N,S) = \frac{M}{M+N}(\mathrm{S}\,\pi_P + (1-\mathrm{S})\pi_{N1}) + \left(1-\frac{M}{M+N}\right)(\mathrm{S}\,\pi_{N0} + (1-\mathrm{S})\pi_P)$$

$$+\,\delta\left(V1_{t+1}[\mathrm{M}+1,\mathrm{N},L]\frac{M}{M+N} + V1_{t+1}[\mathrm{M},\mathrm{N}+1,L] * \left(1-\frac{M}{M+N}\right)\right).$$

Having these two models of decision making allows us to calculate the value of information generated by the soybean rust sentinel plot system. Conceptually, the value of information is the expected cost of managing soybean rust with information minus the cost without information. Mathematically, this is equivalent to $V1_1[M,N,L,S]$ - $V_1[M,N]$.

An interesting question is, when no monitoring network is present, when will farmers choose the preventative strategy and stop learning for all remaining time periods? We approximate an infinite horizon problem assuming a 90 period finite horizon problem and using backwards induction to solve. For the example below, we assumed a soybean price of $8.00/ bushel and a yield of 37 bushels/acre. Figure 2.1 shows the first five years of optimal strategies at each time period given the farmer's beliefs over the probability of infection. The vertices of the decision are labeled according to the state variables in the dynamic programming problem. For instance, the first vertex is labeled {1,1} which represent the parameter values of a beta distribution that makes it equivalent to a uniform distribution. Branches to the left represent observed infections and branches to the right are non-infections. This figure shows that the first node where it is optimal for farmers to choose the preventative strategy is upon reaching a prior belief described by beta {4,1}. This corresponds to the first 3 observations being infections. Alternatively, if a farmer reaches a state space of {1,6}, then it is optimal for him or her to choose the no action strategy. This corresponds to the first 5 observations being non-infections. For the majority of the states it is optimal for the farmer to choose the curative strategy. These results lend credibility to our model of producer behavior because, in regions of the country where there have been 3 years of non-infections, anecdotal information from extension agents indicates that the majority of farmers are not applying preventative fungicides for soybean rust and are spending little if any time scouting for soybean rust. Conversely, farmers in many areas in the Deep South are regularly applying a fungicide before soybean rust is specifically detected, though it should be noted there are several other fungal pests in the Deep South that can motivate fungicide treatment.

Next, we turn our attention to the decision tree when a monitoring network is present. Figure 2.2 shows optimal management decisions when a monitoring network will be present for 4 years and there is a signal with quality 0.75. At the start of this tree, the optimal management decision is to use the within season signal. This is the case when the probability of infection is closer to 0.5 than either extreme and there is a sufficiently high signal quality. Using the within season signal allows farmers to outsource their scouting costs and obtain very good yield protection. Looking at the vertex labeled {2,1}, we see that the best management decision is to apply a preventative fungicide application. Because the monitoring network is present for the first 4 years in this example, the farmer can continue to learn while doing so. Once the monitoring network ceases to exist, choosing the preventative strategy halts learning as can be seen at the vertex {5,1}.

Figure 2.1: Optimal Management Decisions for a Farmer without a Monitoring Network

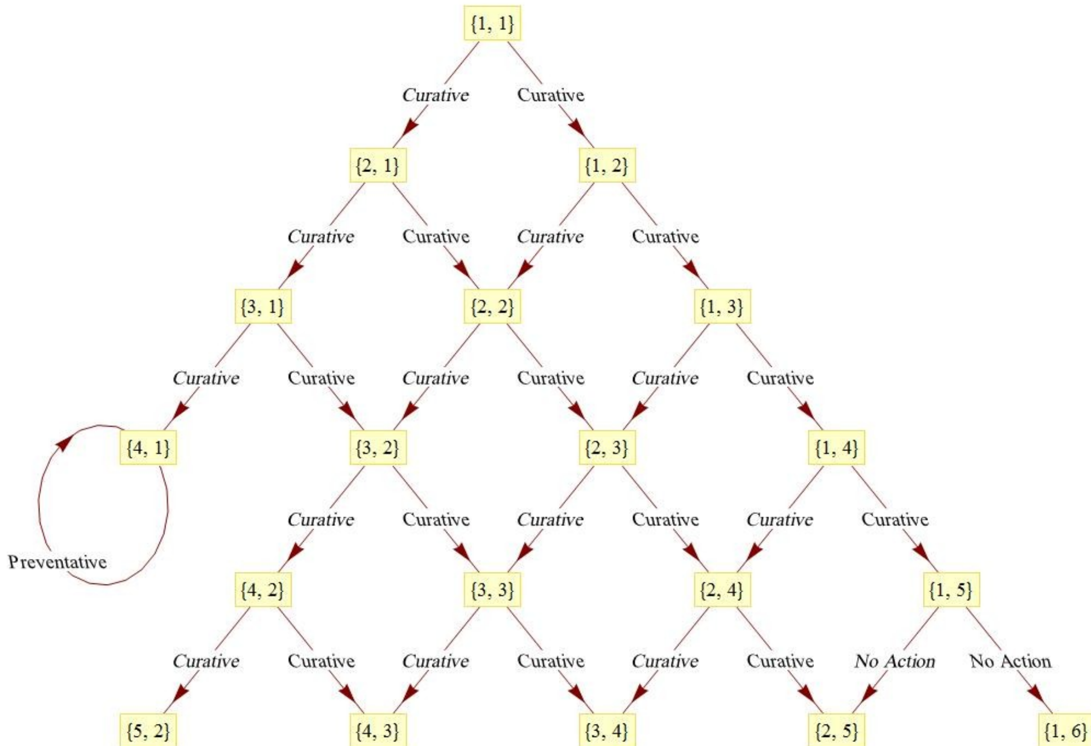

Having presented a dynamic model capable of valuing a monitoring network, we apply this model to the problem of placing a limited number of sentinel plots to maximize the value of information. Using county level data on average soybean acreage and yield from 2005-2007. This data set allows us to construct a representative farmer for each county that will form the basic unit of the model. We use county level estimates of the probability of soybean rust infection from Bekkerman, Goodwin and Piggott (2008) for farmers' prior beliefs on the likelihood of a rust infection in their county. To transform the estimates from Bekkerman, Goodwin and Piggott (2008) into a valid beta distribution for use as a prior, we calculate the M and N such that $\frac{M}{M+N}$ was equal to the probability estimate and that $M+N = 6$ which ensures that all the counties treat their priors based upon the same number of hypothetical observations.

Figure 2.2: Optimal Management Decisions for a Farmer with a 4-Year Monitoring Network with Signal Quality of 0.75

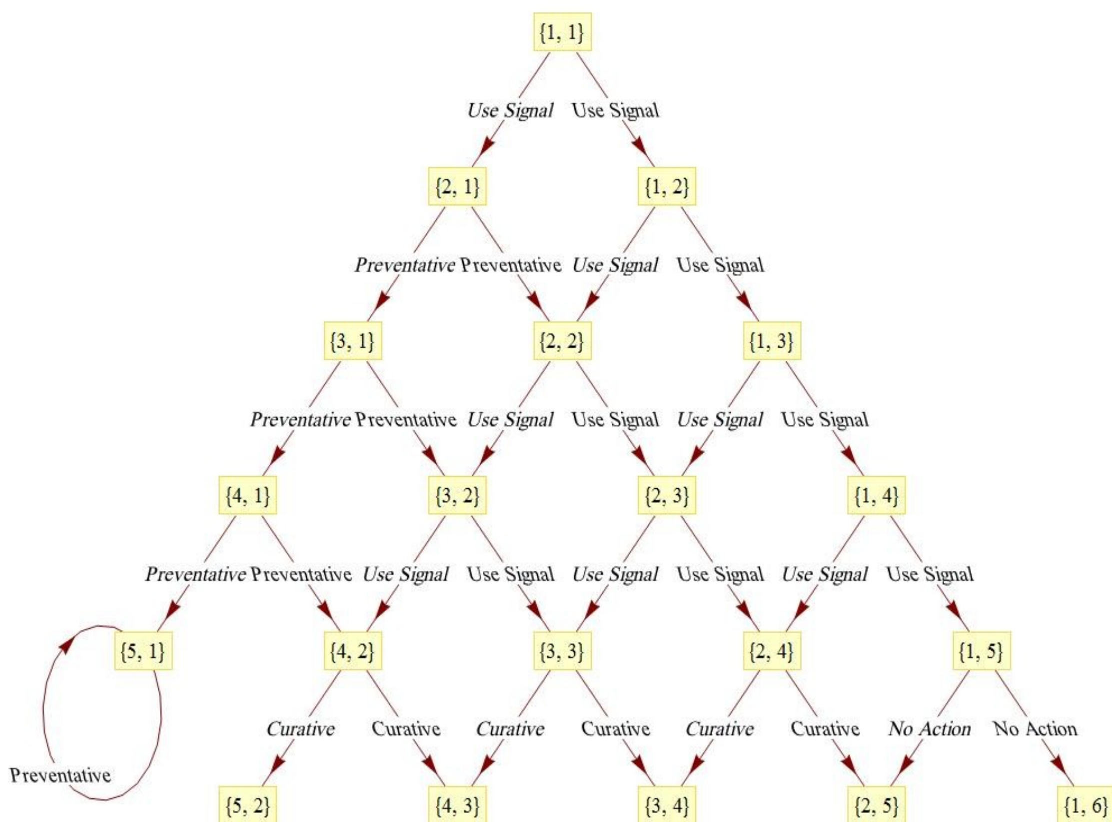

The last piece of information that needs to be discussed is how the signal quality that farmers receive depends on the spatial arrangement of sentinel plots. We choose to model signal quality as a function of the distance between the farmer and the sentinel plot nearest to them. The rate at which signal quality declines with distance should primarily be a function of the spatial autocorrelation of the disease spread. Given the wind borne nature of the disease, there is reason to expect that significant spatial autocorrelation exists. However, we are unaware of any procedure in the literature to measure spatial auto-correlation of point processes. Because of this we consider three scenarios (high, medium, and low) where signal quality decreases at different rates as a function of distance. These scenarios are shown in Figure 2.3 where the top line (blue) denotes a high signal quality function, the middle line (red) denotes the medium signal quality function and the bottom line (yellow) denotes the low signal quality function. Note that the signal quality function stops at 0.5 because, below that, the signal quality would be negatively correlated with the outcome of interest.

Figure 2.3: Signal Quality as a Function of Distance

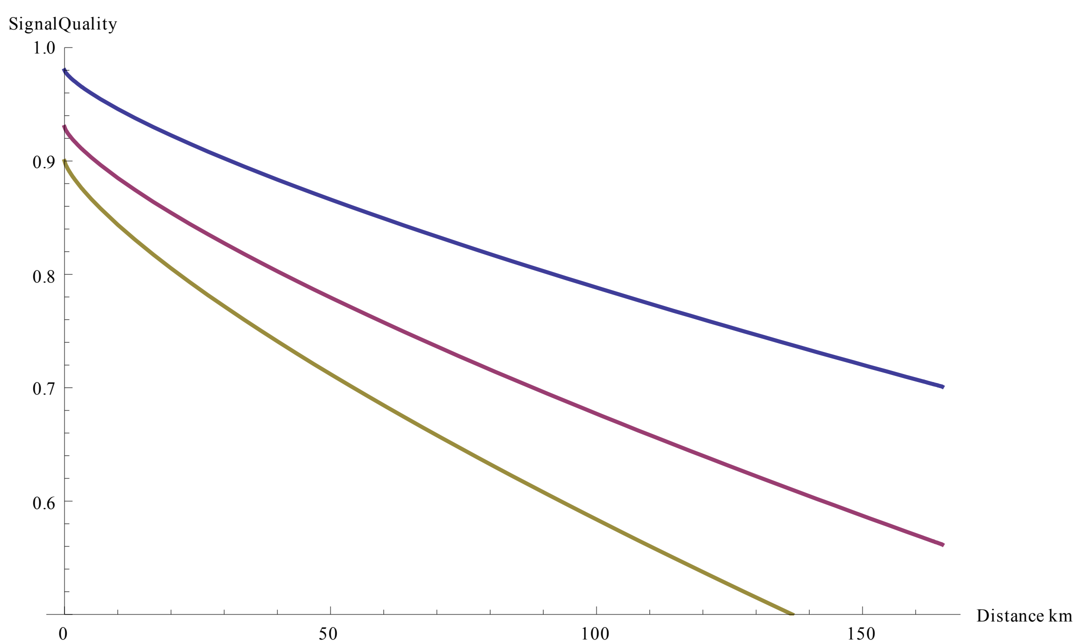

Using the county level data and the model of farmer decision making described above, we ran a series of optimizations with the goal of placing a fixed number of sentinel plots to maximize the value of information generated from them. The first step in this process was to calculate the value function $V_1[M, N]$ for each county taking into account county level differences in soybean yield acreage, and probability of infection. Next, for each county we determined the closest sentinel plot in the trial solution. Using the distance between the county and the nearest sentinel plot we calculated the signal quality experienced by each county. With this information it is now possible to calculate $V_1[M, N, L, S]$ for each county. The value of information produced by a particular trial solution then is the quantity $V_1[M, N, L, S]$- $V_1[M, N]$ summed across all of the counties in the model.

Having described how to calculate the value of information for a particular trial solution, we next turn to the algorithm for generating trial solutions. The algorithm used for generating trial solutions is a type of genetic algorithm as described in Eiben and Smith (2003). The idea of the algorithm is to represent each candidate solution as an array of numbers representing the counties where a sentinel plot is to be placed. In our dataset, we identified approximately 1600 counties with significant soybean production, thus the numbers in the array are constrained to be in the range 1 - 1600. The number of elements in the array denotes the number of sentinel plots that are part of the solution.

The basic idea of the genetic algorithm is to start with an initial population and to then create a next generation by applying a mutation operation to each of the candidate solution in the initial population. Specifically, the mutation operation we used was to drop *R* sentinel plot locations from the candidate solution and then add *R* random sentinel plot locations. Checks were put into place to ensure that no candidate solution had any duplicate numbers indicating that two sentinel plots should be placed in the same location. For each candidate solution, we applied the mutation operator to it twice yielding two new "offspring" solutions. The value of information for all of the offspring solutions was calculated and only the top 50% were kept for the next iteration. In this way the evolutionary force of survival-of-the-fittest is used to guide the population toward the optimal solution.

The number of mutations, *R*, that are made can be thought of as a search radius. If many sentinel plot positions are dropped and new ones added this can be thought of as a wide search radius, while if only a few are changed this is a smaller search radius. When the genetic algorithm first starts, it is likely that the population is relatively far from the optimal solution and thus a wide initial search radius is advisable. As the algorithm runs and the population gets closer to the optimal solution, a wide search radius is no longer optimal because changing several sentinel plot positions at a time is likely to move you further from the optimal solution than closer to it. Because of this we use an algorithm that starts with a wide search radius (number of sentinel plot changes) and decreases the search radius if the number of iterations without an improvement in the best solution found exceeds a threshold. When the number of sentinel plot changes done by the algorithm reaches 0, then the algorithm terminates and returns the best solution found.

The genetic algorithm described above was used to search for the optimal arrangement of a 200, 300, 400 and 500 sentinel plots. The algorithm was run three times for each number of sentinel plots to ensure a close approximation to the optimal solution. Before conducting these runs several exploratory runs were conducted to identify the best parameters to govern the genetic algorithm. These runs showed that starting with an initial population of 200 organisms and 10 mutations provided good results. After 200 consecutive iterations without improvement the number of mutations was decreased by 1.

Results

The value of extending the soybean rust sentinel plot system for one additional year beyond the six years of historical monitoring using the optimal spatial arrangement is shown in Table 2.2 for a number of different signal quality and plot density combinations. Benefit estimates range from $57-117 million dollars. This is in the middle of the range estimated by Roberts *et al.* who estimated the benefits of a year of monitoring to be between $5 - 300 million dollars. It is important to point out that our benefits estimate is the *present value* of an additional year of monitoring. In our model of farmer decision making, knowledge is a capital good and can yield benefits many years after it is acquired.

Table 2.2: Present Value of an Additional Year of Soybean Rust Monitoring

	High Quality Signal	Medium Signal Quality	Low Signal Quality
500 plots	$116,849,525	$80,801,407	$66,346,082
400 plots	$112,533,053	$79,638,784	$65,128,624
300 plots	$106,998,465	$77,237,245	$62,652,926
200 plots	$99,381,341	$71,512,389	$56,997,464

The optimal spatial arrangement of 200 sentinel plots under the high signal quality function scenario is shown in Figure 2.4. The counties where the model recommends a sentinel plot be place are highlighted. This corresponds closely to counties that have the highest soybean production and is substantially different from allocation of sentinel plots for the 2010 growing season shown in Figure 2.5. In 2010, the majority of sentinel plots were placed in the Southern US where the probability of infection is relatively high, but where only a small fraction of US soybean production occurs. In contrast, our model recommends placing substantial resources in the corn-belt where the probability of infection is low, but where production levels are higher. The optimal spatial arrangements of sentinel plots under different signal quality functions were qualitatively similar. The different signal quality scenarios seemed to primarily affect the value of the network, not the optimal arrangement.

Figure 2.4: High Signal Quality 200 Sentinel Plot Locations

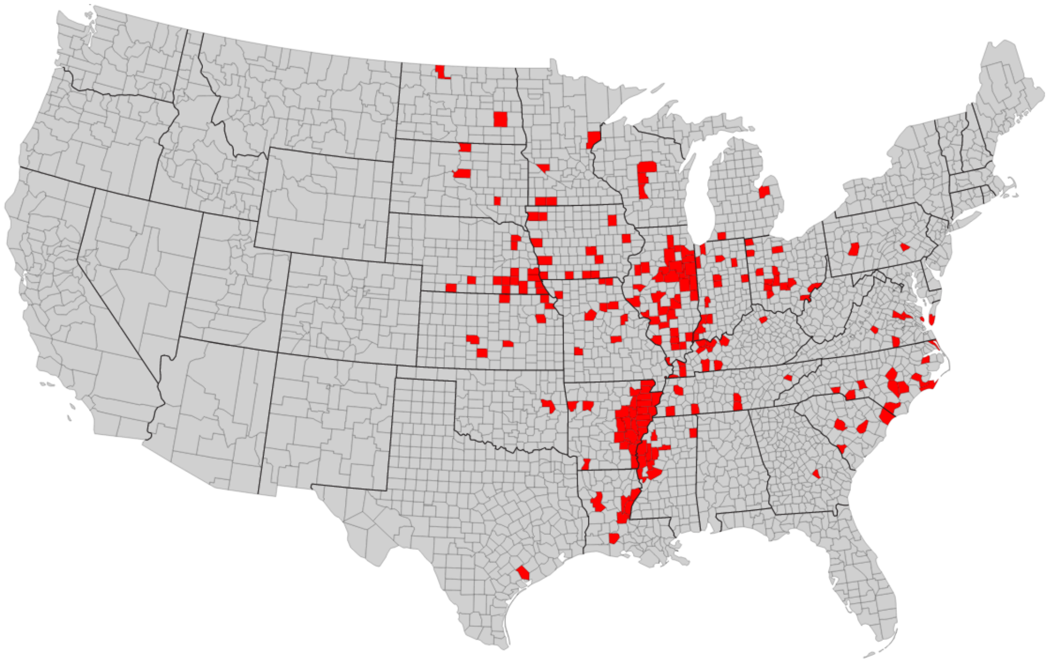

Figure 2.5: Sentinel Plot Locations for 2010 Growing Season

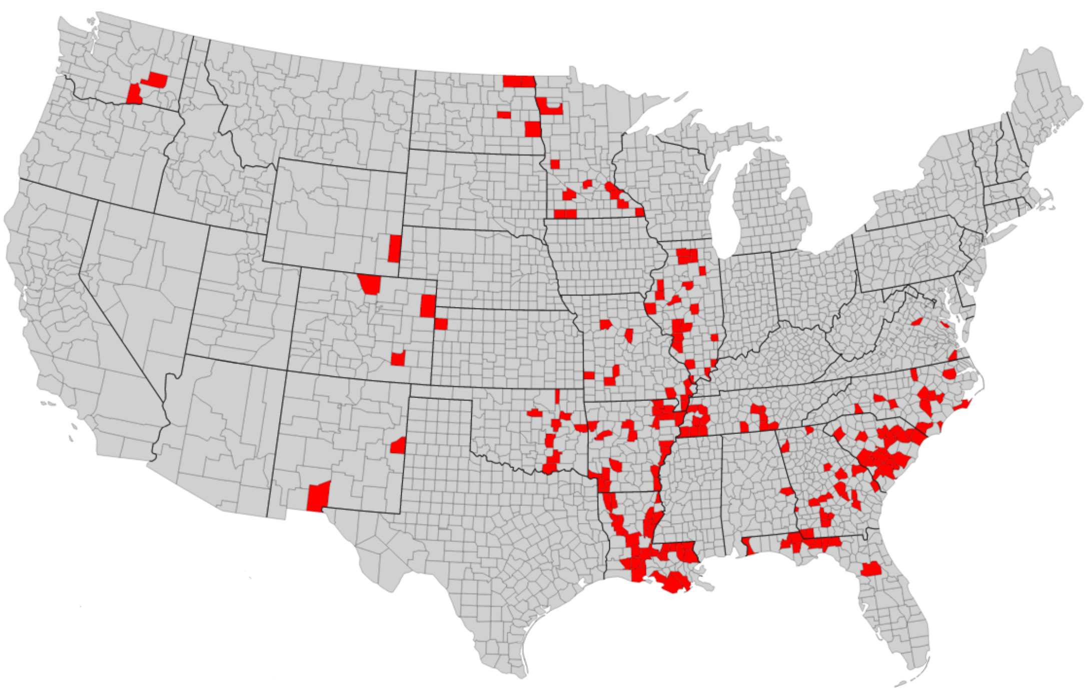

Lastly, we investigate how the marginal value of extending the sentinel plot system changes as more knowledge is accumulated. The marginal benefit of extending the life of the network from L years to $L + 1$ years is $V1_t[M, N, L + 1, S]$- $V1_t[M, N, L, S]$. The marginal benefit of extending the life of the network for three different levels of signal accuracy is shown in Figure 2.6. This figure was constructed assuming that farmer's had an initial prior described by a uniform distribution, \$8 per bushel soybean price and a yield of 37 bushels per acre. From this figure we see that the marginal benefit of extending the network depend critically on the accuracy of the within season signal. When the accuracy of the within season signal is 0.5, then the only benefit of extending the monitoring network is the increased opportunity for learning while applying the preventative fungicide because the within season signal does not provide any additional information beyond the farmer's initial prior belief. If limited resources were available for processing leaf samples, the monitoring network might not be able to publish results in time for farmers to effectively condition their management strategies upon the results, but farmers could still update their prior beliefs from the results when they are eventually published if they choose to apply a preventative fungicide. In this case, the benefits of extending the network decline rapidly. When there is a within season signal, however, the marginal benefits of extending the network decline much more slowly.

An interesting feature of this graph is that the marginal value of additional information with a high signal quality is always declining, this is not the case with a low signal quality. We believe that this is because with a high signal quality the majority of the benefits come from being able to follow the within season signal and that these benefits decline because of discounting. With a low signal quality, however, the value of learning is the primary source of gains. For the marginal returns to increase means that the returns from knowledge exceed the discount rate in the model. While surprising at first, that the marginal value of information may increase and then decrease is not unusual. For instance, the marginal returns from completing the third grade may exceed those of completing second in some cases.

Figure 2.6: Marginal Value of Extending Soybean Rust Sentinel Plot System

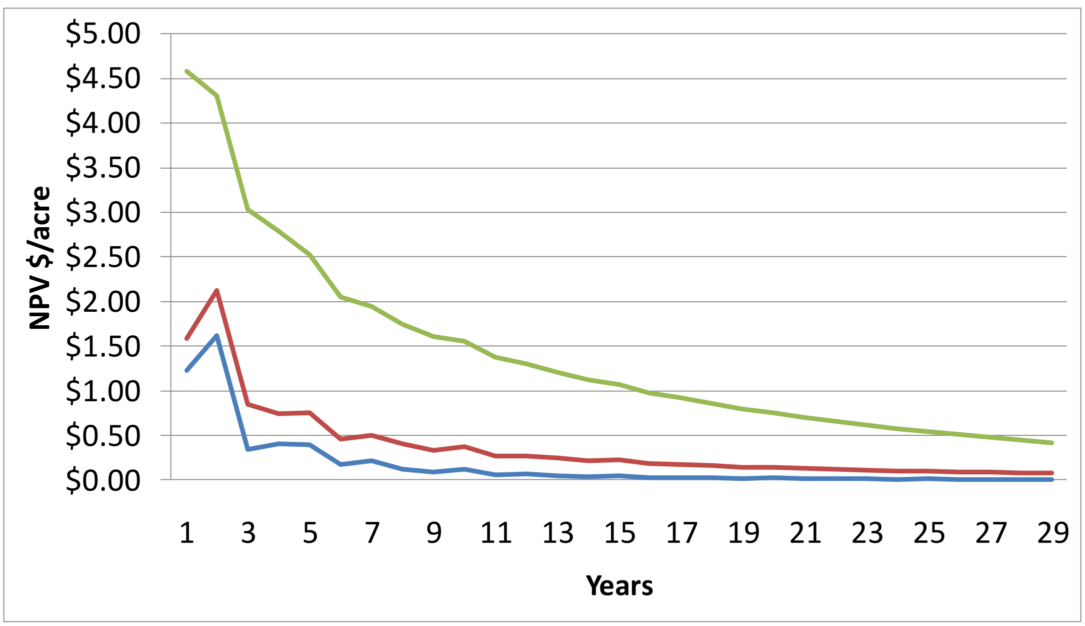

Conclusion

In this paper, we developed a dynamic model of the farmer's decision regarding the management of soybean rust. We used this model to value the US soybean rust sentinel plot system and to optimize the spatial arrangement of sentinel plots. We estimate that the value of extending the life of sentinel plot system for one year, given the optimal spatial arrangement of plots, to be between \$52 - 106 million dollars depending on the number of plots and assumptions about the quality of the within-season signal. This is an order of magnitude greater than the cost of maintaining the system for another year. Also, the optimal spatial arrangement of sentinel plots indicated by our model is substantially different from the planned placement for the 2010 growing season. The current plan places no sentinel plots in the corn-belt region of the country, while our model indicates that some coverage is advantageous. Looking forward, we estimated the marginal value of extending the life of the sentinel plot system for multiple years. We find that the marginal value declines as more knowledge is accumulated, but the rate of decline is highly sensitive to the within-season signal quality parameter. When there is a low quality within-season signal, there are few benefits to having the network beyond the first decade.

Future work in this area should focus on better understanding the spatial auto-correlation of infections. The problem is a challenging one because a number of environmental factors must coincide for an infection to occur; soybeans must be present, innoculum must be deposited and conditions for the days following inoculation must be damp and cloudy as ultra violet radiation and dry conditions can kill soybean rust spores. If the data could be collected, the problem of measuring the spatial auto-correlation of a point processes would entail new work. A possibility is to treat the problem using a latent variable framework and positing a continuous underlying distribution as is commonly done in spatial statistics work but assuming that we only observe an infection when this continuous latent variable exceeds a threshold. While attractive in a theoretical sense, this would be computationally demanding as estimating the spatial-auto correlation via maximum likelihood would require the evaluation of integrals with high dimensionality.

Conclusion

This dissertation has presented two papers on the topic of risk-modifying public goods. The first paper shows that while providing an efficient level of a traditional public good while balancing a budget is an impossibility, there exist a family of social choice functions implementable in dominant strategies that can provide an efficient level of a risk-modifying public good while balancing the budget. There are two key features that make this outcome possible. The first is that when an agent is fully insured, they are indifferent to the state-of-the-world that occurs and thus to the level of the public good that is provided. The second is that the presence of a competitive financial market that is willing to be counter-party to any actuarially fair bets is essential to providing full insurance in all states-of-the-world.

The second paper presented empirical work relating to the value provided by a real-world example of a risk-modifying public good. A key result of this work is that the optimal placement of sentinel plots recommended by the model is significantly different from the plan adopted by policy makers. The model's recommended sentinel plot placement achieves a relatively even coverage across the Eastern United States while the 2010 plan adopted by policy makers focuses on Southern States. An advantage of the model is that it explicitly considers the tradeoff between low-frequency high-severity events (infections in the corn-belt) and high-frequency low-severity events (infections in southern states).

References

Abreu, Dilip and Arunava Sen, “Virtual Implementation in Nash Equilibrium,” *Econometrica*, 59 (1991), 997-1021.

Arrow, Kenneth J., “The Role of Securities in the Optimal Allocation of Risk-Bearing,” *Review of Economic Studies*, 31 (1964), 91-96.

Barbera Salvador and Matthew Jackson, “A Characterization of Strategy-Proof Social Choice Functions for Economies with Pure Public Goods, “ *Social Choice and Welfare*, 11 (1994), 241-252.

Bekkerman, A., B. Goodwin, N. Piggott. 2008. “Spatio-temporal Risk and Severity Analysis of Soybean Rust in the United States” *Journal of Agricultural and Resource Economics*. 33(3):311-331.

Bergemann, Dirk and Juuso Valimaki, “Information Acquisition and Efficient Mechanism Design”, *Econometrica*, 70 (2002), 1007-1033.

Chavas, Jean-Paul and Jay S. Coggins, “On Fairness and Welfare Analysis Under Uncertainty,’" *Social Choice and Welfare*, 20 (2003), 203-228.

Duggan, John, “Virtual Bayesian Implementation,” *Econometrica*, 65 (1997), 1195-1199.

Ehrlich, Isaac and Gary S. Becker, “Market Insurance, Self-Insurance, and Self-Protection,” *Journal of Political Economy*, 80 (1972), 623-48.

Eiben, A.E. and J.E. Smith. 2003. Introduction to Evolutionary Computing. Berlin: Springer-Verlag.

Federal Report. “Economic and Environmental Principles and Guidelines for Water and Related Land Resources Implementation Studies.” Washington, D.C. 1983.

Graham, Daniel A., “Cost-Benefit Analysis Under Uncertainty.” *American Economic Review*, 71 (1981), 715-725.

Isard, S.A. and Gage, S.H. and Comtois, P. and Russo, J.M. 2005. “Principles of the Atmospheric Pathway for Invasive Species Applied to Soybean Rust.” *BioScience*. 55(10):851-861.

Isard, S.A. and Russo, J.M. and Ariatti, A. 2007. “The Integrated Aerobiology Modeling System applied to the spread of soybean rust into the Ohio River valley during September” *Aerobiologia* 23(4): 271-282.

Jackson, Matthew O., “Bayesian Implementation.” *Econometrica*, 59 (1991), 461-477.

Krishna, Vijay and Motty Perry, “Efficient Mechanism Design,” Unpublished, 2000.

Kuchler, F., Duffy, M., Shrum, R. D., and W. M. Dowler. 1984. “Potential economic consequences of the entry of an exotic fungal pest: the case of soybean rust.” *Phytopathology* 74:916-920.

Livingston, M.J., R. Johansson, S. Daberkow, M. Roberts, M. Ash, and V. Breneman. “Economic and Policy Implications of Wind-Borne Entry of Asian Soybean Rust into the United States”. Washington DC: U.S. Department of Agriculture, Economic Research Service, OCS04D02, April 2004.

Neeman, Zvika, “The Relevance of Private Information in Mechanism Design,” *Journal of Economic Theory*, 17 (2004), 55-77.

Pazner, Elisha A. and David Schmeidler, "Egalitarian Equivalent Allocations: A New Concept of Economic Equity," *Quarterly Journal of Economics*, 92, (1978), 671-687.

Roberts, M.J., D. Schimmelpfennig, E. Ashley, and M.J. Livingston. 2006. "The Value of Plant Disease Early-Warning Systems." Economic Research Report No. 18, Economic Research Service, USDA, April 2006.

Rothschild, M. 1974. "A two-armed bandit theory of market pricing." *Journal of Economic Theory* 9(2):185 – 202.

Samuelson, Paul A., "The Pure Theory of Public Expenditure," *Review of Economics and Statistics*, 36 (1954), 387-89.

Serizawa, Shigehiro, "Strategy-Proof and Symmetric Social Choice Functions for Public Good Economies," *Econometrica*, 67 (1999), 121-145.

Walker, Mark, "On the Nonexistence of a Dominant Strategy Mechanism for Making Optimal Public Decisions," *Econometrica*, 48 (1980), 1521-1540.

Zhou, Lin, "Nash Implementation in Pure Public Good Economies," *Economics Letters*, 99 (2008) 470--473.