UC Davis
UC Davis Previously Published Works

Title

Learning about a Moving Target in Resource Management: Optimal Bayesian Disease Control

Permalink

https://escholarship.org/uc/item/83d8c4ck

Journal

American Journal of Agricultural Economics, 99(1)

Authors

MacLachlan, Matthew J
Springborn, Michael R
Fackler, Paul L

Publication Date

2017

Peer reviewed

Learning about a moving target in resource management: Optimal Bayesian disease control

Corresponding Author: Matthew J. MacLachlan[1]
(matthew.maclachlan@ers.usda.gov)
Economic Research Service,
United States Department of Agriculture,
355 E Street SW
Washington, DC 20024
Phone: (510) 847-9297

Michael R. Springborn (mspringborn@ucdavis.edu)
Department of Environmental Science & Policy,
University of California,
2104 Wickson Hall, One Shields Avenue,
Davis, CA 95616

Paul L. Fackler (pfackler@ncsu.edu)
Department of Agricultural & Resource Economics,
North Carolina State University,
4344 Nelson Hall, 2801 Founders Drive,
Raleigh, NC 27607

Resource managers must often make difficult choices in the face of imperfectly observed and dynamically changing systems (e.g. livestock, fisheries, water and invasive species). A rich set of techniques exists for identifying optimal choices when that uncertainty is assumed to be understood and irreducible. Standard optimization approaches however cannot address situations in which reducible uncertainty applies to either system behavior or environmental states. The adaptive management literature overcomes this limitation with tools for optimal learning, but has been limited to highly simplified models with state and action spaces that are discrete and small. We overcome this problem by using a recently developed extension of the Partially Observable Markov Decision Process (POMDP) framework to allow for learning about a continuous state. We illustrate this methodology by exploring optimal control of bovine tuberculosis in New Zealand's cattle. Disease testing—the control variable—serves to identify herds for treatment and provides information on prevalence, which is both imperfectly observed and subject to change due to controllable and uncontrollable factors. We find substantial efficiency losses from both ignoring learning (standard stochastic optimization) and from simplifying system dynamics (to facilitate a typical, simple learning model), though the latter effect dominates in our setting. We also find that under an adaptive management approach, simplifying dynamics can lead to a belief trap in which information gathering ceases, beliefs become increasingly inaccurate and losses abound.

Uncertainty and change are two hallmarks of natural resource management. The dual challenges of imperfect knowledge and a dynamically changing system are present in a wide range of settings, including fisheries, forests, water and livestock disease control to name a few. Where uncertainty is reducible, it is also natural to incorporate learning into management, an approach that has been called adaptive management (Walters, 1986, Walters, 1974, Walters and Hilborn, 1976). Although the term adaptive management is sometimes used only to apply to learning about system behavior, we adopt a more inclusive use of the term that applies to learning about both the behavior and the state of a dynamic system (i.e. encompassing both structural and observational uncertainty). We use adaptive management broadly in this paper to include any situation in which management choices are updated when new information becomes available that changes beliefs concerning imperfectly observed features of a system. Examples of adaptive management applied to uncertain system dynamics include management with learning about a water pollution threshold (Bond and Loomis, 2009), the persistence likelihood of a particular species (Chades, et al., 2012), the invasive species infestation rates of agricultural imports from different exporters (Springborn, 2014), fish survivorship (Springborn and Sanchirico, 2013), the population dynamics of a harvested population (Johnson, 2011, U.S. Fish and Wildlife Service, 2013) and the parameters governing a generic payoff process in an economic environment (Wieland, 2000). Each of these examples involves learning about an unobserved model component that is time invariant (e.g. a fixed biological or economic parameter).

Until recently however, this approach was, for the most part, limited to situations with a small number of discrete state levels. A recently developed approach (Zhou, et al., 2010) extends the POMDP approach to continuous variables using an approximation method known as "density projection" for achieving model tractability. The method was applied to a situation in which a fixed model parameter was uncertain in Springborn and Sanchirico (2013); here it is applied to a situation in which an uncertain state variable is changing over time.

2

This continuous state POMDP approach is well suited to problems of control of infectious disease or invasive species. Typically the true prevalence of an infection is unknown, making planning difficult. This is especially true when testing for the presence of disease is costly and/or subject to implementation constraints due to limited personnel and equipment. Previous literature has identified a number of ways to determine optimal strategies or a good rule of thumb for identifying affected individuals or units. For example in the context of bTB (Gramig and Horan, 2011) and sudden oak death (Mbah and Gilligan, 2010), authors have assumed that managers know the true level of prevalence with certainty and the purpose of testing is solely to identify which particular units are infected and should be treated. Similarly, in Filipe, et al. (2012) and Atallah, et al. (2014) the purpose of testing is to identify which management units (e.g. spatial cell or individual grape vine) are infected for treatment. Although these two studies do not assume that true prevalence is known, in both cases the testing strategy does not depend on the overall prevalence in the system.[2] Testing serves a dual role however; in addition to identifying which units are infected it also provides information about the extent of the infection. In this paper, we account for the value of information in both identifying individuals to treat *and* in honing beliefs about uncertain prevalence for optimal endogenous learning.

The problem of using monitoring to better understand a dynamically changing system state, specifically the extent of an invasive infestation, was examined in Haight and Polasky (2010) (and extended in Fackler and Haight (2014)). The authors use a POMDP approach but the set of possible states was discrete and limited to three levels. Other similarly constrained examples include Regan, et al. (2006) and Regan, et al. (2011). To date most studies that apply adaptive management to problems involving either structural or observational uncertainty have assumed that the underlying uncertain variables are discrete. When the underlying source of uncertainty involves a continuous variable difficulties arise if a conjugate family or belief distributions is not available. Recently a projection approach has been developed to overcome this difficulty. It was applied first to wildlife management in Moore (2008) and put on a rigorous footing in Zhou, et al. (2010). The projection approach was applied to a structural uncertainty problem in Springborn and Sanchirico (2013) and is currently being developed in another study involving observational uncertainty in fisheries management by Kling, et al. (2016). This optimal learning approach has not yet been applied to problems of infectious disease monitoring and control, and more specifically given uncertainty over prevalence.

In this paper we use the continuous POMDP approach to determine an optimal strategy for testing cattle herds that are potentially infected with bovine tuberculosis (bTB). bTB is an infectious and potentially fatal disease of both animals and humans that persists throughout much of the world, including the United Kingdom, Japan, Mexico, much of Central and South America, as well as much of Asia and Africa (Cousins, 2005). Our numerical application is based on bTB control efforts in New Zealand, where the disease has been rigorously studied over its long history.

The disease has remained endemic in New Zealand's cattle herds since 1893 (de Lisle, 1993). Controlling spread and permanent eradication is encumbered by abundant disease vectors (e.g. the brushtail possum and ferrets) and subtle clinical signs that cause infections to remain subclinical (unobservable) to both herd owners and managers until the very late stages of decline, resulting in substantial uncertainty regarding true prevalence across herds (Morrison, et al., 2000). In New Zealand, the central animal health authority[3] applies tuberculin skin tests to identify latently infected

---

[2] In Atallah, et al. (2014), the level of testing is responsive to prevalence, in part. Diseased grape vines in a vineyard are divided between those that are costlessly observed (infective), those that can be observed only with testing (exposed-detectable) and those that cannot be observed (exposed-undetectable). The testing choice is the number and configuration of vines neighboring an infective vine that is to be tested. This choice is not responsive to overall prevalence, however the number of tests overall is proportional to the number of infective vines, and thus to the subset of the overall prevalence that is costlessly observable.

[3] The Animal Health Board (AHB) delegates administrative duties to other public and private entities.

herds. Once identified, infected herds are placed under quarantine and those animals within the herd that test positive are selectively culled. Test results provide value both in identifying infected herds (through avoiding further infections and speeding recovery (without immunity)) and in providing information about prevalence, which can be used to enhance future management decisions.

Disease control problems typically involve both substantial uncertainty and clear opportunities for learning about disease prevalence. For example, Shea, et al. (2014) examine *passive* adaptive management of foot-and-mouth disease and provision of measles vaccinations. Under their passive learning approach information arrival is assumed to be independent of the management decision. Learning is highly stylized— it occurs exogenously in one time step, is perfect (the manager learns the truth with certainty) and costless. The decision problem is also highly simplified—control is chosen in just two stages with the first control decision made under model uncertainty and the second with perfect information. In contrast, in our model learning does not perfectly resolve uncertainty, occurs repeatedly over time, and is determined endogenously based on an explicit balancing of costs.

Our results show that both incorporating learning and accounting for the nature of the moving target (transmission) have implications for efficiency and interact in interesting ways to determine optimal policy. While both provide improvements in outcomes, the gains from considering transmission exceed those from learning. We find that a manager who ignores learning can hold more accurate beliefs about true disease prevalence than a manager who learns but ignores transmission. We also find that focusing on learning at the expense of incorporating transmission leaves the manager at risk of falling into a belief trap where overly optimistic beliefs go uncorrected and prevalence takes off. We discuss in the conclusions how these results might shift if perturbations to the unobserved underlying state are less unidirectional than occurs with disease transmission. Overall, the results show the potential pitfalls of ignoring one or both of these components, especially the moving target element.

We begin below by specifying a model of disease prevalence, spread, testing and control. Then we describe in detail how belief dynamics are made to incorporate both learning from testing and movement in the true underlying unobserved state. Finally we specify the economic decision problem in a dynamic programming framework and present results from the numerical application.

Methods

We adapt the approach of Gramig and Horan (2011) for modeling transmission and control of bTB in cattle populations within a particular region to account for the uncertainty regarding prevalence that may be resolved through testing. The modeling approach makes use of a metapopulation framework, which describes the number of facilities within each health and economic state or "compartment" and the system-wide dynamics governing movement between the compartments. Following Gramig and Horan, between-herd transmission occurs via animal movements within the region. We extend the transmission model to explicitly include infection from non-farm disease vectors (e.g. possums and ferrets). The unit of observation is the herd and the focus is on between-herd (rather than within-heard) disease dynamics.

Gramig and Horan use testing to identify infected herds for treatment but ignore the value of testing information for understanding overall prevalence, which they assume is known. Within our modeling framework, test results are used to identify herds for treatment *and* provide information regarding the unknown prevalence. We model the management problem as a Markov decision process (MDP) in which the manager chooses a level of testing conditional on the current state. We introduce Bayesian learning to model a common feature of disease control problems: the manager is uncertain regarding the true prevalence. The manager may use test results to (partially) resolve this uncertainty.

When system state variables are not perfectly observed, the optimization problems can be cast

in a POMDP framework (Fackler and Haight, 2014). The earliest example of applying a POMDP approach to such optimization problems was completed by Lane (1989), who models learning about fish stock levels in different patches through harvest levels. In previous models that use this approach simplifying assumptions are imposed for tractability. In particular states are assumed to fall into a small number of discrete levels (e.g. 2-3) and the belief distribution is taken to be a discrete probability distribution over these levels.

For many problems using a small number of discrete states does not provide enough flexibility for good management. If the uncertain variable is treated as continuous one can instead represent the belief distribution using a convenient family of density functions. The original but uncertain variable can then be replaced by the parameters of that density function (e.g. the mean and variance of a normal distribution). One difficulty that arises in using this approach is that for many problems the updated (posterior) belief distribution does not have a convenient functional form, i.e. the prior belief and the system dynamics do not represent a conjugate pair (Schlaifer and Raiffa, 1961). Later in this section we describe how we build on a recently developed approach (Zhou, et al., 2010) that extends POMDP methods to allow for continuous variables and addresses the issue of non-conjugacy using projection methods.

At the core of the decision problem, the manager chooses the level of testing to apply within a single period given beliefs regarding prevalence, an understanding of the underlying physical dynamics and expectations of how information may be used to improve the efficiency of future testing choices. The time-span of a single period is defined as one month. We assume the following order of events as depicted on the left side of figure 1: at the beginning of the period, the manager selects the number of herds to test and observes the results; movement controls are imposed on those herds found to be infected; new infections occur and some fraction of facilities recover from movement control status; and then payoffs accrue when the period ends.

In figure 1, we also illustrate the combined dynamics for Bayesian belief updating and changes in herd health or management status. At the beginning of each period, a set of N herds is divided into three subsets: S herds are susceptible, I herds are latently infected and M herds have been identified as infected and are already under movement controls. Both N and M are observed and known to the manager. Gramig and Horan (2011) assume that the division of remaining herds between S and I is observed, i.e. the proportion of uncontrolled herds that are latently infected, $p = I/(S+I)$, is known with certainty. Under this assumption, the value of testing is limited to the identification of infected herds. We relax this assumption to capture the notion that the true level of disease prevalence is typically not known with certainty such that testing also serves as a tool for monitoring. Next, we describe the dynamics of the true (unobserved) level of prevalence and then specify a model for capturing the manager's beliefs and learning.

As depicted in figure 1, after testing a herds and observing K positive tests, the true number of herds in each health and trade status group is given by:

$$
\begin{aligned}
\tilde{S} &= S - (a - K) \\
\tilde{I} &= I - K \\
\tilde{M} &= M + K \\
\tilde{\widetilde{RT}} &= a - K
\end{aligned}
\tag{1}
$$

During testing, the susceptible population decreases by the number of "recently tested", RT, herds that test negative $(a - K)$, the infected group *decreases* by the number that test positive (K) and the movement control group *increases* by the number that test positive. We also relax Gramig and
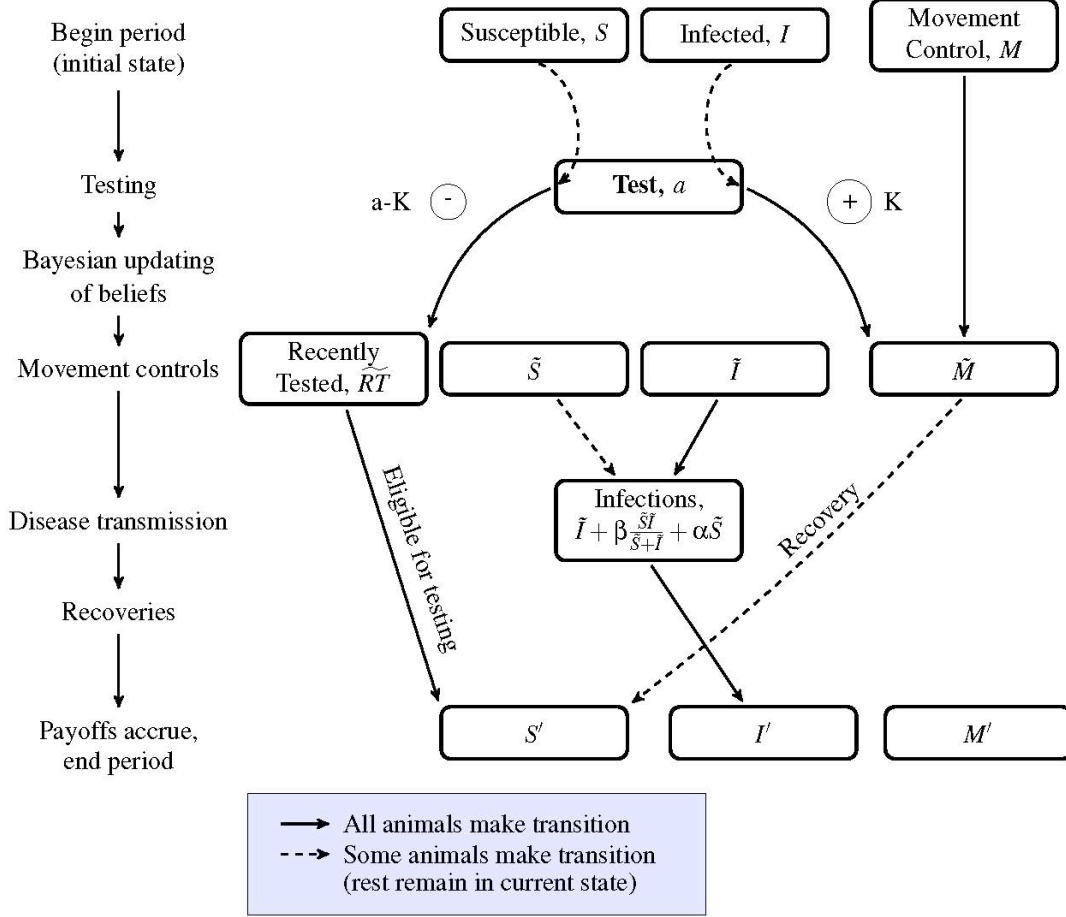
Figure 1: Dynamics for testing, Bayesian belief updating, herd management, transmission and recovery over one period

Horan's assumption that testing outcomes are deterministic: We model K as a binomial random variable conditional on a trials and a "success" probability given by the unknown prevalence, p. We assume that testing is perfect, which is not a strong assumption at the herd level in this context.[4] After testing, the true prevalence of bovine tuberculosis in the $N - \tilde{M}$ herds not in movement control is

$$\tilde{p} = \frac{\tilde{I}}{\tilde{S} + \tilde{I}} \qquad (1)$$

After testing, new infections result from trade between the remaining susceptible herds, $\tilde{S}$, and the remaining latently infected herds, $\tilde{I}$, and some fraction of herds in movement control recover. The physical or compartmental dynamics are captured in the following system of equations:

$$
\begin{aligned}
S' &= \tilde{S} - \beta \tilde{S}\tilde{p} - \alpha \tilde{S} + \gamma \tilde{M} + \widetilde{RT} \\
I' &= \tilde{I} + \beta \tilde{S}\tilde{p} + \alpha \tilde{S} \\
M' &= (1 - \gamma)\tilde{M}
\end{aligned}
\qquad (2)
$$

---

[4] Norby, et al. (2005) show that herd-level specificity is approximately 1 for herds with greater than 150 cattle. In our numerical case study, described in detail below, the average herd size is greater than 200.

where the prime modifier denotes state variables at the end of the current period (and beginning of the next), β is the herd-to-herd transmission coefficient,[5] α is the exogenous vector-to-herd transmission coefficient and γ is the recovery rate of infected herds under movement controls.[6,7] Herds transition from infected to susceptible only if they are placed under movement controls. We assume that the number of new infections in each period, i, is a deterministic function of the number of susceptible and infected facilities. New infections arise from two different sources and we denote these separate process as i_1 and i_2, where $i = i_1 + i_2$. First, we employ the common frequency-dependent transmission function for new infections, $i_1 = \beta \tilde{S} \tilde{p} = \beta \tilde{S} \tilde{I} / (\tilde{S} + \tilde{I})$, where β accounts for the joint probability of interaction and infection.[8,9] Second, infections may arise when susceptible herds come into contact with infected disease vectors such as possums and ferrets. We model exogenous infections as a time-invariant process that depends only on the number of susceptible herds, $i_2 = \alpha \tilde{S}$. Because our emphasis is on learning-by-doing we focus on the choice of testing and assume that vector (e.g. possum) management, and thus α, is constant. Incorporating i_2 is important since it reflects the reality that the disease can be managed but not fully eradicated from the larger system, which has strong implications for optimal control. In the system of equations in (1) – (3), infection risk is endogenous since infections are driven by true prevalence which is reduced when the management choice of testing level (a) identifies and selectively culls animals within infected herds (K). Because infected animals within a herd are culled and replaced with new (susceptible) stock, when herds do "recover" they do not become immune (as in some disease applications) but rather transition back to susceptible status.[10]

We assume that the manager knows the functional form of transmission dynamics but does not directly observe the actual changes in the underlying number of susceptible and infected herds which determine prevalence.

Learning and information dynamics

---

[5] As discussed by Reeling and Horan (2015), private investment in biosecurity could lead to a time varying beta term that responds to the perceived prevalence. Gramig and Horan (2011) find that private investment in biosecurity is, however, negligible when public testing levels lie within a range that is relevant for this paper.

[6] In reality it might be the case that the probability of infection for recently tested herds is lower than the overall prevalence. To account for this, in part, we assume that recently tested herds are not part of the new transmission process for one period. Note that infections with respect to $\tilde{S}$ (Equation 2.3) exclude recently tested herds (Equation 2.1).

[7] γ is known with certainty, and α and β are known and certain to managers that account for disease transmission.

[8] This functional form is a departure from Gramig and Horan's choice of the density dependent or mass-action transmission function: $\beta \tilde{S} \tilde{I}$. While the mass-action or density dependent transmission function performs well as a representation of within-herd disease transmission (Barlow, et al., 1997, McCallum, et al., 2001), the frequency-dependent transmission function is more appropriate for infection resulting from the trade of a finite number of animals (Cross, et al., 2013). The true transmission function that accounts for all forms of transmission likely lies between these two functional forms – density dependent at the local scale and frequency dependent at a regional scale.

[9] It is straightforward to extend this system to allow for stochastic transmission. For example, instead of replacing i_1 in Equation (7) with the deterministic value $\beta \tilde{S} \tilde{p}$, we could assume that $i_1 \sim Bin(\beta, \tilde{S} \tilde{p})$.

[10] Animals do not develop immunity to bovine tuberculosis independent of a vaccination, even if effectively treated.

We model uncertainty and learning with respect to infection prevalence using Bayesian updating. At the beginning of each period, the manager has initial beliefs regarding the proportion of facilities not under movement controls that are latently infected. Let g represent a probability density function characterizing beliefs over the true level of prevalence, p. We model these beliefs using a beta distribution because p lies on the unit interval, the beta distribution is relatively flexible and the shape parameters used to define the distribution may be updated based on testing results in straightforward way. The probability density function of a beta distribution is typically specified in terms of two shape parameters, s and f, for a given value of the unknown state: $g_{Beta}(p; s, f)$. We assume that the number of positive tests is a binomial random variable that depends on the number tests, a, and the unobserved true prevalence: $K \sim \text{Bin}(a, p)$.[11] While the process is one of sampling a finite population without replacement (hypergeometric sampling), the binomial assumption is a reasonable and tractable approximation when the number of tests is small relative to the population (Brunk, et al., 1968).

Given our assumptions regarding a beta distribution and binomial sampling, beliefs are updated using Bayes rule, with posterior beliefs given by $g_{Beta}(p; s, f \mid a, k) = g_{Beta}(p; s+K, f+a-K)$ (Gelman, et al., 2014). For our purposes, it will be convenient to specify beliefs in terms of transformed parameters, specifically the mean $\text{E}[p] \equiv \mu = s/(s+f)$ and $C = s+f$. The latter parameter is sometimes referred to as a "sample size" or "concentration" parameter since with Bayesian updating it grows by the number of observations.

$$\tilde{C} = C + a. \tag{3}$$

This convenient relationship results from the fact that the s parameter increases by the number of positive test results, K, and the parameter f increases by the number of negative results, $a-K$, during Bayesian updating.[12] Their sum, therefore, increases by the number of tests, a, independent of the results.

The updating process for μ is given by:

$$\tilde{\mu} = \frac{\mu C + K}{\tilde{C}}. \tag{4}$$

Conditional on initial beliefs (μ, C), after testing and observing results (a, K) posterior updated beliefs are given by $g_{Beta}\left(\tilde{p}; \mu, C \mid a, K\right) = g_{Beta}\left(\tilde{p}; \tilde{\mu}, \tilde{C}\right)$.

In addition to updating beliefs with new information as above, beliefs must also be adjusted to reflect the fact that the true level of prevalence changes according to the disease transmission and recovery dynamics in the system of equations in (3). In the parlance of control theory the former is an observation update in which the beliefs about the current state are updated using new observable information whereas the latter is a time update in which the beliefs about the future state are computed by combining the beliefs about the current state (the prior) with the transition model (the likelihood). Next, we derive an expression for the transmission dynamics of the true, unobserved prevalence p

---

[11] In our application testing rates are almost always below 10%. The binomial approximation of a hypergeometric process becomes problematic when testing rates are consistently above this value.
[12] This updating process relies on an assumption of perfect testing. We illustrate how to adjust the Bayesian updating step to incorporate imperfect testing in the Appendix.

and derive the implied transmission dynamics of the belief state variables, μ and C. First, we rearrange Equation (2) to express $\tilde{I}$ as a function of $\tilde{p}$ and known values:

$$\tilde{I} = \left(\tilde{S} + \tilde{I}\right)\tilde{p} = \left(N - M - a\right)\tilde{p} \tag{5}$$

Using equations (2), (3) and (6), end-of-period prevalence can be expressed as a function of $\tilde{p}$ and known values:

$$\begin{aligned} p'(\tilde{p}) &= \frac{I'}{S' + I'} = \frac{\tilde{I} + i}{N - M'} = \frac{\tilde{I} + \beta\tilde{S}\tilde{p} + \alpha\tilde{S}}{N - M'} \\ &= \frac{N - M - a}{N - (1 - \gamma)(M + K)}\left(\alpha + (1 + \beta - \alpha)\tilde{p} - \beta\tilde{p}^2\right) \end{aligned} \tag{6}$$

In Equation (7), the true end of period prevalence is increasing in $\tilde{p}$ since it sets the pre-transmission starting point and this value (along with α and β) determines new infections.

Approximating belief dynamics using density projection

Accounting for the fact that the true unobserved state is not fixed (as commonly assumed in learning models) but rather a moving target—according to the dynamics for p in Equation (7)—results in end-of-period beliefs that are not standard beta like the prior, but a quadratic transformation of a beta random variable. This loss of conjugacy (i.e. a prior and end-of-period distribution in the same family) presents a challenge for solving the decision model (described below). Beliefs given by a beta distribution can be fully characterized by two state variables for the two hyper-parameters (μ, C), but an analogous set of parameters are not readily identifiable for the quadratic beta function. Extending simple learning models to account for biophysical realities as we have done here has led to similar challenges in other resource management contexts. For example, in a fisheries model (Springborn and Sanchirico, 2013) show how this challenge of lost conjugacy can be overcome using a belief approximation approach known as density projection.

Density projection has been applied in the Bayesian empirical literature (e.g. Chen and Shao (1997)), the broader optimal control literature outside of economics (e.g. Maybeck (1982) and Zhou, et al. (2010)) and only recently in resource management (Springborn and Sanchirico, 2013). The approach involves approximating the true end-of-period distribution with a proxy that is close to the true distribution and in the same family as the prior. For our problem, this entails using a beta distribution to approximate the true distribution. The practical challenge is then to identify the best approximate to the true distribution. This is achieved by identifying parameters of the approximate distribution that minimize the Kullback-Leibler (KL) divergence between the true and approximate distributions. Zhou, et al. (2010) show that for distributions in the exponential family (including the beta), this approach is equivalent to matching the sufficient statistics of the true and approximate distributions.

We select the belief parameters of the approximate distribution to equate the sufficient statistics of the true and approximate distributions:

$$\begin{aligned} \mathrm{E}\left[\ln\left(p'\right)\right] &= \psi\left(\hat{\mu}\hat{C}\right) - \psi\left(\hat{C}\right) \\ \mathrm{E}\left[\ln\left(1 - p'\right)\right] &= \psi\left(\left(1 - \hat{\mu}\right)\hat{C}\right) - \psi\left(\hat{C}\right) \end{aligned} \tag{7}$$

where ψ is the digamma function and $\hat{\mu}$ and $\hat{C}$ represent the posterior shape parameter values that most closely fit a standard beta to the true distribution. The left-hand-sides of the equations in (8) represent the geometric means of the true distribution, which are numerically evaluated as follows:

$$E[\ln(p')] = \int_0^1 \ln\left(p'(\tilde{p})\right) g_{Beta}\left(\tilde{p}; \tilde{\mu}, \tilde{C}\right) d\tilde{p}$$
$$E[\ln(1-p')] = \int_0^1 \ln\left(1 - p'(\tilde{p})\right) g_{Beta}\left(\tilde{p}; \tilde{\mu}, \tilde{C}\right) d\tilde{p}$$

(8)

We integrate over $\tilde{p}$, instead of directly over p, using Equation (7) and the post-Bayesian updating distribution for $\tilde{p}$ specified by Equations (4) and (5). In the case of imperfect testing, g_{Beta} would be replaced by a posterior distribution that is no long conjugate (i.e. not beta).[13] The right-hand-sides of the equations in (8) are closed-form expressions for the geometric means of the approximating beta distribution. The sufficient statistic conditions in (8) are similar to the well-known conditions solved for maximum likelihood (ML) estimates of the beta distribution with the difference being that in ML the left-hand terms are given by the sample statistics generated using a data set (Nguyen, 2004).

In figure 2, we present an example of the combined Bayesian and physical updating. Suppose that initially the manager does not know if prevalence is low or high and has a relatively high degree of uncertainty ($\mu = 0.5, C = 3$). Next, suppose that a single test is conducted with a negative result ($a = 1, K = 0$). After Bayesian updating, the belief distribution shifts sharply to the left (the negative test suggests a lower prevalence in this case) and beliefs become more concentrated given the additional information ($\tilde{\mu} = 0.38, \tilde{C} = 4.1$). Finally, the true prevalence changes given physical dynamics: the belief distribution shifts back to the right (since there are new infections and no recoveries here) ($\hat{\mu} = 0.39, \hat{C} = 4.1$). Even in cases where the physical dynamic effect is modest, over a long time horizon accumulated error from failing to account for these physical dynamics is likely substantial.

The projected end of period beliefs presented in figure 2 closely track the actual distribution of beliefs, showing that the density projection approach approximates the true beliefs well.

The economic decision problem

We consider a manager who selects a level of testing, a, in each period to maximize the present value of the stream of expected profits. This optimal policy is a function of the current state of the system given by the number of herds in movement control and beliefs on prevalence: $X = \{M, \mu, C\}$. To simplify the problem, we examine the case in which any herds found to be infected are placed under movement controls for the remainder of the period in which they were found to be infected and reenter the pool of susceptible herds in the next period , $\gamma = 1$. This implies that $M = M' = 0$ and $\tilde{M} = K$.

---

[13] See Appendix for an expression of the posterior distribution when testing is imperfect, i.e. involves false positives and false negatives.
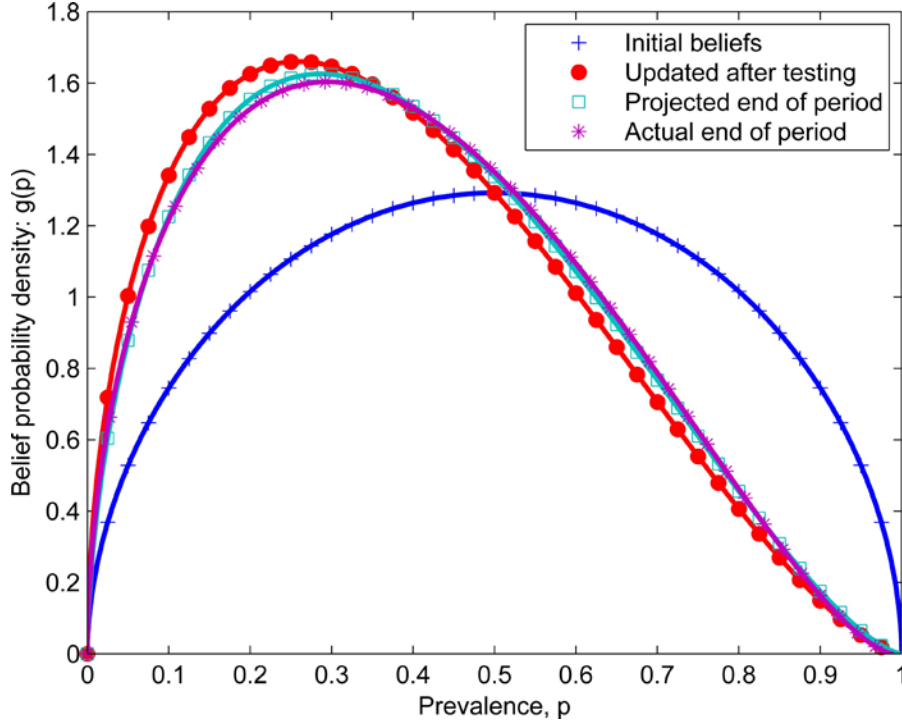
Figure 2: An example of updating an initial belief distribution using a single negative test result ($a = 1$, $K = 0$) and then again to account for physical dynamics

We generally follow Gramig and Horan (2011) in formulating the manager's payoff function. First, owners are assumed to be homogeneous except in the health and movement control status of their herd. Second, a susceptible herd generates profit from production independent of livestock trade, π_0 (e.g. from meat or dairy products). The baseline profit is deflated by direct production losses experienced when infected animals are present (i.e. if latently infected or under movement controls) and further losses result when the herd manager is *unaware* that infected animals are present (i.e. if latently infected). Third, managers experience gains from trade of magnitude π_1 if they are not under movement controls. Fourth, the coefficients of the quadratic testing and linear cleanup costs (or removing infected animals) are ω and r, respectively. True (unobserved) welfare from animal production in a period is thus given by

$$W(X) = \pi_0\left(S' + \delta\phi I' + \delta M\right) + \pi_1\left(N - M\right) - \omega a^2 - rK. \qquad (9)$$

The parameters δ and ϕ capture losses from infection: $\delta < 1$ is the proportion of π_0 remaining for a herd if infected, and $\phi < 1$ is an additional discount factor from being unaware of infection.

In the case of complete within-period recovery ($\gamma = 1$), Equation (10) may be rewritten as

$$W(X) = \pi_0\left(S' + \tilde{\delta}I'\right) + \pi_1 N - \omega a^2 - rK, \qquad (10)$$

where $\tilde{\delta} = \phi\delta$. Expected welfare conditional on an action and state is given by

$$\begin{aligned}
EW(X) &= \pi_0\left(E[S'|\mu,C,a] + \tilde{\delta}E[I'|\mu,C,a]\right) + \pi_1 N - \omega a^2 - rE[K] \\
&= \pi_0\left(\left(1 - E[\mu'|\mu,C,a]\right)N + \tilde{\delta}E[\mu'|\mu,C,a]N\right) + \pi_1 N - \omega a^2 - rE[K].
\end{aligned} \qquad (11)$$

11

The dynamic optimization problem is specified by the Bellman equation and state transition dynamics:

$$J(X) = \max_a \left\{ \text{EW}(X) + \rho \text{EJ}(X') \right\}$$

(12)

$$\text{s.t.} \begin{cases} \mu'(X,a,K) \\ C'(X,a,K) \end{cases} \begin{array}{l} \text{determined by Bayesian} \\ \text{updating and density projection} \end{array}$$

where $\rho \in [0,1]$ is a discount factor.[14]

Disentangling the effects of learning and state dynamics

The full management model specified in (13) incorporates two innovations yet to be combined in disease management models: (1) Bayesian learning, which facilitates accounting for the future value of current investment in information, and (2) adjustment for transmission dynamics to account for expected changes in the unobserved underlying state. In order to disentangle the implications of these two components, we also consider three alternative manager types which omit one or both of these elements. All four manager types account for recovery of herds under movement controls, because there would otherwise be no incentive to test. We outline the four manager types and the distinguishing features of the how they handle information in figure 3. For ease of reference, we denote the full model manager type as Bayesian with transmission, B_T. The non-Bayesian with transmission manager type, NB_T, treats current uncertainty about prevalence as fixed but accounts for new infections due to transmission. The Bayesian with no transmission manager type, B_NT, bases beliefs on testing but does not account for new infections. Finally, the non-Bayesian with no transmission manager type, NB_NT, treats current uncertainty regarding prevalence as fixed and ignores the potential for new infections.

Assessing the alternative manager types above also facilitates a comparison between the full model we propose (B_T) and existing approaches in two strands of the management literature, which also correlate with observed management approaches. These comparisons are summarized in the final column of figure 3. The NB_T approach is indicative of the approach taken in previous bio-economic models of disease, where uncertainty regarding disease prevalence—and the opportunity to reduce that uncertainty—has typically been ignored (e.g. Bicknell, et al. (1999) and Gramig and Horan (2011)). Management decisions therefore ignore the value of information regarding prevalence provided by testing, i.e. managers do not account for the way in which current testing provides knowledge that improves future decisions. The value of testing is attributable exclusively to prevalence reduction. The NB_T manager considered here does acknowledge uncertainty over prevalence but, in keeping with the literature just described, forgoes the opportunity to update that imperfect information about prevalence using test results. Thus, this manager type represents the standard non-learning stochastic optimization approach to resource management. The NB_T manager type is the closest approximation to the current regulatory strategy that infrequently changes background testing rates over larger areas in responses to unanticipated changes in prevalence.

---

[14] In cases where $\gamma \neq 0$ the number of facilities under movement controls change according to the equation of motion: $M_0 = (1-\gamma)(M+K)$.

Manager type	Manager uses test results to improve estimate of prevalence via Bayesian updating	Bayesian updating rules:	Manager accounts for disease trans-mission	Physical state updating rules: updates according to equation (7), where	Policy analog
B_T	Yes (B)	Hyper-parameters update according to equations (4) and (5)	Yes (T)	$\alpha > 0, \beta > 0$	Proposed optimal manager
B_NT			No (NT)	$\alpha = \beta = 0$	Active adaptive management approach with simplified dynamics; focuses on optimal endogenous learning at expense of assuming the target for learning is fixed; approximate analog to historical approach in New Zealand.
NB_T	No (NB)	Hyper-parameters remain fixed $\tilde{\mu} = \mu$ $\tilde{C} = C$	Yes (T)	$\alpha > 0, \beta > 0$	Typical approach to bioeconomics of disease (and other natural resource systems); focuses on system dynamics but does not pursue optimal endogenous learning; approximate analog to status quo in New Zealand.
NB_NT			No (NT)	$\alpha = \beta = 0$	Benchmark naïve manager; serves as point of reference

Figure 3: Overview of manager types with implications for updating and policy analogs

The second strand of related literature focuses on adaptive resource management, which spans from early foundations in ecology (Walters and Hilborn, 1978) to more recent applications in resource economics (e.g. Springborn (2014)). Fully optimal adaptive management models account for the expected value of information when choosing current actions, given the capacity to update beliefs and improve future decisions. In the present disease management context, this would imply ignoring the role of transmission, and treating initial prevalence as a fixed and unknown parameter.[15] The B_NT manager type thus emulates the approach of the bulk of adaptive management literature to date, which simplifies system dynamics to maintain conjugacy and thus complexity in belief dynamics. This manager is also historically relevant as it (imperfectly) represents New Zealand's disease management approach in the late 1970's, when disease spread was severely underestimated (due to the delayed onset of clinical symptoms and an ignorance about the important role disease vectors in transmission) (TBfree New Zealand, 2015). The manager in this example was very sensitive to information (i.e. negative test results) and unable or unwilling to acknowledge biophysical realities.

Lastly, the NB_NT manager neither utilizes Bayesian learning nor accounts for transmission. This manager is the least sophisticated of the alternatives and only accounts for changes in beliefs about prevalence resulting from recovery of herds identified as infected (as do all others). The NB_NT manager provides a "naïve" benchmark from which to compare the incremental improvements from Bayesian adaptive management and accounting for the moving target dynamic created by transmission.

---

[15] bTB transmission had historically been ignored because it spread and manifested slowly, and vector transmission was poorly understood.

Case study: Bovine Tuberculosis in New Zealand

We apply our methodology to the problem of managing bTB infections in a subgroup of herds within the Waikato region of New Zealand. Waikato – among other regions of New Zealand – has struggled with bTB monitoring and control since the end of the nineteenth century, when it became a notifiable[16] disease, due to the expansive nature of its livestock production and abundance of vector populations. We consider a system that includes 100 herds where a maximum of 25 herds may be tested in a given month.[17] Furthermore, parameter estimates were selected from the economic and epidemiological literature to characterize producer profits, testing and cleanup costs and transmission. Parameter choices and their sources are provided in table 1 in the Appendix. We solve the full problem specified in (13) and for each of the alternative manager types using value function iteration (Judd, 1998).

We estimate the value function over a discrete grid of points selected from the continuous domain of the state space given by the hyper-parameters. Grid points were selected with higher density over regions for which value function nonlinearity was pronounced, specifically for values of C near the bottom of its range $(C = 2)$ and values of μ near each bound of its unit interval range.

While C is technically unbounded from above, further increases in concentration beyond 40 or 50 had very little effect on optimal policy, thus this parameter was bounded at 61.

Results

Policy functions for each of the four manager types are presented in figure 4. These results depict optimal testing rates conditional on the current belief state, represented by μ and C. Corresponding solutions for the value functions are presented in the Appendix. One of our over-arching questions is, which of the two key components—accounting for Bayesian learning or the moving target dynamic— has a more pronounced effect on optimal actions and payoffs? From figure 4 we see that, relative to the baseline (NB_NT), accounting for transmission (NB_T) clearly leads to a more striking change in policy than accounting for learning (B_NT). This indicates that optimal management responds more strongly to the prospect of new infections than to the opportunity to learn. Accounting for transmission has the expected impact of increasing optimal testing—when the manager accounts for the fact that infected herds generate further infections, the value of their identification through testing increases.

As expected, optimal testing is increasing in expected prevalence, μ, for all manager types, until the upper bound on testing is reached. In the other dimension, optimal testing is only sensitive to concentration (confidence) at low levels of C. The direction of this effect depends on the manager type and μ. It is surprising that optimal testing (information seeking) is *increasing* as uncertainty falls in some situations. In adaptive management models it is typical for the level of the informative action to be decreasing as confidence grows: as uncertainty falls, so does the value of additional information. Counter to these expectations, optimal testing under B_T is weakly increasing as confidence grows, as indicated by convex curvature in figure 4, suggesting a precautionary testing approach. Hauser and Possingham (2008) find that such precautionary approaches may exist in the short and medium time horizons. This result can be understood by considering the baseline model

---

[16] Notifiable livestock diseases are defined as those that must be reported to the relevant government agency. Diseases are generally categorized as notifiable if they pose substantial risk to other livestock producers or consumers of animal products.

[17] 25 (out of 100 herds) tests per month represents a high upper-bound. Herds identified as infected and placed under movement controls are only tested at six month intervals, suggesting that at most approximately 17% of facilities are being tested in a given month (TBfree New Zealand, 2014).

(NB_NT) and the incremental effects of incorporating Bayesian learning (B_NT) and transmission (NB_T). From the baseline results (NB_NT) we observe testing rates that are increasing in C over much of the domain, showing that the surprising outcome is driven in part simply by the direct effect of accounting for uncertainty over prevalence (irrespective of its effect on transmission and the opportunity to reduce uncertainty). Second, as noted above, transmission has pronounced effect on policy, here serving to expand the domain over which convex curvature shows optimal testing that is increasing in C (see NB_T in figure 4).

Ultimately as we move from the NB_T to B_T policy function, only the bottom left corner (low confidence and expected prevalence) displays convex curvature attributable to the Bayesian learning component. While the value of information effect again suggests that testing should be decreasing in confidence, here we see the opposite. This likely emerges due to less of a need for precaution when taking an adaptive approach—with Bayesian learning (B_T), if true prevalence turns out to be high, this will be learned and policy adjusted accordingly. In contrast, the NB_T manager knows that he will not be aware of any such a deviation from expectations and thus compensates accordingly (with higher testing). Overall, we see that accounting for both Bayesian learning and the moving target dynamic is crucial for developing intuition and because they interact in a complex fashion.
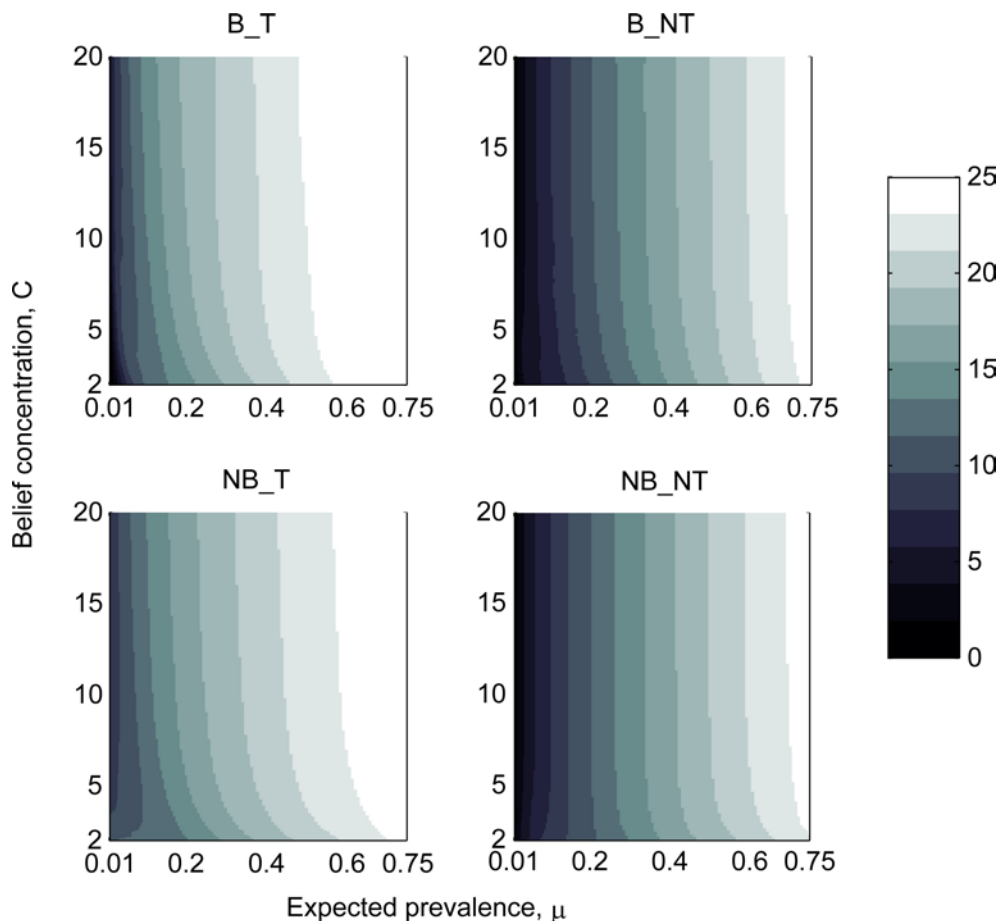

Figure 4: Policy functions (optimal testing rates, given by shading) for each man- ager type, over the belief space specified by expected prevalence (*μ*, horizontal axis) and concentration (*C*, vertical axis)

Policy functions are informative but ultimately only provide a snap-shot of the immediate response to a particular belief state. To evaluate differences in the dynamic paths of key variables—testing,

prevalence, beliefs and welfare—we use Monte Carlo simulations governed by each of the four manager types.

Policy comparison

To simulate, we must first specify an initial belief distribution that is shared by each manager type. Because a learning methodology is of most interest when uncertainty is moderate to high and since most disease systems of interest feature a relatively low prevalence, we specify initial beliefs for all manager types as $\mu_0 = 0.1$, and $C_0 = 5$. In the Appendix we present results for other initial belief distributions which yield qualitatively similar results. In each of 5,000 Monte Carlo simulations, the true initial prevalence, p_0, is drawn randomly from the initial belief distribution. This process for random selection of the initial, unobserved underlying state follows from the assumption that initial beliefs reflect uncertainty but are not biased. We simulate the system over a 30 year time horizon. In each period, the manager considers beliefs, chooses the testing level, responds to the results, transmission occurs and payoffs accrue. The repetition of this sequence generates dynamic paths for testing intensity, true prevalence and any divergence or convergence between expected and true prevalence. Averages for each of these series across Monte Carlo simulations are shown in figure 5 for each manager type. Since changes in these particular variables are small after a decade, figure 5 presents the first 10 years. Estimates of testing, prevalence, average belief error after 10 year and cumulative welfare after 30 years for each of the manager types are provided in the Appendix in table 2.

Figure 5a shows a consistent ordering in the average rate of testing after the first two years: NB_T > B_T > B_NT > NB_NT.[18,19] These testing rates are statistically significantly different after two and 10 years based on a Wilcoxon rank-sum test $(p < 0.01)$. As expected, the NB_NT manager tests least since transmission and the value of information are ignored. On average in the beginning of the time window, the B_T manager pursues moderate testing: more than the B_NT manager (to account for prevalence growth via transmission) but less than the NB_T manager (given the greater need of the non-learner to be precautionary).[20] In the latter periods, the B_T manager is able to sustain low prevalence with average testing rates slightly lower than those of the B_NT manager.

The ordering of testing intensities above predictably inverts to determine the ordering of average prevalence (figure 5b), and these differences were statistically significant $(p < 0.01)$. The NB_T and B_T managers test sufficiently to suppress prevalence; the B_NT and NB_NT managers fail to suppress prevalence. The suppression achieved by the B_T and NB_T managers is plausible. There exist numerous cases outside of New Zealand where eradication has been achieved,[21] while New Zealand eradication or near-eradication has been achieved in pockets where the natural vector of reintroduction (possum populations) are low (TBfree New Zealand, 2013). It may seem counterintuitive that both testing and prevalence are higher on average under NB_NT than B_NT for the first year and a half. This dynamic emerges because under B_NT learning occurs from the outset, error in beliefs falls relative to NB_NT, and testing is responsive—elevated for simulations with high

---

[18] After the 10th year, the B_T manager conducts slightly fewer tests than the B_NT manager.
[19] Note that while the NB_T manager consistently exceeds a 10% testing rate, it also does not use the binomial approximation during learning.
[20] Note that the NB_T manager is not precautionary by design but rather, in seeking an optimal policy under uncertainty that is not expected to attenuate under Bayesian learning, the manager responds to the ever-present possibility of high prevalence with a level of testing that is precautionary, relative to other managers.
[21] For example, Australia has been able to sustain eradication since 1997 with only minor violations and the United States has been able to maintain disease free status in all but a few regions (Animal Health Australia, 2015, USDA-APHIS, 2015).

initial prevalence and reduced in the opposite case. This adaptive response allows B_NT to achieve great reduction in prevalence with less testing, on average.

While it might not be surprising that suppression is not achieved with the least savvy approach (NB_NT), a question arises as to why the B_NT manager fails to achieve strong control of the disease. Insight comes from considering the evolution of how wrong each manager's beliefs are, on average. In figure 5c we present the average absolute belief error (i.e. difference between expected and actual prevalence) for each manager type. The NB_NT manager falsely concludes that early testing and treatment permanently suppress the disease: testing falls, transmission takes off, beliefs are not updated and error is large. Bayesian learning on its own provides an escape from this trap but only partially. The B_NT manager is consistently overly optimistic with respect to future prevalence (since transmission is ignored). Learning partially corrects for this consistent dynamic bias. As a result B_NT management allows for a substantial prevalence.

The shortcomings of the B_NT model are compounded by its own special trap. In some problematic cases (approximately 28% of simulations with our selected prior distribution), the B_NT manager observes results that rapidly cause beliefs to converge on an extremely low prevalence. Ignoring transmission, the B_NT manager believes a state of permanent near-eradication has been achieved. Testing ceases and prevalence takes off, leading to infection of all herds. The B_NT manager fails to correct course given the particular form of this trap in which informative testing is deemed no longer necessary. This is of course an extreme case—in reality with ubiquitous infection, such a strategy would eventually be challenged. However, it highlights the danger of developing management models in which the complexity of system dynamics have been ironed out in order to simplify or make tractable an adaptive management approach.

As shown in figure 5c, we find that the non-learning strategy NB_T actually shows less error in beliefs than the Bayesian learning strategy B_NT. This result arises as follows. The NB_T manager compensates for irreducible uncertainty and transmission with higher testing than the others. This effectively suppresses long-run prevalence to a very low level across a wide range of true initial prevalence starting points. Because the NB_T manager expects and achieves a very low prevalence level, there is very little room for error in beliefs regarding prevalence. We find that in what little error there is, the NB_T manager believes prevalence to be worse than truth and so continues to act with sufficient aggression to keep prevalence tamped down. If it was infeasible to achieve a low prevalence level, the NB_T manager would likely not be able to maintain such a limited error in beliefs.

From a methodological perspective, this reduction in error provides additional evidence that that our projection approach, which minimizes KL-divergence, serves to improve the accuracy of beliefs. To our knowledge, there does not exist intuition for an acceptable range for the KL-divergence statistic, and defining "high" divergence is subjective. We are ultimately concerned with the effectiveness of the approach for updating a manager's beliefs such that they are more informative about the underlying truth. The average absolute belief error (AAE) of the B_T manager consistently falls over time, to 0.033 after 2 years and 0.022 after 5 years. The consistently lower AAE of the B_T relative to the B_NT manager reflects the improved accuracy of beliefs from learning *and* accounting for a moving target using density projection.
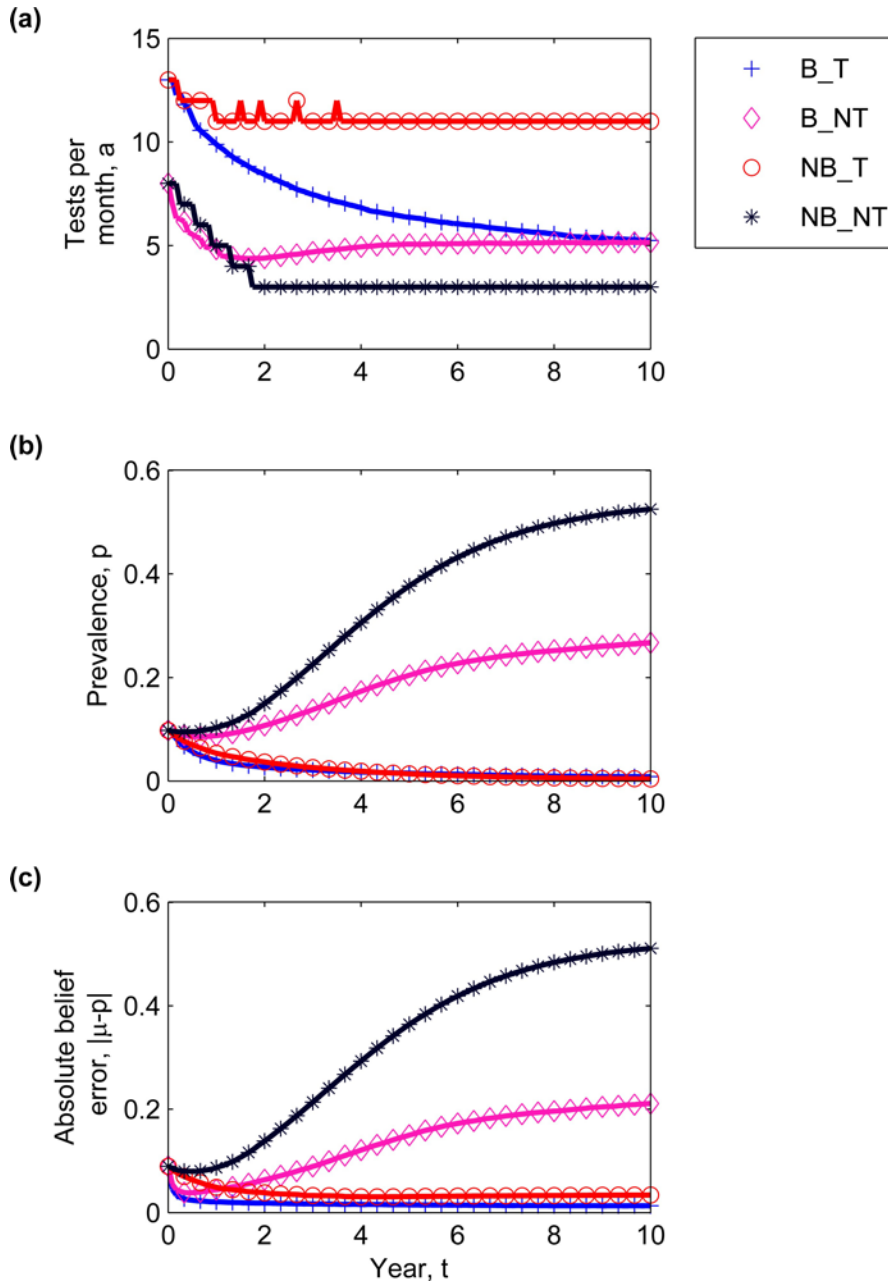
Figure 5: Average outcomes over Monte Carlo simulations for (a) testing rates, (b) prevalence levels and (c) absolute belief error (i.e. difference between expected and actual prevalence) for each manager type

Welfare comparison

The efficiency of public control of infectious disease is ultimately informed by the long-run cumulative welfare that is provided, as measured by producer profits, which are reduced by infection, and testing and cleanup costs[22] in this case. The trajectories of the average across Monte Carlo

---

[22] The central role of testing in this paper motivated a more thorough examination of our specification of testing costs. A 50% increase or decrease in testing costs did not qualitatively change our results, including ordering of testing levels, prevalence and welfare. The NB_T manager maintains a precautionary testing approach even at high test costs.

simulations of the cumulative profits accrued through year $t \in [0,30]$ are shown in figure 6. Because the absolute scale of the vertical axis depends on the application of interest and time horizon, we show the percent difference between each of the manager types and the NB_NT manager (the 0 line indicates performance equal to NB_NT).

As expected, the testing choices of the B_T manager yield the highest cumulative value in the medium and long-run, while the NB_NT manager realizes the lowest. For the other manager types (B_NT and NB_T), more substantial gains are realized from an acknowledgment of epidemiological factors than through the use of information to update beliefs. These differences are all statistically significant $(p < 0.01)$. This difference is largely driven by consistent under-testing among managers that ignore transmission (B_NT and NB_NT) relative to those managers that do account for transmission (B_T and NB_T), which leads to significantly higher levels of costly prevalence.

The B_T manager pursues moderate testing on average: more than the B_NT manager in the early stages to account for prevalence growth (via transmission) but less than the NB_T manager since any long-run growth in prevalence is identified, prompting an increase in testing. The B_T testing on average leads to low prevalence and yields higher welfare than the alternatives, and is the first-best alternative.

The NB_NT manager is consistently over-optimistic about prevalence (given the singular direction of transmission and an inability to update beliefs through learning), under-tests, and as a result experiences significant losses to productivity. By construction, this is the least savvy of the four managers, and is used as a baseline to show relative welfare gains. However, it should be noted that the NB_NT approach may yield larger profits in the short-run (due to low testing expenditures) than the other three alternatives.

Ex ante, the relative magnitudes of the welfare gains from learning or accounting for transmission are unclear, particularly given the substantial uncertainty in this example. Transmission dynamics allow for more accurate accounting for the future value of testing and control, while learning reduces uncertainty and updating may partially compensate for the downward bias from omitting transmission. For this application, we find that accounting for transmission (or more generally, the upward shift in the state variable caused by physical dynamics) is more important than Bayesian learning. However, this may be largely attributable to the manager's ability to drive the prevalence to very low levels, and a Bayesian learning approach may be more desirable in contexts in which the state variable consistently takes on intermediate values.

The shortcomings of the B_NT model are compounded by problematic scenarios in which the manager rapidly moves to a belief that prevalence is low or zero.[23] Failing to account for reintroduction or transmission leads to a non-negligible propensity for the B_NT manager to believe a state of permanent eradication has been achieved. Under this belief, it is optimal to cease testing, which eliminates any possibility of updating beliefs. The disease subsequently takes off, leading to infection of all herds and a severe reduction of producer profits. As acknowledged earlier the similarities between this subset of results and the historical example of public management failure lead us to acknowledge it rather than discard it as unrealistic or problematic.

The NB_T manager compensates for irreducible uncertainty with higher testing in each belief state than the others. On average, the response is overly aggressive (relative to B_T), but has the benefit of suppressing prevalence to compensate for uncertainty. This aggressive level of testing has

---

[23] Experiencing the illusion of permanent eradication is more likely when the initial concentration of beliefs (c_0) is small, the initial expected prevalence (μ_0) s small, the randomly chosen initial prevalence is low (p_0) and early testing yields more negative results than expected (κ is a random variable). See the Appendix for a discussion of an extended set of initial beliefs that highlight this relationship.

two additional benefits. First, it leads to a lower average prevalence than the testing schedule used by the B_T manager in the long-run. Figure 5b shows that the average prevalence experienced by the B_T manager is lower in the first 2-3 years, but experiences a slightly higher prevalence thereafter. Second, it ensures the accuracy of the expected prevalence. Figure 5c shows that, on average, for B_T and NB_T the absolute differences between expected and actual prevalence are comparable after the first few years. However, the raw residuals generated by the NB_T model (not shown) have a positive bias, indicating a consistent overestimate of prevalence.
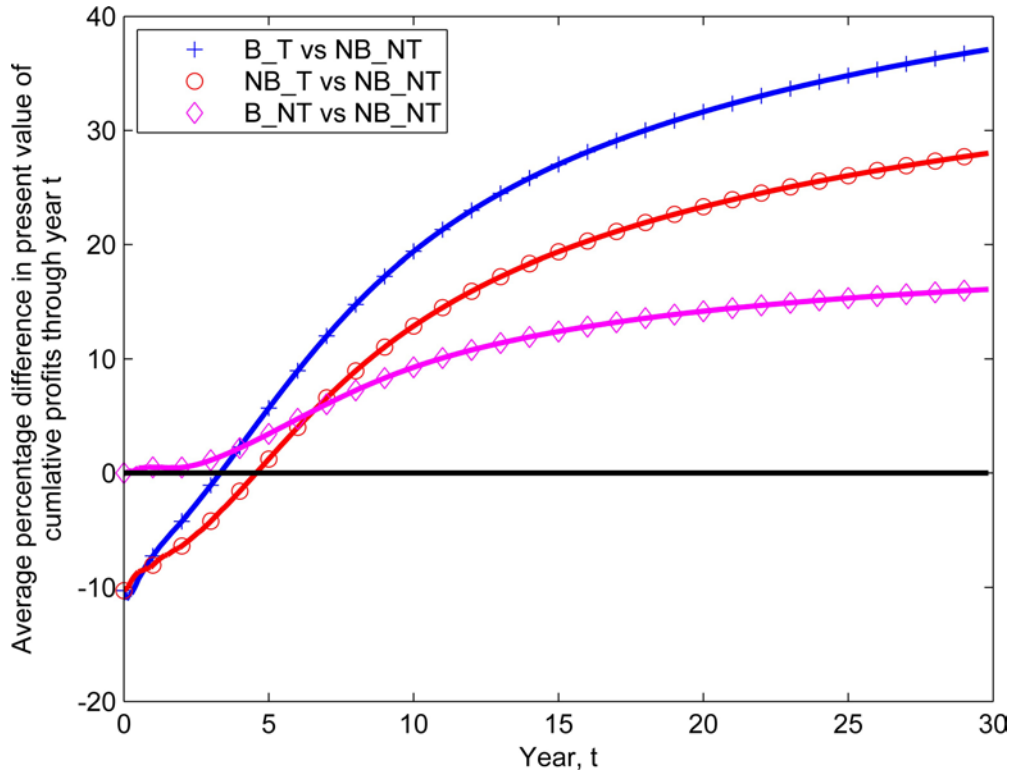

Figure 6: Average percentage increase in the present value of cumulative profits (relative to the NB_NT baseline)

Over the 30 year time span, the gains from either Bayesian learning or accounting for transmission are substantial (approximately 16% and 28%, respectively). Additionally, introducing a consideration of epidemiological forces to the B_NT manager results in substantial improvements (approximately 21 percentage points). The incremental gain from introducing Bayesian learning to a manager who already considers epidemiological forces yields a more modest welfare improvement (approximately 9 percentage points).

As shown in figure 5, the B_T and NB_T managers effectively drive prevalence to low levels. However, the more moderate and adaptive strategy of the B_T manager leads to welfare improvements, which are primarily accrued in the first few years of the program. The average present value of cumulative profits realized by the B_T manager exceed those realized by the NB_T manager within the first year. This gap continues to grow throughout the remaining periods. This consistent divergence suggests that there are continual gains to accounting for learning even after those variables shown in figure 5 reach consistent values.

The managers that do not account for transmission (B_NT and NB_NT) consistently under-test, providing short-run cost savings at the cost of dampened long-run producer profits attributable to infections. This undesirable outcome is inevitable for the NB_NT manager, but may be avoided the B_NT manager if the trap of ceasing testing is avoided. For those simulations in which the B_NT

manager continues to test in each period, the testing rate remains below but is substantially closer to that selected by B_T manager. These simulations result in less severe under-testing, and substantially higher cumulative profits. The interested reader can see a supplementary appendix online for a scenario $\left(\mu_0 = 0.1, C_0 = 61\right)$ in which the all of B_NT simulations avoid the trap of ceasing testing.

To alleviate concerns that the gains realized from accounting for information through Bayesian learning are driven by outliers as opposed to more broadly-based improvements, in figure 7a we present the proportion of simulations in which the present value of cumulative profits realized by the B_T manager exceeds those realized by the NB_T manager over a 30 year period. We also show in figure 7b the distribution of the percentage differences in the present value of cumulative profits between the B_T and NB_T models by year 30.
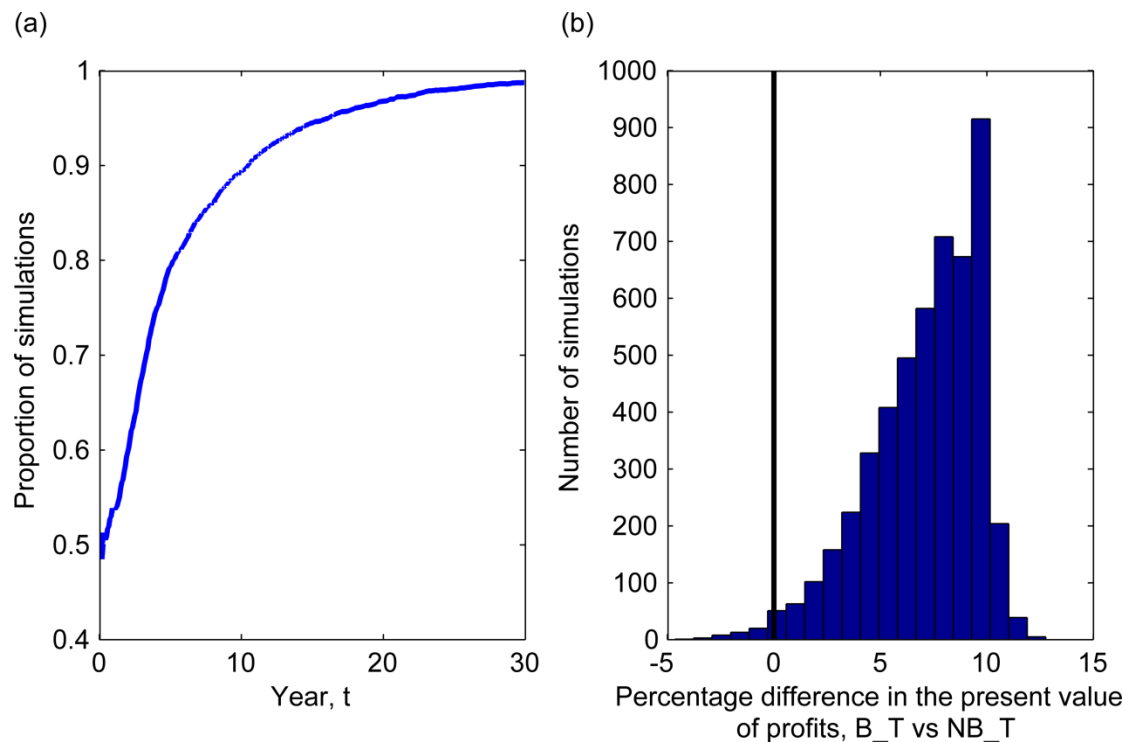
(a)
(b)

Figure 7: (a) Proportion of Monte Carlo simulations in which the present value of cumulative profits realized by the B_T manager type exceed those realized under NB_T by year *t*; (b) percentage difference in the present value of cumulative profits realized by the B_T versus NB_T manager type after 30 years across Monte Carlo simulations

The fraction of Monte Carlo simulations for which B_T outperforms NB_T is initially close to 0.5 (i.e. reflects random test results), but grows throughout the time horizon. The monotonicity of the path and large value in the final period suggests that the welfare dominance of B_T over NB_T is not driven in Monte Carlo simulation by a few extreme differences, but rather a consistent increase in performance: in 98.7% of the simulations B_T outperforms NB_T.

Additional evidence supporting the theory that B_T consistently outperforms NB_T is provided in the histogram shown in figure 7b. The bulk of simulations (approximately 98.7%) lie in the positive domain, and the distribution of these differences is single-peeked. The p-value associated with a t-test and a Wilcoxon rank-sum test for the hypothesis that the mean of the present value of the cumulative profits for the B_T manager is greater than that for the NB_T manager are both <0.01.

Discussion and conclusion

In this article, we develop practical modeling tools for adaptive management of resource systems that are not only uncertain but also in flux. The integration of learning and accounting for dynamics facilitates assessment of adaptive management when an important state is both uncertain and changing over time. This methodology advances a previously unexplored application in disease management, providing insights into the value of adaptive control and a warning against smoothing over important physical dynamics. Potential inaccuracies arising from incomplete specification of the physical dynamics also highlight the importance of accurate estimates of key physical and economic features when managing infectious diseases.

This research contributes to a broader literature at the intersection of traditional resource management under uncertainty (without learning) and adaptive management. The true, underlying conditions of resource systems are often imperfectly observed. In this context, we find that incorporating (1) learning and (2) an understanding of the true physical dynamics have important, interdependent effects on optimal management and payoffs. In the context of endemic disease management, it is more costly to ignore physical dynamics (transmission) than it is to ignore learning. Furthermore, abstracting away from underlying system dynamics to focus on learning can lead to consistent error in beliefs and even belief traps in which temporary management success ends information gathering, which in turn allows for the system to deteriorate, unbeknownst to managers. Our findings on the relative costs of ignoring learning versus physical dynamics should be generalized beyond disease systems with caution. Ignoring perturbations (like transmission) that push the unobserved state in a particular direction is likely to be more problematic than ignoring perturbations that are less unidirectional.

The current, less responsive broad-based disease testing strategy of New Zealand is an artifact of historical technological limitations, which reduced the opportunities to learn and adapt. Managers have needed to provide clear, concise guidelines, while tabulation of information from producers has been challenging and slow. These restrictions have, however, been reduced through New Zealand's investment in infrastructure, making cost reductions from and adaptive approach feasible. Specifically, recently developed data collection and existing information provision systems (National Animal Identification and Tracing System and TBfree New Zealand) allow for the necessary rapid communication between regulators and producers. This paper provides evidence that a more responsive approach to broad-based disease testing would be welfare enhancing *if* the cost savings exceed the costs of implementing a more complex management scheme.

The welfare gains realized from accounting for transmission—even in the absence of Bayesian learning—provide additional evidence supporting Gramig and Horan's (2011) call for high-quality estimates and information to model disease. Gramig and Horan note the absence of rigorous epidemiological studies linking herd management and trade to disease spread, which leads modelers to use simplified and potentially erroneous transmission functions. Our results indicate that the use of inaccurate estimates of key transmission parameters leads to inefficient testing—even when a manager learns and accounts for transmission—attributable to incorrect assessments of future disease spread and erroneous beliefs about prevalence.

This approach could be extended to incorporate additional epidemiological and economic realism. Several such extensions were identified throughout the article: adding stochasticity to the transmission function; allowing herds to remain under movement controls for longer periods; introducing risk aversion; and expanding system scale (i.e. the number of herds). Further, less-straightforward extensions would require additional theoretical development. For example, in reality the unobserved state dynamics (e.g. transmission function) may not be well understood and an object for learning in its own right. Incorporating additional data sources in the decision process would more realistically model the learning process and would allow managers to avoid Bayesian traps.

References

Animal Health Australia. 2015. *Australian Bovine Tuberculosis Surveillance Project.*

Atallah, S.S., et al. 2014. "A plant-level, spatial, bioeconomic model of plant disease diffusion and control: Grapevine leafroll disease." *American Journal of Agricultural Economics*:aau032.

Barlow, N., et al. 1997. "A simulation model for the spread of bovine tuberculosis within New Zealand cattle herds." *Preventive veterinary medicine* 32:57-75.

Beef & Lamb New Zealand. 2014. *New Season Outlook 2014-2015.*
Bicknell, K.B., J.E. Wilen, and R.E. Howitt. 1999. "Public policy and private incentives for livestock disease control." *Australian Journal of Agricultural and Resource Economics* 43:501-521.

Bond, C.A., and J.B. Loomis. 2009. "Using numerical dynamic programming to compare passive and active learning in the adaptive management of nutrients in shallow lakes." *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie* 57:555-573.

Brunk, H., J.E. Holstein, and F. Williams. 1968. "The Teacher's Corner: A Comparison of Binomial Approximations to the Hypergeometric Distribution." *The American Statistician* 22:24-26.

Chades, I., et al. (2012) "MOMDPs: a solution for modelling adaptive management problems." In *Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12).*

Chen, M.-H., and Q.-M. Shao. 1997. "Performance study of marginal posterior density estimation via Kullback-Leibler divergence." *Test* 6:321-350.

Cousins, D., Florisson, N. 2005. *A Review of Tests Available for Use in Diagnosis of Tuberculosis in Non-bovine Species.*

Cross, P.C., D. Caillaud, and D.M. Heisey. 2013. "Underestimating the effects of spatial heterogeneity due to individual movement and spatial scale: infectious disease as an example." *Landscape ecology* 28:247-257.

de Lisle, G. 1993. *Bovine Tuberculosis: The New Zealand Problem.* Proceedings of the Grassland Association.

Fackler, P.L., and R.G. Haight. 2014. "Monitoring as a partially observable decision problem." *Resource and Energy Economics* 37:226-241.

Filipe, J.A., et al. 2012. "Landscape epidemiology and control of pathogens with cryptic and long-distance dispersal: sudden oak death in northern Californian forests." *PLoS computational biology* 8:e1002328.

Gelman, A., et al. 2014. *Bayesian data analysis*: Taylor & Francis.

Gramig, B.M., and R.D. Horan. 2011. "Jointly determined livestock disease dynamics and decentralised economic behaviour." *Australian Journal of Agricultural and Resource Economics* 55:393-410.

Haight, R.G., and S. Polasky. 2010. "Optimal control of an invasive species with imperfect information about the level of infestation." *Resource and Energy Economics* 32:519-533.

Hauser, C.E., and H.P. Possingham. 2008. "Experimental or precautionary? Adaptive management over a range of time horizons." *Journal of Applied Ecology* 45:72-81.

International, L. 2013. "New Zealand Dairy Statistics 2012-2013."

Johnson, F.A. 2011. "Learning and adaptation in the management of waterfowl harvests." *Journal of environmental management* 92:1385-1394.

Judd, K.L. 1998. *Numerical methods in economics*: MIT press.

Kling, D.M., J.N. Sanchirico, and P.L. Fackler. 2016. "Optimal monitoring and control under state uncertainty: application to lionfish management." *Working Paper*.

Lane, D.E. 1989. "Operational research and fisheries management." *European Journal of Operational Research* 42:229-242.

LIC International and TBfree New Zealand. 2013. *New Zealand Dairy Statistics 2012-2013.*

Maybeck, P.S. 1982. *Stochastic models, estimation, and control*: Academic press.

Mbah, M.L.N., and C.A. Gilligan. 2010. "Optimization of control strategies for epidemics in heterogeneous populations with symmetric and asymmetric transmission." *Journal of theoretical biology* 262:757-763.

McCallum, H., N. Barlow, and J. Hone. 2001. "How should pathogen transmission be modelled?" *Trends in ecology & evolution* 16:295-300.

Ministry of Primary Industries. 2012. "Farm Monitoring 2012: National Sheep and Beef."

Moore, A. 2008. "Managing Populations in the Face of Uncertainty: Adaptive Management, Partial Observability and the Dynamic Value of Information." University of Science.

Morrison, W., et al. 2000. "Pathogenesis and diagnosis of infections with Mycobacterium bovis in cattle." *Veterinary Record* 146:236-242.

Nguyen, T.T. 2004. "Maximum likelihood estimators of the parameters in a beta distribution."

STATISTICS TEXTBOOKS AND MONOGRAPHS 174:229-236.

Norby, B., et al. 2005. "Use of simulation modeling to estimate herd-level sensitivity, specificity, and predictive values of diagnostic tests for detection of tuberculosis in cattle." *American journal of veterinary research* 66:1285-1291.

Office of Management and Budget. 2003. *Circular A-4: Regulatory Analysis.*
Porphyre, R. 2008. "Factors Associated with the Transmission Dynamics of Bovine Tuberculosis in New Zealand." Massey University.

Reeling, C.J., and R.D. Horan. 2015. "Self-Protection, Strategic Interactions, and the Relative Endogeneity of Disease Risks." *American Journal of Agricultural Economics* 97:452-468.
Regan, T.J., I. Chadès, and H.P. Possingham. 2011. "Optimally managing under imperfect detection: a method for plant invasions." *Journal of Applied Ecology* 48:76-85.

Regan, T.J., et al. 2006. "Optimal eradication: when to stop looking for an invasive plant." *Ecology Letters* 9:759-766.

Rosvear, M., and D. Ulrich. 2010. *Bovine Tb Strategy: Review of Costs.*

Schlaifer, R., and H. Raiffa. 1961. "Applied statistical decision theory."

Shea, K., et al. 2014. "Adaptive management and the value of information: Learning via intervention in epidemiology." *PLoS Biol* 12:e1001970.

Springborn, M. 2014. "Reducing damaging introductions from international species trade through invasion risk assessment." *Invasive Species in a Globalized World: Ecological, Social, and Legal Perspectives on Policy*:163.

Springborn, M., and J.N. Sanchirico. 2013. "A density projection approach for non-trivial information dynamics: adaptive management of stochastic natural resources." *Journal of Environmental Economics and Management* 66:609-624.

Statistics New Zealand. 2012. "Agricultural Production Statistics."

TBfree New Zealand. 2013. *Animal Health Board Annual Report 2012/2013.*

---. 2015. *A brief history of bovine TB.*

---. 2014. *Moving Infected Stock.*

U.S. Fish and Wildlife Service. 2013. *Adaptive Harvest Management: 2013 Hunting Season.* Washington D.C.

USDA-APHIS. 2015. *Bovine Tuberculosis Disease Information.*

Walters, C. 1986. "Adaptive management of renewable resources."

---. 1974. "An interdisciplinary approach to development of watershed simulation models." *Technological Forecasting and Social Change* 6:299-323.

Walters, C.J., and R. Hilborn. 1976. "Adaptive control of fishing systems." *Journal of the Fisheries Board of Canada* 33:145-159.

---. 1978. "Ecological optimization and adaptive management." *Annual review of Ecology and Systematics*:157-188.

Wieland, V. 2000. "Monetary policy, parameter uncertainty and optimal learning." *Journal of Monetary Economics* 46:199-228.

Zhou, E., M.C. Fu, and S. Marcus. 2010. "Solving continuous-state POMDPs via density projection." *Automatic Control, IEEE Transactions on* 55:1101-1116.

Appendix
Bayesian updating under imperfect testing

In the context of bTB in New Zealand's cattle, the large average herd size implies a negligible probability of false negatives for all infected animals within an average herd. Furthermore veterinary pathologists can rule out false positives during necropsy. However, herd-level specificity and sensitivity may be less than 1 for other diseases or for bTB in other settings. For example we expect lower testing precision if herd sizes are small or diagnostic technology is imprecise.

When imperfect testing is present, the simple Bayesian updating rules used in (4) and (5) no longer hold. Instead we obtain a posterior distribution, g_{imp}, that is no longer beta and thus the prior is no longer conjugate. To derive g_{imp} we revise the likelihood function portion of the posterior to adjust for inaccuracies.

The probability of observing a positive test result under imperfect testing

$$\overline{p}(p;TPR,TNR) = p \cdot TPR + (1-p) \cdot (1-TNR), \tag{13}$$

depends on the true prevalence, $p,$ the true positive rate, $TPR,$ and the true negative rate $TNR.$ The probability of observing K positive tests result from a tests is still binomial but with success parameter $\overline{p}(\cdot)$:

$$\Pr\left(K \mid a, \overline{p}(p)\right) = \binom{a}{K} \overline{p}(p)^K \left(1-\overline{p}(p)\right)^{a-K}, \tag{14}$$

where the parameters of $\overline{p}(\cdot)$ have been suppressed for simplicity. The posterior distribution given by Bayes' rule is proportional to the product of the likelihood function and the prior (beta) distribution:

$$g_{imp}\left(\tilde{p} \mid a, K\right) \propto \overline{g}(\tilde{p} \mid a, K) \equiv \left[\overline{p}(\tilde{p})^K \left(1-\overline{p}(\tilde{p})\right)^{a-K}\right]\left[\tilde{p}^{\mu C-1} \left(1-\tilde{p}\right)^{(1-\mu)C-1}\right], \tag{15}$$

where p has been replaced with $\tilde{p}$ to indicate beliefs over post-testing prevalence. Since this expression no longer simplifies to one that is proportional to a beta distribution, we do not have the convenience of a conjugate prior. However, the true posterior after testing can be calculated numerically as $g_{imp}\left(\tilde{p} \mid a, K\right) = \overline{g}(\tilde{p} \mid a, K) \Big/ \int_0^1 \overline{g}(\tilde{p} \mid a, K) d\tilde{p}$ and substituted for $g_{beta}(\cdot)$ in the density projection step found in (9).

Imperfect testing also changes the efficacy of targeted controls that follow positive test results. While the forms of these changes are context-specific, they amount to fewer herds being returned to the susceptible population for bTB. Rather than K herds being returned, only $K*TPR$ are successfully restored. $K(1-TPR)$ are treated without any actual change in health status.

Parameter estimates used in simulation

Table 1: Parameter Estimates Used in Numerical Analysis

Parameter	Symbol	Value	Source(s)
Transmission Rate	β	0.1	Porphyre (2008) and TBfree New Zealand (2013)
Vector transmission rate	α	3.9E-5	(TBfree New Zealand (2013))
Baseline profits	π_0	\$11,168	Beef & Lamb New Zealand (2014); International (2013) and TBfree New Zealand (2013) ; Ministry of Primary Industries (2012); Statistics New Zealand (2012)
Production losses from infection	$1-\tilde{\delta}$	0.65	Bicknell, et al. (1999)
Testing cost per herd	ω	\$3.52	Rosvear and Ulrich (2010)
Cleanup cost per herd	r	\$7.04	Rosvear and Ulrich (2010)
Monthly discount factor	ρ	0.997	Office of Management and Budget (2003)

Table 1 provides the essential parameters used during our numerical analysis, described in the Results section. Because a direct estimate of the herd-to-herd transmission rate was not available for New Zealand, we estimate a rate as follows. Given an estimated annual rate of replacement per facility at 21% (Porphyre, 2008)   and an average herd size of 259 (LIC International and TBfree New Zealand, 2013, Statistics New Zealand, 2012), the average monthly influx of new animals to a facility is estimated at 4-5. Facilities identified as infected have an average of approximately 15 reactor animals are present (TBfree New Zealand, 2013). A high estimate of the probability of transmission (conditional on trade with an infected facility) is given by assuming that all replacement animals are purchased externally, infected individuals are no less likely to be traded, and one or more infected animals in trade guarantees transmission. In that case, the probability of transmission—as given by one minus the hypergeometric density of zero infected animals in a trade of four to five animals—is 0.21-0.26. This is likely an over-estimate since infected animals are less likely to be selected for trade, one infected animal in a trade does not guarantee infection of the new herd and buyers may take precautionary measures. If we roughly account for these factors by assuming that the number of infected animals that are truly candidates for trade is about half of infected individuals, we arrive at the transmission rate used of 0.1.

Value function

Only the value function for the B_T manager is represented in figure 8 which shows the present value of a stream of *expected profit* in each possible belief state. The value functions for the other managers follow similar patterns, but which differ slightly based on how expectations about future profits are developed. The optimized objective function shown is associated with the policy function shown in figure 4a in the main text.
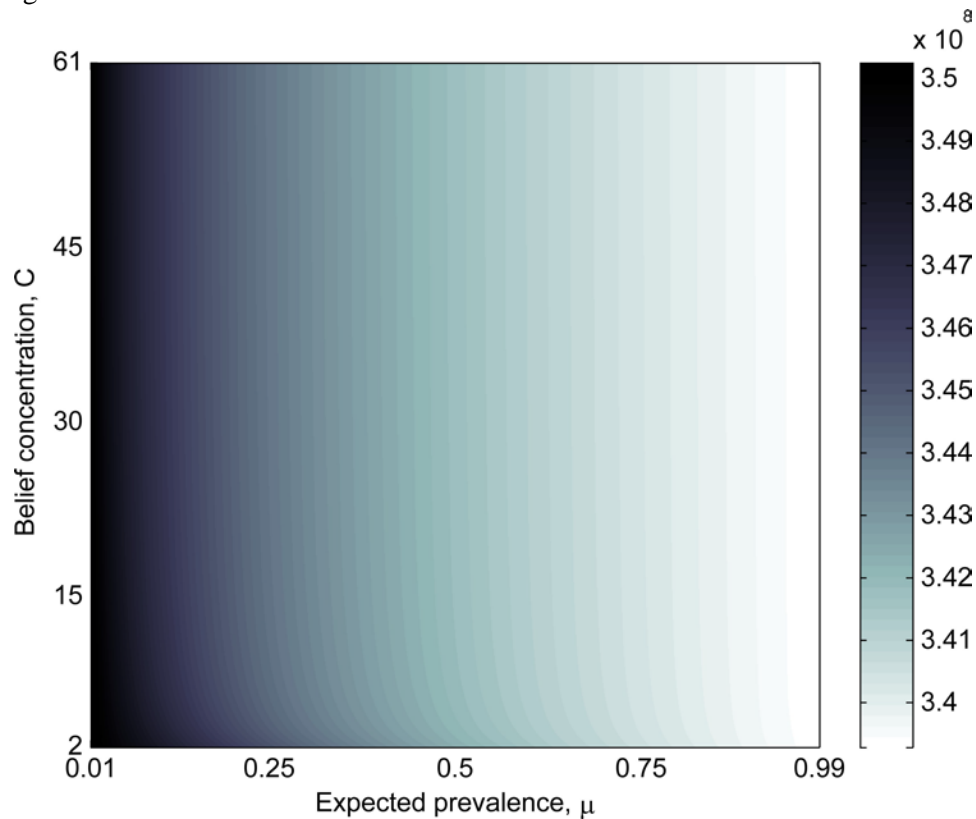

Figure 8: Value function for all values of *μ* and *C* for the B_T manager

The value function follows a predictable pattern. Most saliently, expected profits are decreasing in expected prevalence, μ. When information is scarce (i.e. C is small), expected value is increasing in C, suggesting a direct value to the precision of the prevalence estimate. This value is attributable primarily to more efficient testing strategies, and to a manager's ability to rule out high prevalence scenarios.

The other managers' value functions follow similar patterns with an important difference. The decision makers that do not account for transmission experience higher expected profits because their expectations fail to capture losses due to future disease spread.

Testing and outcome estimates

Table 2. Mean and standard deviation (in parentheses) for variables depicted in Figures 5a-5c (testing level, prevalence and belief error) and Figure 6 (profit increase) after the terminal period in each figure (year 10 or 30)

Variable	Manager type			
	B_T	**NB_T**	**B_NT**	**NB_NT**
Tests	5.2	11.0	5.2	3.0
(year 10)	(2.6)	(0)	(3.7)	(0)
Prevalence (year	0.0085	0.0042	0.27	0.53
10)	(0.020)	(0.014)	(0.32)	(0.28)
Belief error (year	0.013	0.034	0.21	0.51
10)	(0.010)	(0.0071)	(0.34)	(0.26)
*Profit increase**	41.44	32.20	20.55	-
(year 30)	(21.80)	(20.88)	(30.01)	-

* Average percentage increase in the present value of cumulative profits relative to the NB_NT baseline

Table 2 contains means and standard deviations for each of the action and outcome variables represented in the terminal periods of figures 5 and 6: testing, prevalence and belief error after 10 years and the percentage increase from the baseline (NB_NT) after 30 years. Statistics for each manager type are calculated across the 5,000 Monte Carlo simulations. Note that testing for the non-learning mangers is zero because their testing strategies area not responsive to observed differences in testing outcomes.