Ecological Economics 86 (2013) 78–85

Contents lists available at SciVerse ScienceDirect

Ecological Economics

journal homepage: www.elsevier.com/locate/ecolecon

Analysis

Optimal strategies for the surveillance and control of forest pathogens: A case study with oak wilt

Tetsuya Horie [a], Robert G. Haight [b], Frances R. Homans [c,*], Robert C. Venette [b]

[a] Sophia University, Graduate School of Global Environmental Studies, 7-1 Kioi-cho, Chiyoda-ku, Tokyo 102-8554, Japan
[b] USDA Forest Service, Northern Research Station, 1992 Folwell Avenue, St. Paul, MN 55108, USA
[c] Department of Applied Economics, University of Minnesota, 1994 Buford Ave., 217G Ruttan Hall, Saint Paul, MN 55108, USA

ARTICLE INFO

Article history:
Received 30 November 2011
Received in revised form 26 September 2012
Accepted 30 September 2012
Available online xxxx

Keywords:
Forest pathogens
Surveillance
Optimal management
Uncertainty
Ceratocystis fagacearum

ABSTRACT

Cost-effective strategies are needed to find and remove diseased trees in forests damaged by pathogens. We develop a model of cost-minimizing surveillance and control of forest pathogens across multiple sites where there is uncertainty about the extent of the infestation in each site and when the goal is to minimize the expected number of new infections. We allow for a heterogeneous landscape, where grid cells may be differentiated by the number of trees, the expected number of infected trees, rates of infection growth, and costs of surveillance and control. In our application to oak wilt in Anoka County, Minnesota, USA, we develop a cost curve associated with saving healthy trees from infection. Assuming an annual infection growth rate of 8%, a \$1 million budget would save an expected 185 trees from infection for an average of \$5400 per tree.
We investigate how more precise prior estimates of disease and reduced detection sensitivity affect model performance. We evaluate rules of thumb, finding that prioritizing sites with high proportions of infected trees is best. Our model provides practical guidance about the spatial allocation of surveillance and control resources for well-studied forest pathogens when only modest information about their geographic distribution is available.

1. Introduction

Non-native fungal pathogens have widespread effects on forests and threaten many tree species. Chestnut blight (*Cryphonectria parasitica*) virtually eliminated the American chestnut (*Castanea dentata*), and now the Dutch elm disease fungus (*Ophiostoma novo-ulmi*) and the oak wilt fungus (*Ceratocystis fagacearum*) are spreading in the eastern United States (Hansen, 2008; Loo, 2009). For many forest pathogens, managers do not know the extent of infection because the disease is difficult to detect (Meentemeyer et al., 2008). Further, an effective means of controlling spread is prompt identification and removal of diseased trees. Because budgets for surveillance and removal are typically limited, managers must decide how to allocate those funds among sites with different features. For example, should locations with low expected numbers of infected trees be given priority since these locations can more easily become disease-free, or should locations with low surveillance costs be given precedence because more diseased trees can be identified with the same budget? In addition, managers face a tradeoff between surveillance and control activities. Greater spending for surveillance enables managers to identify more infected trees, but increased surveillance limits the budget remaining for removal of diseased trees.

Researchers use both simulation and optimization tools to evaluate strategies for surveillance and control of non-native forest insects and pathogens. Insights from these models depend on the underlying assumptions about the invasion process and the decision-making environment, which differ across studies. Simulation models include detailed representations of the processes likely to affect insect or pathogen spread, including human-mediated dispersal, and predict the effectiveness of a few pre-defined surveillance and control strategies (Harwood et al., 2011; Kovacs et al., 2011; Økland et al., 2010). For example, Harwood et al. (2011) use a spatial-dynamic simulation model to reconstruct an epidemic of a newly introduced Dutch elm disease fungus, which swept through Great Britain in the 1970s and killed the majority of mature elm trees (*Ulmus* spp.). They predict the impact of proposed counterfactual policies of surveillance and control at a national scale and conclude that more rapid and intense efforts would not have prevented elm decline because of the scale and density of the elm population nationwide. While simulation models such as this include the spatial structure needed to predict insect or pathogen spread, they are often difficult to parameterize because of a limited understanding of the spatial processes. Further, they are designed to examine the relative value of a predefined set of policies rather than find the optimal policy.

Optimization models, in contrast, typically include the dynamics of an infestation while simplifying or ignoring spatial structure. These models determine the optimal timing and intensity of surveillance

* Corresponding author.
E-mail address: fhomans@umn.edu (F.R. Homans).

0921-8009/\$ – see front matter
http://dx.doi.org/10.1016/j.ecolecon.2012.09.017

assuming that the timing of detection is uncertain and eradication or suppression can be conducted only after the infestation is found (Bogich et al., 2008; Homans and Horie, 2011; Mehta et al., 2007; Ndeffo Mbah and Gilligan, 2010). The implications are that damage caused by the infestation continues to occur during surveillance and that older invasions are larger and more difficult to suppress or eradicate. Higher intensities of surveillance enable earlier control and less damage but cost more. This is a central tradeoff involved in determining the optimal intensity of surveillance (Epanchin-Niell and Hastings, 2010).

Another type of dynamic optimization uses a partially observable Markov decision process framework to select surveillance and control activities when the level, observation, and dynamics of the infestation are uncertain and the objective is to minimize total discounted costs (Haight and Polasky, 2010). In this framework, the decision (e.g., monitor, treat, or do nothing) is conditional on the beliefs about the level of infestation. The beliefs, in turn, are updated using the transition probabilities of the infestation, the observation probabilities associated with the action, and the observations themselves. For example, an optimal policy may be to choose no action when there is a large probability of no infestation, monitoring alone with intermediate probability values and treatment alone when the probability of moderate or high infestation is large. While this framework effectively models uncertainty about the level and dynamics of the infestation, finding optimal policies is difficult and currently focuses on single-site problems.

As noted in a recent review (Meentemeyer et al., 2012), researchers are beginning to develop models of optimal disease control in a dynamic landscape. These models include spatial optimization models to allocate detection effort over space. Detection effort is shown to be related to species occurrence rates, potential damage, and surveillance costs, which may vary across the landscape. Epanchin-Niell et al. (2012) develop a dynamic model of pest colony establishment and growth and design optimal long-term equilibrium surveillance programs to minimize the total costs of surveillance and eradication. They use the model to optimize long-term surveillance effort across heterogeneous landscapes subject to region-wide surveillance budgets. Hauser and McCarthy (2009) develop a static model to optimize one-time surveillance effort across multiple sites when species' presence is uncertain prior to detection and probability of occurrence varies across sites. In contrast to the equilibrium analysis of Epanchin-Niell et al. (2012), the static model of Hauser and McCarthy (2009) is appropriate for optimizing surveillance when many local populations are thought to have established prior to the initiation of a surveillance program.

Similar to the model of Hauser and McCarthy (2009), we develop a model to optimize one-time surveillance effort across multiple sites except that we include uncertainty about the extent (rather than simply the presence) of the infestation in each site. We handle this uncertainty by splitting the management decision into two stages. In the first stage, sites are selected for surveillance given their expected levels of infestation. In the second stage, treatments are prescribed within the surveyed sites contingent on the levels of infestation found. The objective is to minimize the expected growth of the infestation subject to the total budget for surveillance and treatment. The model is a mixed-integer linear program adapted from Snyder et al. (2004) who develop a similar model to maximize the expected number of species represented in sites selected for preservation.

We apply the model to a forested landscape in Anoka County, Minnesota, where native oaks (*Quercus* spp.) are affected by the non-native oak wilt fungus, and where the eradication or suppression of the pathogen is possible only on the sites where surveillance is conducted. We use the model to inform decisions about sites to select for surveillance and removal of infected trees and provide information about cost-efficient tradeoffs between detection and control activities. Because prior information about the extent of infestation is limited, we evaluate how the level of uncertainty about infestation extent affects optimal strategies and estimates the value of gathering more information. Because oak wilt diagnosis from direct tree inspection is difficult, we investigate how reduced detection sensitivity affects the allocation of funds between surveillance and control. Finally, we use the model to evaluate simple rules of thumb to choose sites for surveillance.

2. Model Development

We develop a model of a forest management area composed of a number of distinct sites. The manager's general goal is to control an invasive pathogen in the management area, and the overarching objective is to minimize the number of newly infected trees following treatment. The choice variables are (1) a yes–no variable for each site indicating whether surveillance is undertaken in the first stage, and (2) the number of infected trees removed in each site in the second stage. We have two constraints. First, the budget constraint ensures that the total costs of surveillance and treatment do not exceed the budget level. Second, we ensure that treatment cannot occur in a given site unless the site has been surveyed; the number of trees removed in each site is bounded above by the number of trees that have been identified as infected. Broadcast treatments such as aerial spraying are not an option for forest pathogens: treatments must be applied to infected trees. Therefore, forest managers must know which trees are infected before removing them. The second constraint is introduced to reflect this requirement.

Prior to conducting detailed surveillance, the forest manager has some idea of the proportion of infected trees on each site. We consider two characterizations of their information set: either they know the exact proportion of infected trees on each site but not which specific trees are infected, or they do not know the exact proportion of infected trees but have an estimate based on a small sample of trees in each site. We first develop the model with a known proportion, and then generalize by incorporating a distribution of possible infection proportions for each site that reflects the manager's uncertain knowledge about this key piece of information.

2.1. A Forest Pathogen Growth Model

Our forest landscape is composed of distinct sites, each containing N_j host trees. A proportion, θ_j, of host trees in any given site j are infected by an invasive pathogen. The number of infected trees (I_j) at the beginning of the period can be written: $I_j = \theta_j N_j$. In the absence of management, the number of newly infected host trees, Q_j, is an increasing function of the number of infected trees in the site: $Q_j = g_j I_j$. The growth rate for infected hosts (g_j) may be site-specific because growth may depend on characteristics such as tree density or soil type that vary over the landscape. Managers can slow or stop infections by removing infected trees. Therefore, the number of newly infected trees, Q_j, on site j is a function of how many infected trees remain after the manager removes R_j trees from the site[1]:

$$Q_j = g_j\left(I_j - R_j\right). \tag{1}$$

Any amount of removal will be effective in lowering the number of newly infected trees.

It is important to note that the number of newly infected trees in each site is independent of the number of infected trees in neighboring sites. We make this assumption because many tree pathogens, including oak wilt, spread primarily via root transmission or insects that travel short distances. It is possible to extend our model to

[1] With the oak wilt pathogen, infective spore mats are produced under the bark of dead trees in the spring after trees are killed by the pathogen. Only dead trees left standing are a source of pathogen for new infections (Juzwik, 2009; Juzwik et al., 2008).

allow for the long-distance spread of tree pathogens, such as by human transport, by making the number of newly infected trees in a site dependent on the numbers of infected trees and removals in surrounding sites. We leave that extension for further work.

2.2. A Management Model with Known Proportions of Infected Trees

The management area consists of J sites with a known number of possible host trees, N_j, on each site, $j=1,\cdots,J$. The manager's objective is to minimize the number of newly infected trees in the entire management area:

$$\min\sum_{j=1}^{J}Q_j = \min\sum_{j=1}^{J}g_j\left(I_j-R_j\right). \tag{2}$$

The manager has two sets of choice variables. The first is a set of binary variables (X_j) indicating whether a site is inspected ($X_j=1$) or not ($X_j=0$). The second is a set of continuous variables (R_j), the number of infected trees removed in each site. The minimization is subject to the constraint that tree removal can only occur in sites that have been inspected, and the number of trees removed in a given site is bounded between zero and the number of infected trees that are detected:

$$0\le R_j\le I_jX_j\gamma. \tag{3}$$

The parameter γ is detection sensitivity, which is the probability that an infected tree is actually detected ($0\le\gamma\le1.0$). The value of this parameter depends on sampling technology and methods and characteristics of the pathogen and host.

The budget constraint is:

$$\sum_{j=1}^{J}c_1N_jX_j+\sum_{j=1}^{J}c_2R_j\le B, \tag{4}$$

where c_1 is the per-tree cost of surveillance in site j, c_2 is the per-tree cost of removal in site j, and B is the size of the budget. This budget constraint reflects the notion that, once a site is selected for surveillance, all trees in that site must be examined and evaluated for the presence of the disease. The assumption of comprehensive surveillance creates a fixed cost of site surveillance that is an increasing function of the total number of trees on the site.

2.3. A Management Model When the Proportions of Infected Trees are Unknown

2.3.1. Forming Beliefs about the Number of Infected Trees on Each Site

If the true proportion of infected trees in the population of trees in a given site, θ_j, is known, the number of infected trees can be calculated as the product of θ_j and the number of trees on that site, N_j. If this proportion is not known, a forest manager can estimate θ_j by inspecting a small random sample of n trees on each site and finding the sample proportion of infected trees, θ'_j, by dividing the number of infected trees, α_j, by the sample size so that $\theta'_j=\alpha_j/n$. While the sample proportion provides information about the true proportion, the true proportion remains unknown. However, the forest manager can formulate a belief about the distribution of the true proportion of infected trees (θ_j) based on the sample. This distribution translates into a belief about the distribution of the number of infected trees in each site when the parameter is multiplied by the number of trees on the site.

Bernoulli trials result in successes (1) or failures (0), according to some underlying probability. In this context, a "success" is finding an infected tree and a "failure" is finding a healthy tree. The beta distribution characterizes the distribution of the true probability of success (here, the true proportion of infected trees) based on the number of successes and failures in a sample. In particular, the probability density function of θ_j, defined over a compact interval [0,1], is characterized by two site specific parameters α_j (the number of infected trees in a sample) and β_j (the number of healthy trees in a sample):

$$f\left(\theta_j;\alpha_j,\beta_j\right)=\frac{\Gamma\left(\alpha_j+\beta_j\right)}{\Gamma\left(\alpha_j\right)\Gamma\left(\beta_j\right)}\theta_j^{\alpha_j-1}\left(1-\theta_j\right)^{\beta_j-1}. \tag{5}$$

$\Gamma(\cdot)$ is the probability density function of the gamma distribution. The expected value of θ_j is $\frac{\alpha_j}{\alpha_j+\beta_j}$, and the variance of θ_j is $\frac{\alpha_j\beta_j}{(\alpha_j+\beta_j)^2(\alpha_j+\beta_j+1)}$. In this conception of the problem as a one period model, the manager formulates a single belief about the distribution based on a single set of samples in each site. While this structure is amenable to a Bayesian updating approach in which new information can be incorporated and new beliefs formed, we confine ourselves to a single episode of belief formation.

2.3.2. Approximating Continuous Distributions with Sets of Discrete Scenarios

After formulating beliefs about the distribution of the true proportion of infected trees on each site j, a forest manager can approximate the true distributions by generating a set of scenarios of infection states in the management area. A scenario of infection in the management area is a vector $\theta=(\theta_1,\cdots,\theta_j,\cdots,\theta_J)$ of proportions of infected trees in all J sites. Each element θ_j of the vector θ is randomly drawn from [0,1] with the belief $f(\theta_j;\alpha_j,\beta_j)$. A manager randomly and independently draws S vectors to generate a set of S scenarios. Let $\theta_j(s)$ denote the sth draw of θ_j for site j and $\theta=(\theta_1(s),\cdots,\theta_j(s),\cdots,\theta_J(s))$ denote the sth scenario for all J sites. Together, the set of scenarios $\Theta=(\theta(1),\cdots,\theta(s),\cdots,\theta(s))$ reflects a range of possible infection proportions in all J sites according to the distribution characterized by site-specific parameters. We assume that the manager considers that each scenario $\theta(s)$ is equally likely: $\theta(s)$ occurs with probability $1/S$.

2.3.3. The Management Model

When the proportion of infected trees in each site is unknown, and the manager formulates a belief about the proportion based on a sample, the objective of the manager is to minimize the expected number of newly infected trees. This expected number of newly infected trees in the management area is the sum over S of the products of the realization of the random variable and their associated probabilities. In this case, since each scenario is equally likely with probability $1/S$, the management problem becomes:

$$\min\frac{1}{S}\sum_{s=1}^{S}\sum_{j=1}^{J}Q_j(s).$$

As before, removal can only occur in sites that have been inspected, and this has to be true for every scenario. Further, the number of trees removed in a given site and scenario is bounded between zero and the number of infected trees that are detected. Finally, the budget constraint must also hold for each scenario. Formally, the manager solves the following problem by selecting sites to inspect (X_j) and the number of trees to remove in each site for each scenario ($R_j(s)$) to solve the following problem:

$$\min\frac{1}{S}\sum_{s=1}^{S}\sum_{j=1}^{J}Q_j\left(R_j(\theta(s))\right) \tag{6}$$

subject to

$$0\le R_j(\theta(s))\le X_j\theta_j(s)N_j\gamma,\ \forall j=1,\cdots,J,\ \forall s=1,\cdots,S, \tag{7}$$

$$\begin{gathered}\sum_{j=1}^{J}c_1X_jN_j+\sum_{j=1}^{J}c_2R_j(s)\le B,\ \forall s=1,\cdots,S,\\ Q_j\left(R_j(\theta(s))\right)=g_j\left(\theta_j(s)N_j-R_j(\theta(s))\right),\end{gathered} \tag{8}$$

$$\forall j = 1, \cdots, J, \forall s = 1, \cdots, S, \text{and} \tag{9}$$

$$X_j \in \{0, 1\}, \forall j = 1, \cdots, J. \tag{10}$$

3. Application

We apply our model to the case of oak wilt in Anoka County, Minnesota, a 1156 km^2 county with 327,000 people (as of 2008) in the Minneapolis–Saint Paul metropolitan region (Fig. 1). Oak wilt is the most significant disease of oak trees in the eastern and central United States (Juzwik, 2009). Infection almost always causes mortality in red oaks (*Quercus* spp., Section Lobate) and less frequently in white oaks (Section Quercus). Anoka County was estimated to have 5.92 million oak trees and 885 active oak wilt pockets covering 5.47 km^2 in 2007 (Haight et al., 2011). Timely tree removal is the most common management option for controlling oak wilt infestation. While other treatment options exist, tree removal is the only management option that simultaneously addresses the need to remove dead trees, often required by municipal ordinance (Kokotovich and Zeilinger, 2011), and to prevent the infection of additional trees (Koch et al., 2010).

3.1. Defining Parameters for the Model

We chose to analyze a portion of Anoka County where oak wilt is a particularly severe management problem. As shown in Fig. 1, the area is in the northwest portion of the county, divided into 90 hexagonal 1.1 km^2 grid cells. Our first step is to calculate the number of host trees in each site. In a previous study of the economic damage of oak wilt in Anoka County (Haight et al., 2011), we estimated oak density in six primary land cover types of the Minnesota Land Cover Classification System (Minnesota Department of Natural Resources, 2009): forest, woodland, shrubland, herbaceous, cultivated vegetation, and artificial surface. For this study, we overlaid a map of the primary cover type polygons obtained from MLCCS database on the 90 sites and calculated the area of each cover type in each site. Multiplying the density of oaks in each cover type by the area in each cover type resulted in the number of trees per cover type in each site. Summing over all cover types in a site resulted in the total number of trees, N_j, for each site.

Next, we estimated the proportion of trees that are infected with oak wilt in each site using information obtained from the ReLeaf database, a statewide inventory of oak wilt pockets maintained by the Minnesota Department of Natural Resources (MN-DNR). The database includes the location and size of each pocket, the year in which the pocket was detected, whether the pocket was treated, and the types of treatments applied. The database includes 4283 pockets in Anoka County recorded from 1992 to 2007, and the boundaries of the pockets are represented by polygons in a digital map. We chose only active infection pockets reported in 2005, when the most intense monitoring was conducted. Overlaying this information on the grid cell layer, we calculated the fraction of each grid cell occupied by the infected pockets. We subsequently used this proportion to represent the sample fraction of infected trees, θ'_j, in the computation of the infection scenarios.

In areas where oak wilt is known to occur, foliar symptoms may be used to identify infected trees; however, laboratory testing may be needed to provide an accurate diagnosis (reviewed in Juzwik et al., 2011). Thus, infected trees may not be identified in the field with perfect certainty. Because the accuracy of oak wilt diagnoses from direct tree inspections has not been formally evaluated, we performed sensitivity analysis by varying the detection sensitivity parameter γ. In the base case, we assume perfect detection ($\gamma = 1.0$), and then we look at the case where 80% of the infected trees are correctly identified ($\gamma = 0.8$) to see how reduced detection sensitivity affects the allocation of funds between surveillance and control.

There are two common means of oak wilt transmission between trees: underground and overland spread (Juzwik, 2009). Root systems of related and adjacent oak trees are frequently grafted together, and the disease can be transmitted underground through the root grafts between diseased and healthy trees. This is the most common transmission mechanism and is responsible for the characteristic landscape pattern of oak wilt infection known as an infection center or "pocket," which expands in size over time. New pockets are established when beetles (family Nitidulidae) feed on the fungal mats of *C. fagacearum*, acquire spores, and move short distances to trees that have been freshly wounded. This type of transmission is known as "overland spread."

Haight et al. (2011) developed a model of oak wilt pocket expansion and new pocket formation to simulate the progression of the disease and estimate the number of infected trees over time in Anoka County. To simplify their growth model, we assume an exponential rate of infection growth within each hexagonal cell j, where the population of newly infected trees equals the population of originally infected trees less removals multiplied by the growth rate (g_j) (Eq. (1)). To find growth

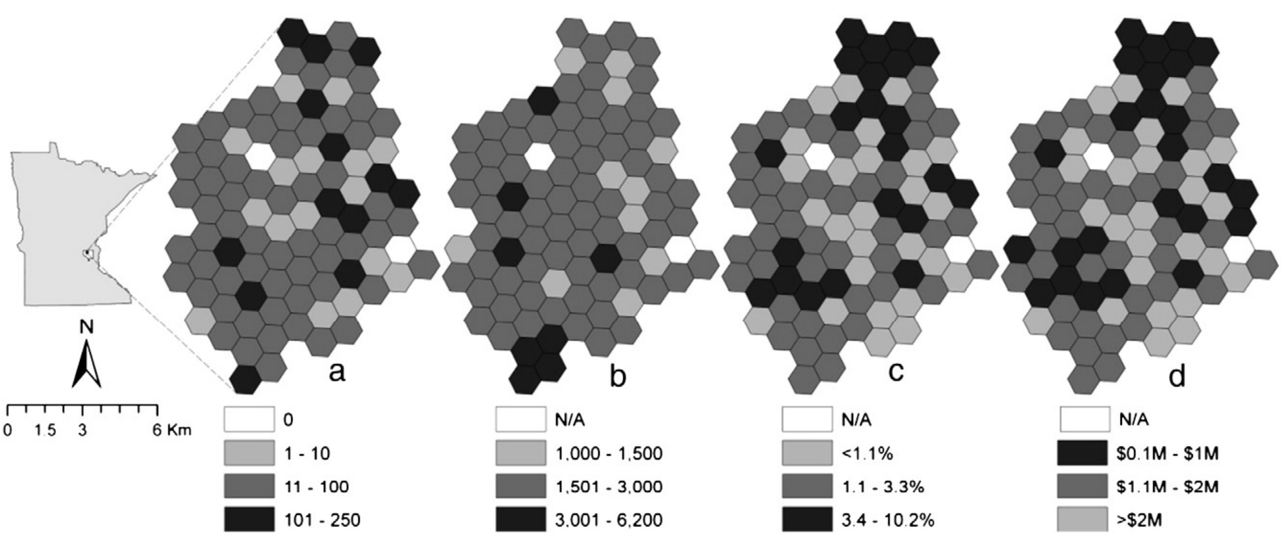

Fig. 1. Location of the study area within Anoka Co., MN, divided into 90 hexagonal 1.1 km^2 grid cells, showing: (a) number of oaks infected with oak wilt; (b) number of healthy oaks; (c) percentage of all oak trees that were infected with oak wilt; and, (d) budget level required for the optimization model to select a cell for oak wilt inspection and management. Two cells in (a) had no trees with oak wilt infection and were not considered in subsequent analyses. Cells in (d) are shaded from highest to lowest priority (dark–light).

rates consistent with the literature, we take the number of infected trees in Anoka County at the beginning of the period and the number of infected trees after a 10 year period from the oak wilt growth model in Haight et al., 2011 (Table 3a, third row, 2.42 m/yr radial growth rate). We then use the formula $[\ln(N_{10}/N_0)]/10$ to compute the implied exponential growth rate 8%. We assume a constant growth rate in our application even though our optimization model allows for site-specific growth rates because we lacked information about how growth rates might vary across sites in our study area.

We calculate the surveillance cost, c_1, from the wage rate of arborists employed by city and county governments. Arborists employed by governments tend to be ranked in GS6 or 7 in the General Schedule (GS) Locality Pay Table. Taking the average of means of salaries for GS6 and GS7, we calculate the salary for an arborist to be 17.03 dollars per hour. To diagnose whether a tree is diseased or not takes 10 min on average. Thus, surveillance cost per tree results in 2.84 dollars (17.03 dollars × 1/6). We assume that the surveillance cost per tree is uniform over all the sites. The per-tree cost of removal, c_2, usually depends on both the location and the size of a tree. Larger trees cost more to be removed. On average, trees in Anoka County are 26.6 cm diameter at breast height (dbh). We use this average value for all trees removed so that the per-tree cost of removal is assumed to be 360 dollars as a benchmark value.

Finally, we prepared a set of 2000 scenarios, $\Theta_s = (\theta_1(s), \cdots, \theta_j(s), \cdots, \theta_{90}(s))_{s=1}^{2000}$. Using a sample size ($n_j$) of 0.009* N_j and the fraction of infected trees θ'_j calculated above, we computed the parameters for the beta distribution $f(\theta_j;\alpha_j,\beta_j)$ for each site $j = 1, \cdots, 90$, where $\alpha_j = n_j\theta'_j$ and $\beta_j = n_j\left(1-\theta'_j\right)$. Using Matlab (MathWorks, 2009), we drew θ_j from an interval [0,1] following $f(\theta_j;\alpha_j,\beta_j)$ 2000 times for each site j. For each draw $s = 1, \cdots, 2000$, we obtained a vector of scenarios $\theta(s) = (\theta_1(s), \cdots, \theta_j(s), \cdots, \theta_{90}(s))$. Collecting these 2000 scenario vectors, we obtained the set of scenarios Θ. We selected 2000 scenarios because we found that solving the problem with different sets of 2000 scenarios did not affect the optimal solution or objective function value. When we solved the problem with sets containing fewer scenarios, we found some variation in the optimal solutions.

The problem specified in Eqs. (6)–(10) was solved by using the integrated solution package GAMS (GAMS Development Corporation, 1990), which is designed for large and complex linear and mixed integer programming problems. The termination criterion for the optimization runs is a combination of time limit and optimality: the solver is instructed to stop and report the solution after 4 h of runtime or when the relative gap is less than 0.005, whichever happens first. All of the solutions had relative gaps less than 0.005 in less than 4 h of runtime on a Lenovo T60 laptop computer with an Intel Core 2 central processing unit.

3.2. Results

3.2.1. Distribution of Oaks and Oak Wilt Infection

The study area contains 194,798 healthy oaks (20 trees ha^{-1}) with most cells having 13–27 trees ha^{-1} (Fig. 1b). There are 4853 infected oaks (0.5 infected trees ha^{-1}) in the study area with most cells having fewer than 100 infected trees (<1 infected tree ha^{-1}) (Fig. 1a). The proportion of infected oaks per cell varies from 0 to 10%. Cells with the highest proportions of infected oaks are located in on the northern, western and eastern edges of the study area (Fig. 1c).

3.2.2. Total Costs and Surveillance Costs

We solved the optimization model for 25 different budget levels from \$100,000 to \$2,500,000 in \$100,000 increments and mapped the locations of cells that are selected for inspection under the different budgets (Fig. 1b). We found that once a cell is selected for inspection, it is also selected at all higher budget levels. Cells selected at the lower budget levels (<\$1,000,000) are concentrated in the northern, western, and eastern edges of the study area and correspond fairly well to cells with the highest proportions of infected trees (Fig. 1c).

We computed the expected percentage of healthy trees saved from infection for each budget level, and show these results in a cost curve (Fig. 2). We find that the slope of the cost curve increases as the fraction of healthy trees saved from infection increases. Note that the surveillance cost curve flattens several times. This is because saving more healthy trees from infection does not necessarily mean surveying more sites. It is optimal for a forest manager not to survey more sites until she removes all the infected trees on sites already surveyed. We find that the marginal cost of saving trees from infection rises with more trees saved. The intuition for this result is that the lower priority cells have many trees to inspect and few infected trees to find. The cost of preventing a small number of infections is high because of the high surveillance cost in those cells. Optimized surveillance costs comprise between 13% and 22% of the total cost of surveillance and treatment. Note also that the cost of saving trees from infection is quite high. For instance, it would cost \$1 million to save approximately 50% of potentially infected trees in one 99 km^2 area in one county.

To find what might be gained by detecting the pathogen when it is treated at a much earlier stage, we performed the same analysis assuming that the fraction of infected trees was much smaller: 1/100 of its size in 2005. Results of this analysis are shown in Fig. 3. As expected, it is much cheaper to control any given expected percentage of infected trees. However, the fraction of the budget spent on surveillance is much higher – in the 70% range – because there are many more healthy trees to inspect per diseased tree detected.

3.2.3. Rules of Thumb for Selecting Sites for Surveillance

Without recourse to an optimization model, how might a forest manager achieve a high degree of success in limiting the spread of a pathogen? Possible rules of thumb might include: (1) choosing sites with the highest expected number of infected trees; (2) choosing sites with the highest expected proportion of infected trees; and (3) choosing sites with the highest expected number of healthy trees to be saved. The first rule is plausible because the pathogen grows exponentially, and surveillance of sites with large numbers of infected trees would allow for removal of a large source of infection. The second rule is plausible because inspection of a site could identify a large number of infected trees relative to the number of trees inspected. The third rule may be reasonable as a way to protect many healthy trees. Our approach to learning how successful these rules might be was to rank sites according to these three plausible criteria, as well as random selection of sites, and specify the sites to be inspected with limited budgets. After constraining the model so that these sites – and only these sites – were inspected, we ran the optimization model and found the expected percentage of potentially infected trees saved from infection for four budget levels. Results, with the original fraction of infected trees, are shown in Table 1.

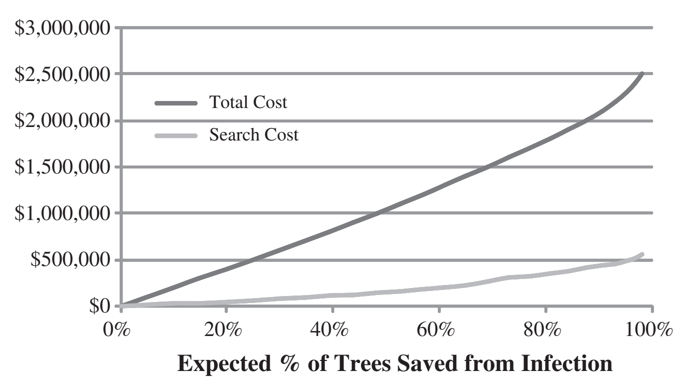

Fig. 2. Cost curve.

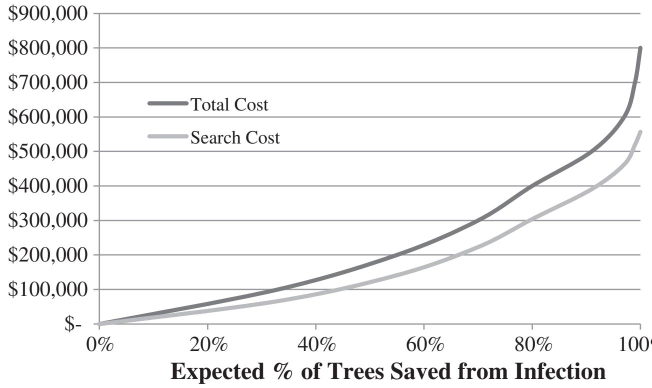

Fig. 3. Cost curve with 1% of the infections in 2005.

Table 1 shows that the rules of thumb perform well relative to the optimization, but that selecting sites at random is a poor strategy. The second rule – selecting sites based on the expected proportion of trees infected – matches the optimization fairly well. However, as Table 1 shows, unless all the sites are inspected, there is still a penalty in terms of a loss in performance if a rule of thumb is used instead of the optimization. For example, with a $1,000,000 expenditure, 48.29% of trees are saved using the optimization, while 46.48% and 47.12% are saved using the rules of thumb. Only 40.36% are saved using random surveillance. The worst strategy is to choose sites based on the expected number of healthy trees. This rule of thumb performs worse than random selection, with 39.01% of trees saved. Fig. 1 illustrates these results in a map of the study area, showing that ranking sites by the percent infected (Fig. 1c) is the closest to the optimal selection of sites for surveillance (Fig. 1d).

Interestingly, optimal surveillance often leads to higher surveillance costs than the first two rules of thumb for achieving a given targeted expected percentage of trees saved. Surveillance may be expensive, but when sites are optimally chosen, surveillance is worth the cost because it leads to more effective removal. However, spending resources on surveillance may be wasteful in comparison with the optimal solution: both the third rule of thumb and random site surveillance lead to higher than optimal surveillance costs. Selecting sites with high numbers of healthy trees leads to inspecting sites with high surveillance costs and relatively few infected trees; the number of infected trees found per dollar spent on surveillance was low with this strategy.

The rules of thumb assume, of course, that the required information is already available. We have not attempted to incorporate the costs of acquiring this information. However, to make the rules worthwhile, the marginal costs to gather this information would need to be less than or equal to the marginal benefits of applying the rule of thumb rather than selecting sites randomly.

3.2.4. Increased Sampling

These results are based on the assumption that 0.9% of trees are sampled before the surveillance-treatment phase in order to formulate beliefs about the expected proportion of trees infected. To assess whether increased pre-optimization sampling would improve performance, we repeated our analysis assuming that 9% of trees are sampled, an assumption equivalent to the original estimates of α_j and β_i being 10-fold greater. This assumption does not affect the expected value of θ_j but reduces the variance by approximately 10%. Results are reported in Table 2.

More precise estimates of infection do improve the performance of the optimization model: the expected number of trees saved increases with the increase in sampling. In general, the more precise estimate of θ_j improves the odds of selecting cells that truly have a greater proportion of infected trees for surveillance. This improvement, in turn, lowers the search cost to find infected trees, leaving more of the budget for tree removal. However, at low budget levels, an increase in the amount of sampling has little effect on the expected number of trees saved. At high budget levels, additional information is more beneficial. For example, at a $1.5 million budget level, increasing sampling costing $46,000 (sampling an additional 16,197 trees) yields an additional 11.37 more trees saved. This is a wiser way to spend additional funds than spending it on surveillance and control in the optimization phase, since spending $46,000 there yields 7 additional trees saved (calculated by optimizing the model with a budget of $1,546,000). These calculations assume that the per-tree cost of identifying infected trees is the same in the sampling phase and in the surveillance phase ($2.84/tree).

3.2.5. Reduced Detection Sensitivity

The results in the base case assume that infected trees are detected with certainty ($\gamma = 1.0$). When the detection sensitivity parameter is reduced to 0.8, the performance of the optimal solutions is reduced: the expected number of trees saved decreases with the decrease in detection sensitivity (Table 3). In general, lower detection sensitivity reduces the number of infected trees that are identified and removed in the sites selected for surveillance. In response, more funds are allocated to surveillance and the number of sites surveyed increases. For example, with a budget of $1 million, $180,269 is used to survey 30 sites when $\gamma = 0.8$ compared with $146,755 to

Table 1
Rule of thumb results. Differences from the optimized results are shown in parentheses.

Budget	Optimized	Rule: expected # of infected trees	Rule: expected proportion of infected trees	Rule: expected # of healthy trees	Random selection of sites
Expected % of trees saved					
$ 500,000	24.9%	23.7% (−1.2%)	24.1% (−0.8%)	18.0% (−6.9%)	21.3% (−3.6%)
$1,000,000	48.3%	46.5% (−1.8%)	47.1% (−1.2%)	39.0% (−9.3%)	40.4% (7.9%)
$1,500,000	69.2%	68.0% (−1.2%)	68.5% (−0.6%)	59.0% (−10.2%)	62.5% (6.7%)
$2,000,000	87.8%	87.4% (−0.4%)	87.5% (−0.3%)	78.5% (9.3%)	82.8% (5.0%)
Search cost					
$ 500,000	$ 61,223	$52,259 (−$8964)	$50,016 (−$11,207)	$144,679 ($83,456)	$90,184 ($28,962)
$1,000,000	$146,755	$130,741 (−$16,014)	$121,187 (−$25,565)	$260,649 ($113,894)	$238,302 ($91,547)
$1,500,000	$269,261	$244,514 (−$24,747)	$241,355 (−$27,905)	$403,889 ($134,628)	$334,209 ($64,949)
$2,000,000	$418,792	$394,447 (−$24,345)	$394,972 (−$23,819)	$456,295 ($37,504)	$480,324 ($61,533)

Table 2
Results with increased pre-optimization sampling. Differences between high and low sampling are shown in parentheses.

Budget	Low Sampling			High Sampling		
	Expected % of saved trees	Expected # of saved trees	Search cost	Expected % of saved trees	Expected # of saved trees	Search cost
$500,000	24.9%	95.1	$61,223	25.5% (+0.6%)	98.7 (+3.6)	$ 55,597 (−$5626)
$1,000,000	48.3%	184.7	$146,755	49.7% (+1.4%)	192.7 (+8.0)	$126,447 (−$20,308)
$1,500,000	69.2%	265.1	$269,261	71.3% (+2.1%)	276.5 (+11.4)	$245,140 (−$24,121)
$2,000,000	87.8%	337.1	$418,792	90.6% (+2.8%)	351.3 (+14.9)	$394,972 (−$23,819)

Table 3
Results with decreased disease detection sensitivity. Differences between low and high detection sensitivity are shown in parentheses.

Budget	High detection sensitivity ($\gamma=1.0$)			Low detection sensitivity ($\gamma=0.8$)		
	Expected % of saved trees	Expected # of saved trees	Search cost	Expected % of saved trees	Expected # of saved trees	Search cost
$500,000	24.9%	95.1	$61,223	24.0% (−0.9%)	91.8 (−3.3)	$ 77,450 (+$16,227)
$1,000,000	48.3%	184.7	$146,755	45.5% (−2.8%)	174.7 (−10.0)	$180,269 (+$33,514)
$1,500,000	69.2%	265.1	$269,261	64.2% (−5.0%)	246.7 (−18.4)	$341,878 (+$72,617)
$2,000,000	87.8%	337.1	$418,792	77.2% (−10.6%)	299.4 (−37.7)	$506,523 (+$87,731)

survey 25 sites selected when $\gamma=1.0$. We also note that reducing the detection sensitivity uniformly across sites does not affect the order in which the sites are selected for surveillance. For example, the 30 sites selected for surveillance with a $1 million budget when $\gamma=0.8$ include the 25 sites selected for surveillance when $\gamma=1.0$.

3.2.6. Infected Tree Removal

It is optimal for a forest manager to exhaust her budget left after surveillance by cutting down as many infected trees as possible with the remaining funds. This approach minimizes the number of new infections for any *s*th scenario. Also, a forest manager would be indifferent about where to cut infected trees. This is simply because we assume that the infection growth rate is uniform over the sites. If growth rates varied over the landscape, the expected number of healthy trees saved from infection would also vary according to where tree removal occurs. Then, in the tree removal step, a forest manager would prioritize sites with high growth rates.

4. Discussion

An effective strategy for limiting the spread of invasive forest pathogens such as oak wilt is to find and remove diseased trees. The detection step is important because diseased trees are often difficult to identify and must be specifically identified before being removed. The problem is further complicated by spatial heterogeneity of host trees, infection levels, and costs across the landscape. There has been scant attention paid to spatial optimization of surveillance and control of invasive species. We present a static, spatial optimization model of surveillance and control where the number of infected trees is uncertain and the number of susceptible trees, the expected number of infected trees, the infection growth rate, and the cost of tree removal vary across sites. Our model is a mixed-integer, linear program, inspired by the site selection literature, designed to choose locations in a grid on which to focus surveillance and control in a setting with budget-constraints and prior information on the expected number of diseased trees. The model offers practical guidance to managers in charge of deciding how and where to spend limited public dollars when the goal is to reduce the spread of a forest disease. Further, the model can be used to construct a curve that reflects the cost of protecting healthy trees from infection. This curve provides the manager with information about the level of budgets required to achieve targeted percentages of healthy trees saved from infection. The marginal cost of saving an additional healthy tree from the pathogen increases as the targeted protection level is set higher.

In our application to oak wilt management in Anoka County, Minnesota, we find that the cost of protecting healthy trees is substantial. For example, assuming an annual infection growth rate of 8%, a $1 million budget for surveillance and control would save an expected 185 trees from infection. This is, on average, $5400 per tree. Changing the assumed growth rate from 0.08 to 0.12 has no effect on the optimal surveillance strategy (which cells to survey), the surveillance cost, or the expected percentage of trees saved from infection. It does have an effect on the expected number of trees saved from infection, however. Because more trees are potentially infected with a higher growth rate, removing the same number of infected trees saves a higher number. For a $1 million budget, assuming a growth rate of 12% increases the expected number of trees saved to 277 for a lower average cost of $3610. These high costs from an annual surveillance and removal program suggest that it may be preferable to wait longer between surveillance and removal efforts to reduce overall costs. On the other hand, as Fig. 3 shows, treating a landscape with few infections leads to low costs of reducing the expected percentage of trees saved, even when the average cost per tree saved is very high. Determining the optimal timing of surveillance and removal activities would require a dynamic model with multiple periods, which is beyond the scope of our study.

We focus on the cost side of the ledger and thus hesitate to conclude whether a particular level of expenditure for surveillance and control is justifiable from a benefit–cost point of view. Nevertheless, results from hedonic property value studies on the value of urban trees give perspective to our estimate of the average cost per tree saved from infection. While most estimates of tree value are lower, a recent study of street trees in Portland, Oregon found a value of $8870 per tree (Donovan and Butry, 2010). Perhaps more relevant is a study by Holmes et al. (2010) who examined the property-value impacts of tree mortality caused by a non-native forest insect (hemlock woolly adelgid). They found that when the cost to neighboring parcels was considered, the cost was about three times as high as when the cost to only a single parcel was estimated. This suggests that the economic cost of an infectious tree can be much higher than just the cost of losing a healthy tree. Expenditures to control oak wilt may be worthwhile if the goal is to preserve a healthy tree canopy.

The primary innovation of our model is its handling of uncertainty in the extent of infestation. Uncertainty is managed by splitting the decision of surveillance and control into two stages. In the first stage, sites are selected for surveillance given their expected levels of infestation. In the second stage, treatments are prescribed within the surveyed sites contingent on the levels of infestation found. This allows us to evaluate the effect of the level of uncertainty on optimal strategies and estimate the value of gathering better prior information. We find that forming more accurate estimates of the proportion of infected trees through increased sampling reduces the cost of surveillance and removal. However, the cost of increased sampling may outweigh the benefits of reduced surveillance and removal costs at low budget levels. Increased sampling takes place over the entire landscape, and the benefit of additional precision in the estimated amount of infection has only small effects on sites chosen. The optimal sampling rate may be an additional outcome of a more complex dynamic model.

We also explored potential rules of thumb that could be used in the absence of an optimization model for applications like this one. Two rules – choosing sites with high expected proportions of infected trees and choosing sites with high expected numbers of infected trees – performed well relative to the optimum. These rules would tend to find cells with a high number of infected trees found per dollar spent on surveillance. Choosing sites with high expected numbers of healthy trees would not be a sound strategy, as this rule yields fewer saved trees than random selection of sites. It is important to recognize, however, that these rules of thumb were evaluated in a setting where the growth rate was assumed constant across the landscape. Other general rules of thumb have been used in oak wilt management programs. In Texas, for example, sites with rapid oak wilt expansion are given higher priority as are sites with high prospects for treatment success (Davies, 1992). Our sense is that rules of thumb may be an implicit part of forest management, and this research suggests that using rules of thumb that are evaluated explicitly with an optimization model may be a helpful way to conduct forest pest management.

From a computational perspective, the two-stage surveillance and control model we develop is tractable: we solved problems with 90 sites and 2000 scenarios with commercial mixed integer programming software in less than 4 h on a laptop computer. While solving models with more sites will likely require additional execution time, it is promising that rules of thumb can provide near-optimal solutions to problems that exceed computer resource limits.

While our optimization model addresses a one-time investment in surveillance and treatment, the performance of the rules-of-thumb for selecting sites for surveillance can be tested over a longer time horizon using a more detailed model of oak wilt establishment and spread. The economic efficiency of these surveillance and treatment strategies can then be compared with sanitation strategies that were developed to slow the spread of Dutch elm disease in U.S. cities in the 1980s (Baughman, 1985). It may also be possible to use a spatial-dynamic optimization model based on partially observable Markov decision processes to develop surveillance and control policies when there are uncertainties in the incidence of infection, infection dynamics, and detection (e.g., Chadès et al., 2011). The performance of these policies can then be compared with rules-of-thumb for selecting sites for surveillance and control obtained from static, spatial optimization models.

Acknowledgments

This article was partially produced under a co-operative agreement (number 58-7000-6-0081) with the Economic Research Service's (USDA) PREISM economics of invasive species management program. This research has also been supported by the U.S. Forest Service, Northern Research Station and the Minnesota Agricultural Experiment Station. The authors thank Susan Burks of the Minnesota Department of Natural Resources for providing information on the distribution of oak wilt, D. Bengston, K. Kovacs, and S. Snyder for valuable pre-submission reviews, and three anonymous referees for valuable comments.

References

Baughman, M.J., 1985. Economics of Dutch elm disease control: a model and case study. Journal of Forestry 83 (9), 554–557.
Bogich, T.L., Liebhold, A.M., Shea, K., 2008. To sample or eradicate? A cost minimization model for monitoring and managing an invasive species. Journal of Applied Ecology 45, 1134–1142.
Chadès, I., Martin, T.G., Nicol, S., Burgman, M.A., Possingham, H.P., Buckley, Y.M., 2011. General rules for managing and surveying networks of pests, diseases, and endangered species. Proceedings of the National Academy of Sciences of the United States of America 108 (20), 8323–8328.
Davies, C.S., 1992. Environmental management of oak wilt disease in central Texas. Environmental Management 16 (3), 323–333.
Donovan, G.H., Butry, D.T., 2010. Trees in the city: valuing street trees in Portland, Oregon. Landscape Urban Planning 94 (2), 77–83.
Epanchin-Niell, R.S., Hastings, A., 2010. Controlling established invaders: integrating economics and spread dynamics to determine optimal management. Ecology Letters 13, 528–541.
Epanchin-Niell, R.S., Haight, R.G., Berec, L., Kean, J.M., Liebhold, A.M., 2012. Optimal surveillance and eradication of invasive species in heterogeneous landscapes. Ecology Letters. http://dx.doi.org/10.1111/j.1461-0248.2012.01800.x (on line).
GAMS Development Corporation, 1990. The general algebraic modeling system (GAMS). Retrieved from http://www.gams.com.
Haight, R.G., Polasky, S., 2010. Optimal control of an invasive species with imperfect information about the level of infestation. Resource and Energy Economics 32 (4), 519–533.
Haight, R.G., Homans, F.R., Horie, T., Mehta, S.V., Smith, D.J., Venette, R.C., 2011. Assessing the cost of an invasive forest pathogen: a case study with oak wilt. Environmental Management 47 (3), 506–517.
Hansen, E.M., 2008. Alien forest pathogens: phytophthora species are changing world forests. Boreal Environment Research 13, 33–41.
Harwood, T.D., Tomlinson, I., Potter, C.A., Knight, J.D., 2011. Dutch elm disease revisited: past, present and future management in Great Britain. Plant Pathology 60, 545–555.
Hauser, C., McCarthy, M., 2009. Streamlining 'search and destroy': cost-effective surveillance for invasive species management. Ecology Letters 12 (7), 683–692.
Holmes, T.P., Murphy, E.A., Bell, K.P., Royle, D.D., 2010. Property value impacts of hemlock woolly adelgid in residential forests. Forest Science 56 (6), 529–540.
Homans, F., Horie, T., 2011. Optimal detection strategies for an established invasive pest. Ecological Economics 70, 1129–1138.
Juzwik, J., 2009. Epidemiology and occurrence of oak wilt in Midwestern, Middle, and South Atlantic States. In: Billings, R., Appel, D. (Eds.), Proceedings of the National Oak Wilt Symposium, pp. 55–66 (Austin, TX).
Juzwik, J., Harrington, T., MacDonald, W., Appel, D., 2008. The origin of *Ceratocystis fagacearum*, the oak wilt fungus. Annual Review of Phytopathology 46, 13–26.
Juzwik, J., Appel, D.N., MacDonald, W.L., Burks, S., 2011. Challenges and successes in managing oak wilt in the United States. Plant Disease 95 (8), 888–900.
Koch, K.A., Quiram, G.L., Venette, R.C., 2010. A review of oak wilt management: a summary of treatment options and their efficacy. Urban Forestry & Urban Greening 9 (1), 1–8.
Kokotovich, A.E., Zeilinger, A.R., 2011. Exploring social and institutional variation across oak wilt risk management programs in Minnesota. Urban Forestry & Urban Greening 10 (1), 39–45.
Kovacs, K.F., Mercader, R.J., Haight, R.G., McCullough, D.G., Siegert, N.W., Liebhold, A.M., 2011. The influence of satellite populations of emerald ash borer on projected economic damage in U.S. Communities, 2010–2020. Journal of Environmental Management 92, 2170–2181.
Loo, J.A., 2009. Ecological impacts of non-indigenous invasive fungi as forest pathogens. Biological Invasions 11 (1), 81–96.
MathWorks, 2009. Matlab, Version 7.
Meentemeyer, R.K., Anacker, B., Mark, W., Rizzo, D.M., 2008. Early detection of emerging forest disease using dispersal estimation and ecological niche modeling. Ecological Applications 18, 377–390.
Meentemeyer, R.K., Haas, S.E., Vaclavik, T., 2012. Landscape epidemiology of emerging infectious diseases in natural and human-altered ecosystems. Annual Review of Phytopathology 50, 379–402.
Mehta, S.V., Haight, R.G., Homans, F.R., Polasky, S., Venette, R.C., 2007. Optimal detection and control strategies for invasive species management. Ecological Economics 61, 237–245.
Minnesota Department of Natural Resources, 2009. Minnesota Land Cover Classification System. Retrieved from http://www.dnr.state.mn.us/mlccs/index.html.
Ndeffo Mbah, M.L., Gilligan, C.A., 2010. Balancing detection and eradication for control of epidemics: sudden oak death in mixed-species stands. PLoS One 5 (9), e12317.
Økland, B., Skarpaas, O., Schroeder, M., Magnusson, C., Lindelöw, Å., Thunes, K., 2010. Is eradication of the pinewood nematode (*Bursaphelenchus xylophilus*) likely? An evaluation of current contingency plans. Risk Analysis 30 (9), 1424–1439.
Snyder, S.A., Haight, R.G., ReVelle, C.S., 2004. A scenario optimization model for dynamic reserve site selection. Environmental Modeling and Assessment 9, 179–187.