Dynamic-Bayesian disease management under state uncertainty: learning and bovine tuberculosis control in New Zealand cattle

Matthew MacLachlan, University of California, Davis, mjmaclachlan@ucdavis.edu
Michael Springborn, University of California, Davis, mspringborn@ucdavis.edu
Paul Fackler, North Carolina State University, pfackler@ncsu.edu

Selected Paper prepared for presentation at the Agricultural & Applied Economics Association's 2014 AAEA Annual Meeting, Minneapolis, MN, July 27-29, 2014.

I. Introduction

The World Organization for Animal Health (OIE) conjectures that animal disease results in an average decline of 20% in livestock production worldwide (Vallat 2014). Furthermore, livestock diseases may infect consumers and local wildlife, resulting in decreased human and ecosystem health. In turn, a decline in consumption may result from aversion to potentially dangerous products or opposition to environmental degradation. Depending on the pathogen of interest, exclusion from foreign markets for either live animals or animal products may result from the presence of a pathogen. These negative impacts collectively provide a strong incentive to prevent the introduction of a pathogen, and to mitigate spread once the pathogen is introduced. Due to the scale and nature of transmission, however, individual producers are often insufficiently motivated or incapable of providing the socially desirable level of control. Extensive governmental intervention is necessary when individual producers are not fully motivated to account for potential costs associated with disease transmission or when a disease vector resides on public lands. An essential consideration of these interventions is how the government collects and utilizes data to more efficiently apply control efforts in future periods.

Controlling many infectious diseases among livestock is challenging due to incomplete information regarding which facilities contain infected livestock and uncertainty regarding diseases prevalence. When a government decides to intervene, the central veterinary authority (which is the Animal Health Board in New Zealand (AHB)) must balance the costs of testing for infected livestock and treatment with the gains realized from reduced prevalence and additional information. The value of information from testing includes enhanced immediate targeted control efforts and more informed investment decisions in future applications of broad-based control measures (levels of testing, self-protection, etc.). The focus of this paper is to examine the process of learning and adaptive control within the context of bovine tuberculosis (bTB) surveillance and control among New Zealand dairy and

beef cattle, accounting for the value of testing for both targeted and more efficient broad-based controls.

bTB is an infectious and potentially fatal disease of both animals and humans that persists throughout much of the world, including but not limited to the United Kingdom, Japan, Mexico, much of Central and South America, as well as much of Africa and Asia (Cousins and Florisson, 2009). Some of these nations have established eradication programs to improve productivity and trade prospects. In those nations with eradication programs, this disease is a consistent threat to trade. Importing nations are likely to avoid exports from nations with high disease prevalence due to transmission via live cattle movements to producers. Once introduced, bTB may be transmitted via milk, particularly raw milk, to consumers.

While eradication or near-eradication has been achieved in many developed countries such as Australia, Canada and the United States, reductions in prevalence in New Zealand have come slowly and at a substantial financial cost. Despite intensive and sustained control efforts in New Zealand, eradication has been encumbered by characteristics of the pathogen and environmental and anthropogenic factors: a long incubation period, a pervasive but elusive wild host (the common brushtail possum), imperfect testing methods, and the diffuse nature of production. These features have allowed bTB to remain endemic among New Zealand cattle herds since the mid-to-late 20th century, and substantially increased the difficulty of determining prevalence. For an endemic disease such as this, there may be an especially high value to the central veterinary authority in understanding the prevalence of the disease. More specifically, this information may be used to better inform future testing choices, private bio-security investments, and vector population control.

We review bTB testing in New Zealand and explore policy options under uncertainty and learning. This research will improve understanding of livestock disease control, and provide general insights into the value of learning in a system in which the conditions are dynamically changing. We

present the concepts that motivate a numerical simulation in which the AHB's beliefs regarding disease transmission and updating are modeled explicitly under a variety of physical conditions.

To address our problem, we establish a novel methodology for decision making under uncertainty. More specifically, we allow for learning about an uncertain prevalence that is partially determined by the level of control applied within the system. Within our framework, learning is predictable and may be used to improve the efficiency of future decisions. Our approach could be applied to a broad array of systems for which a state is unknown and affected by the decision makers choices (fisheries, invasive species, fossil fuel extraction). We find substantial efficiency gains from a testing regime that is motivated both by test results and an understanding of the physical dynamics underlying the system relative to regimes in which either information is not considered, or physical dynamics are ignored.

The remainder of the paper is organized as follows: Section II reviews the available monitoring and control measures available to the AHB; Section III presents our theoretical framework; Section IV details our methodology for implementation; Section V presents results; and Section VI discusses our findings and potential implications of our research.

II. Monitoring and Control of Bovine Tuberculosis

The control of bovine tuberculosis is difficult, and primarily depends on the collection and use of data regarding disease prevalence among livestock and wild host vectors. Additional measures are employed once an infected facility is identified. An understanding of these available options provides insights into those questions that our analysis is able to address.

Monitoring

Information regarding prevalence is primarily obtained through three sources: field tests, tests and necropsies conducted at the slaughter house, and population monitoring and testing of wild hosts. In our context, interactions with infected wildlife is rare and all animals sent to slaughter are tested. Therefore, we choose to focus on the AHB's response to obtaining information from field tests. This abstraction allows us to more cleanly capture the relationship between information obtained in field tests and the future choices regarding the intensity of field tests. It may, however, introduce positive bias for the value of field testing if this work is extended to an empirical approach.

The primary test used to identify infected animals is the caudal fold test (CFT), which is inexpensive and familiar to veterinarians. Commercially viable, alternative testing methods are available, including gamma-interferon assays and comparative cervical tuberculin tests, but these have not proven to be any more effective as a primary screening tool (Whipple et al., 1995; CDFA 2013). Instead, because each of the tests may result in false positives for different reasons, the alternatives are applied when an animal is found to be infected by the CFT. If an animal receives all positive test results, then it is subject to slaughter and a subsequent necropsy, which will ultimately rule out the vast majority of false positives. We are thus able to ignore this possibility in our model.

False negatives pose an ongoing challenge, particularly if the number of infected animals on a given facility is low. For example, the herd-level sensitivity is equivalent to the sensitivity of single test if a single animal is infected. Further complicating an objective measure of false negatives, the sensitivity of a test is a choice of the central veterinary authority. The CFT results in a skin reaction on all subjects, and the AHB has the authority to define the threshold above which the animal is consider infected (OIE, 2009). Because we are unable to capture the within-herd disease dynamics using our model and having sensitivity as an endogenous variable makes the model intractable, we simply model sensitivity as

perfect. This assumption is strong, and subsequent research will be targeted at more rigorously addressing this limitation.

Mitigation

A cost effective treatment of bTB does not currently exist for cattle. Instead, the aforementioned tests are applied to cattle to determine their health status. If an animal is found to be infected, then that animal is immediately marked as "reactor" animal. These animals are sold for a substantial loss (approximately 65% (Bicknell, Wilen and Howitt, 1999)). After a reactor is initially identified, a facility must be free of reactor animals for two sequential, annual tests for movement restrictions to be lifted. Additionally, the facility on which the animal is placed under movement controls, which imposes costly testing requirements for animals sold from the site. Additionally, all animals that come in contact with a reactor on the receiving property within 90 to 120 days must also be tested.

Prevention

There are several preventative measures available to producers. These primarily include bio-security measures that take one of two general forms. First, producers may reduce the interactions between livestock and wild hosts. Interactions between infected wild hosts and livestock are limited within our region of interest. Second, testing and quarantine measures may be taken to reduce the likelihood latently infected cattle brought will be brought onto the facility. The second set of endogenous private controls are important within our context, but beyond the scope of this paper. For a rigorous analysis of the feedbacks between testing and biosecurity, see Gramig and Horan (2011).

III. Methodology: Theoretical Framework

We model the spread and control of bTB among herds within the Waikato region of New Zealand. Within this region, the primary source of infection is intranational trade. For simplicity, we ignore trade from regions outside of Waikato and from disease vectors such as possums or ferrets. Additionally we do not extend our model to capture infections among domesticated deer. The unit of observation is the herd, making it impossible for us to capture within herd disease dynamics and the dependence of transmission on these dynamics. However, this more parsimonious model allows us to more tractably capture the decision maker's learning and decision processes with respect to testing at the cost of precision in modeling the transmission process.

We choose to model our system as a Markov decision process (MDP) to capture the inherent randomness in test results that are applied to a population with susceptible and latently infected individuals. We also introduce Bayesian learning to model an essential feature of all disease control systems: the decision maker is uncertain regarding the true prevalence and uses test result to reduce this uncertainty. Such problems are referred to as a Partial Observability Markov Decision Processes (POMDP) (Fackler and Haight, 2014), and have been applied to problems in which the decision maker learns about a fixed underlying probability. We extend this framework to capture a problem in which the underlying probability evolves over time and is endogenous to capture the relationship between beliefs regarding an unknown prevalence and control efforts. The decision maker learns about a moving target that changes both with disease transmission and her own application of testing and resulting targeted controls.

Order

A well specified order of events is necessary to establish a mathematical framework in discrete time, which will facilitate a deeper analysis. Each time step is equal to one month, and we assume the

following order of events: at the beginning of the period, the decision maker selects the number of herds to test and observes the results, movement controls are imposed on those herds found to be infected, then new infections occur, some fraction of facilities recover from movement control status, herds are returned from "recently tested" status to susceptible, and finally payoffs accrue to producers. While it is unrealistic to assume that all infections occur prior to the advent of recoveries in a given period, the number of recovered facilities is small. Figure 1 provides a visual guide to this ordering. A more complete depiction is provided in Figure 3.

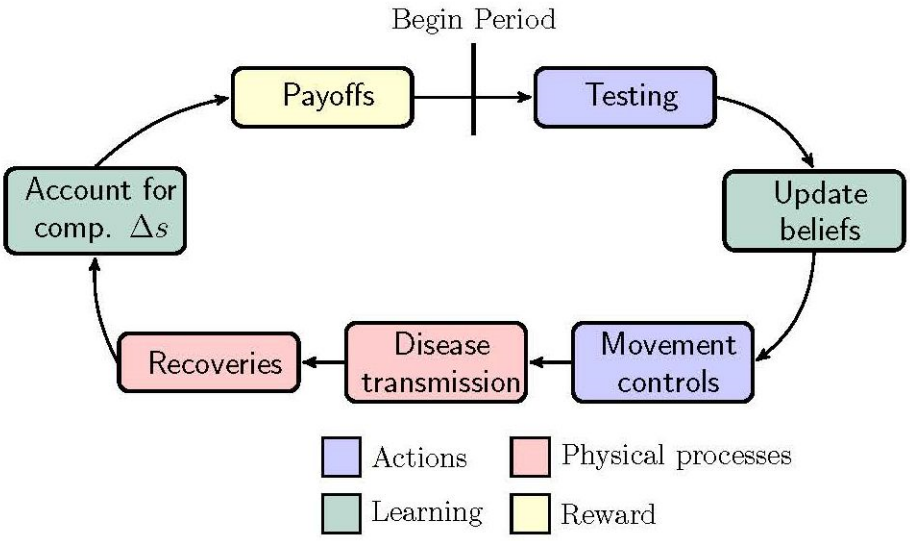

Figure 1. Order of events.

Disease dynamics

We adapt the approach of Gramig and Horan (2011) for modeling transmission and control of bTB among cattle populations with a few modifications. In Figure 3, we illustrate combined dynamics for Bayesian belief updating and changes in 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, *M* herds have been identified as infected and are placed on movement controls. While *N* and *M* are observed, in contrast to Gramig and Horan, the division of remaining herds between *S* and *I* is unobserved. This implies that the proportion of uncontrolled herds that are latently infected,

$p = I/(S+I)$, is unknown. After testing *a* herds and observing *K* positive tests, the *true* number of herds in each health status group is given by:

$$\begin{aligned} \tilde{S} &= S - a + K \\ \tilde{I} &= I - K \\ \tilde{M} &= M + K. \end{aligned} \tag{1}$$

Here, the susceptible population decreases by the number of facilities 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 relax the assumption that testing outcomes are deterministic: we model K as a binomial random variable conditional on a trials and "success" probability, p.[1] Excluding herds in movement control, after testing the prevalence of bovine tuberculosis is

$$\tilde{p} = \frac{\tilde{I}}{\tilde{S} + \tilde{I}}. \tag{2}$$

Next, new infections result from trade between herds in $\tilde{S}$ and herds in $\tilde{I}$. Finally, some fraction of herds in movement control recover. The physical or compartmental dynamics are:

$$\begin{aligned} S' &= \tilde{S} - \beta \tilde{S}\tilde{p} + \gamma \tilde{M} \\ I' &= \tilde{I} + \beta \tilde{S}\tilde{p} \\ M' &= (1-\gamma)\tilde{M}, \end{aligned} \tag{3}$$

where ' denotes state variables at the end of the current period (also the beginning of the next period), β is the transmission coefficient and γ is the recovery rate from movement controls. We employ the common frequency-dependent transmission function for new infections, denoted $i = \beta \tilde{S}\tilde{p} = \beta \frac{\tilde{S}\tilde{I}}{\tilde{S}+\tilde{I}}$, where β accounts for the joint probability of interaction and infection. This functional form is a

[1] This model can be extended to account for the fact that testing is imperfect, e.g. with sensitivity of σ.

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 transmission function performs well as a representation of within-herd disease transmission (Barlow et al., 1997; Neill et al., 1989), the density-dependent transmission function is more appropriate for infection resulting from the trade of a finite number of animals (McCallum et al., 2001).

Learning dynamics

We model uncertainty and learning with respect to infection prevalence using Bayesian updating. At the beginning of each period, the decision maker has initial beliefs regarding the proportion of facilities not under movement controls that are latently infected. Let g represent beliefs over the true level of p. We use a beta distribution to model these beliefs since *p* lies on the unit interval and its density can be updated based on testing results in straightforward way. A beta distribution is typically specified in terms of two shape parameters, $g_{Beta}(p; s, f)$. Conditional on the binomial outcome $K \sim Bin(a, \mathrm{p})$,[2] updating beliefs on *p* using Bayes rule results in posterior beliefs given by $g_{Beta}(p; s, f \mid a, k) = g_{Beta}(p; s+K, f+a-K)$ (Gelman et al. 2004). For our purposes it will be convenient to specify beliefs in terms of transformed parameters, specifically the mean $E(p) \equiv \mu = \frac{s}{s+f}$ 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:

$$\tilde{C} = C + a. \tag{4}$$

The mean updates according to:

$$\tilde{\mu} = \frac{\mu C + K}{\tilde{C}} \tag{5}$$

[2] Sampling is inherently hypergeometric because a facility will not be tested multiple times within a single month, barring the discovery of a reactor animal. However, when the number of tests is small relative to the total population, binomial sampling is a reasonable approximation of hypergeometric sampling (Cryer, 2014).

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

True prevalence transition dynamics

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 dynamics in the system (3). We assume that true transmission is unobserved, deterministic and understood. Next, we derive an expression for the transmission dynamics of the unobserved variable p and show the implications for the 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:

$$\tilde{I} = \left(\tilde{S} + \tilde{I}\right)\tilde{p} = \left(N - M - a\right)\tilde{p} \tag{6}$$

Using the system of equations in (2), (3) and (6), p' can be expressed as a function of $\tilde{p}$ and known values:

$$\begin{aligned} p'(\tilde{p}) &= \frac{I'}{S' + I'} = \frac{\tilde{I} + i}{N - M'} = \frac{\tilde{I} + \beta\tilde{S}\tilde{p}}{N - M'} \\ &= \frac{\left(N - M - a\right)\tilde{p} + \beta\left(N - M - a\right)\left(1 - \tilde{p}\right)\tilde{p}}{N - M'} \\ &= \frac{N - M - a}{N - M'}\left(\left(1 + \beta\right)\tilde{p} - \beta\tilde{p}^2\right). \end{aligned} \tag{7}$$

In the formulation represented in (7), we can see that the expected prevalence in the subsequent period will be strictly decreasing in the number of tests. Similarly, the expected prevalence is increasing in the post-Bayesian updating prevalence for all feasible values of $\tilde{p}$. In this way, control efforts and new infections are working in opposite directions during the physical updating. The magnitude of this effect

depends on both deterministic and random processes, making the conditions under which a force dominates unclear.

Approximating beliefs using density projection

The true dynamics for p' in Equation (7) do not lead explicitly to end-of-period beliefs that are exactly beta[3]. A solution to the problem requires that these beliefs be approximated in some way. We use density projection to identify an approximate characterization of beliefs that is "close" to true beliefs. Specifically, we use a beta distribution to approximate these beliefs and select the belief parameters in a way that minimizes the Kullback-Leibler divergence between the true and approximate beliefs. Zhou et al. (2010) showed that for distributions in the exponential family, this approach is equivalent to matching the first two moments of the true and approximate distributions. These two moment conditions are given by

$$\begin{aligned} E(\ln(p')) &= \psi(\mu' C') - \psi(C') \\ E(\ln(1-p')) &= \psi\left((1-\mu')C'\right) - \psi(C'), \end{aligned} \tag{8}$$

where ψ is the digamma function. Using Equation (7) we numerically evaluate the left hand side terms in Equation (8) as follows

$$\begin{aligned} E(\ln(p')) &= \int_0^1 \ln\left(p'(\tilde{p})\right) g_{Beta}\left(\tilde{p};\tilde{\mu},\tilde{C}\right) d\tilde{p} \\ E(\ln(1-p')) &= \int_0^1 \ln\left(1-p'(\tilde{p})\right) g_{Beta}\left(\tilde{p};\tilde{\mu},\tilde{C}\right) d\tilde{p}, \end{aligned} \tag{9}$$

and then solve for μ' and C'.[4]

[3] The distribution would instead be a generalized beta. A generalized beta has the undesirable property of being able to take on values outside of the $[0,1]$ interval.

[4] It is straightforward to extend this system to allow for stochastic transmission. For example, instead of replacing i in Equation (7) with the deterministic value $\beta\tilde{S}\tilde{p}$, we could assume that $i \sim Bin(\beta,\tilde{S}\tilde{p})$. Then Equation (7) would take the form $p'(\tilde{p},i)$ and the expressions in (9) would require an additional integral over i.

In Figure 2, we provide a representation of a simulated example of Bayesian and compartmental updating. Initially, the AHB has a higher degree of uncertainty. It conducts 10 tests with negative results, which results in the belief distribution shifting sharply left. Subsequently, the distribution shifts slightly back to the right as the AHB accounts for new infections. Because the transmission coefficient is relatively small for bTB, this secondary effect is small in any given period. However, the consequences for failing to account for these physical dynamics over a long time horizon may be severe.

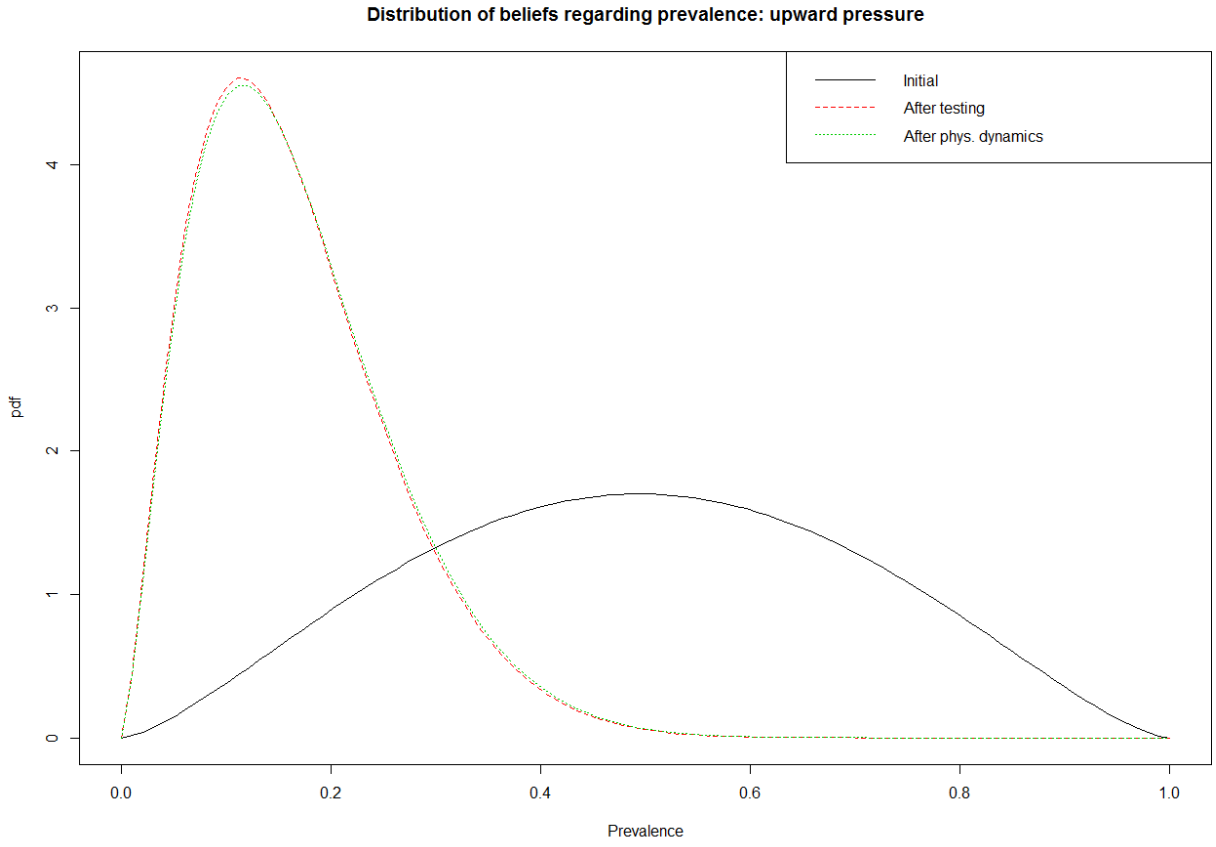

Approximating beliefs by matching mean and variance.

Density projection will ultimately be used to select the posterior parameters. However, the estimation procedure requires the use of initial guesses. By matching the mean and variance of the beta

distribution, we are provide an incorrect but reasonable initial guesses for the values of μ and C[5]. The method used to identify these values is detailed below.

Each of the variables in expression (7) is observed *except* $\tilde{p}$. Taking the expectation of Equation (6) (conditional on observing *K* and *M*), transmission dynamics for the mean prevalence are represented by

$$\begin{aligned}\mu' = E\left[p'\right] &= \frac{\left(N-M-a\right)}{N-M'}\left(\left(1+\beta\right)E\left[\tilde{p}\right]-\beta E\left[\tilde{p}^2\right]\right)\\ &= \left(\frac{N-M-a}{N-M'}\right)\left(\left(1+\beta\right)\tilde{\mu}-\beta\left(\frac{\tilde{\mu}\left(1-\tilde{\mu}\right)}{\tilde{C}+1}+\tilde{\mu}^2\right)\right),\\ &= \left(\frac{N-M-a}{N-M'}\right)\left(\frac{\beta\tilde{C}\tilde{\mu}\left(1-\tilde{\mu}\right)}{\tilde{C}+1}+\tilde{\mu}\right)\end{aligned} \tag{10}$$

since $E\left[\tilde{\mathrm{p}}\right]=\tilde{\mu}.$ The expressions in the second and third lines of (10) are equivalent.

To identify dynamics for C', note that we may find the variance of p' in two ways. By setting these equations equal to each other, we may find a expression for C' in terms of known variables:

$$\begin{aligned}V' &= \frac{\mu'(1-\mu')}{C'+1} = V(p')\\ C' &= \frac{\mu'(1-\mu')}{V(p')}-1\end{aligned} \tag{11}$$

The expression for the variance in terms of $\tilde{p}$ is somewhat complicated and expressed below.

$$\begin{aligned}V(p') &= V\left(\left(\frac{N-M-a}{N-M'}\right)\left(\left(1+\beta\right)\tilde{p}-\beta\tilde{p}^2\right)\right)\\ &= E\left[\left(\left(\frac{N-M-a}{N-M'}\right)\left(\left(1+\beta\right)\tilde{p}-\beta\tilde{p}^2\right)\right)^2\right]-E\left[\left(\frac{N-M-a}{N-M'}\right)\left(\left(1+\beta\right)\tilde{p}-\beta\tilde{p}^2\right)\right]^2\\ &= \left(\frac{N-M-a}{N-M'}\right)^2\left(\begin{array}{l}\left(1+\beta\right)^2\left(E\left[\tilde{p}^2\right]-E\left[\tilde{p}\right]^2\right)\\ +\beta^2\left(E\left[\tilde{p}^4\right]-E\left[\tilde{p}^2\right]^2\right)+2\left(\beta^2+\beta\right)\left(E\left[\tilde{p}\right]E\left[\tilde{p}^2\right]-E\left[\tilde{p}^3\right]\right)\end{array}\right)\end{aligned} \tag{12}$$

[5] Matching the mean and variance will minimize KL divergence for a normal distribution but not a beta distribution.

While somewhat difficult to interpret, these expressions are useful in that they capture the simultaneous upward and downward pressures on prevalence imposed by new infections and testing, respectively. Their inability to minimize KL divergence is noted, but less problematic because they are simply used as initial guesses for the approximations detailed in expression (8).

Complete picture

With the dynamics of both the health states and the belief parameters of interest established, we turn next to the specification of incentives and the decision problem of determining the level of testing administered in each period. A graphic that capture both the learning and physical dynamics is expressed below:

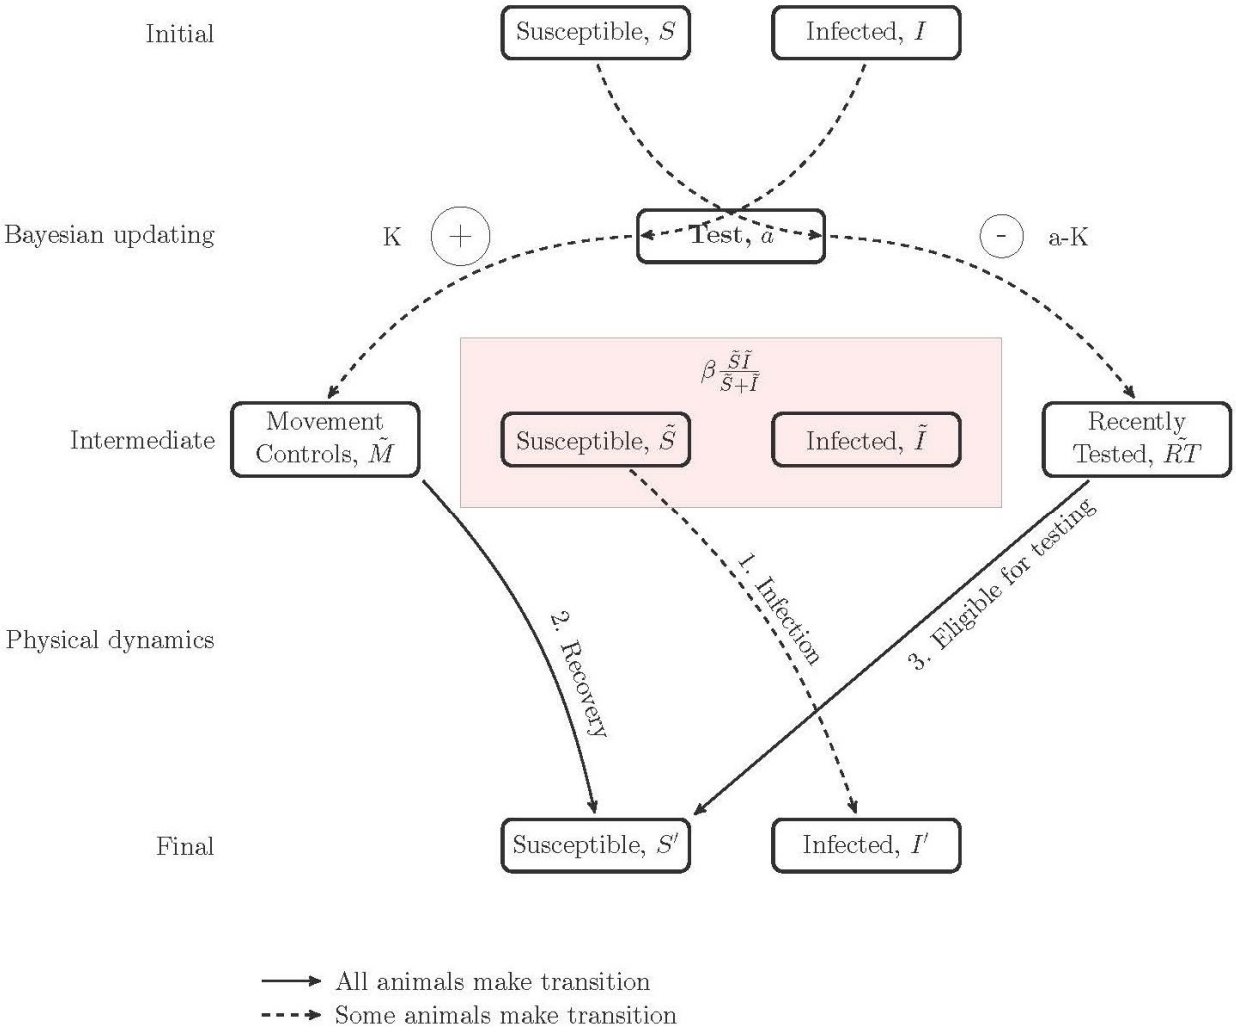

Figure 2. Dynamics for testing, Bayesian belief updating, herd management, transmission and recovery over one period.

Figure 2 details the same order of events expressed in the *Ordering* section, and shows how each of the meta-populations progress through the testing process and physical dynamics.

Farmer payoffs

We consider a social planner who selects a level of testing, a_t, in each period, t, to maximize the present value of the stream of expected profits. This optimal policy will be a function of only the current state of the system, given by the number of herds in movement control and beliefs on prevalence: $X_t = \{M_t, \mu_t, C_t\}$. Further simplifying the problem, we are examining the special case when $\gamma = 1$or herds recover from movement controls instantaneously. Therefore, $M_t = 0 \; \forall \; t$.

We generally follow Gramig and Horan (2011) in formulating the livestock owners' payoff function. First, agents are assumed to be homogeneous except in their health status. Second, we assume that a susceptible herd will generate some level of profit from production independent of trade, π_0, which may be generated by the production of either meat or dairy products. This profit will be deflated by direct production losses experienced when infected animals are present and further losses result when the herd manager is unaware that infected animals are present (i.e. if latently infected). Third, managers will experience gains from trade of magnitude $\pi_1$if they are not under movement controls. Fourth, we have linear testing and cleanup costs, denoted c_1 and c_2 respectively. The aggregate welfare from animal production in a period is thus given by

$$W_t = \pi_0 \left(S_t + \delta\phi I_t + \delta M_t \right) + \pi_1 \left(N - M_t \right) - c_1 a_t - c_2 k_t \tag{13}$$

The parameters δ and ϕ are used to capture losses from infection: δ is the proportion of π_0 enjoyed by a herd which is equal to 1 if healthy and $\delta < 1$ if infected, and ϕ is equal to 1 unless a herd is infected *and* unaware of the infection in which case, $\phi < 1$.

In the case where we have perfect recovery, (refer to equation above) may be rewritten much more simply as

$$W_t = (\pi_0 + \pi_1)(S_t + \tilde{\delta} I_t) - c_1 a_t - c_2 k_t \tag{14}$$

where $\tilde{\delta} = \phi\delta$. In expectation (14) may is rewritten as follows:

$$EW_t = (\pi_0 + \pi_1)\left((1-\mu_t)N + \mu_t \tilde{\delta} N\right) - c_1 a_t - c_2 E\left[k_t \mid a_t, \mu_t, C_t\right] \tag{15}$$

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

$$\begin{aligned} J(X_t) &= \max_{n_t} \left\{ EW_t + \rho E J(X_{t+1}) \right\} \\ s.t. \quad \mu_t &= \left(\frac{N - M - a}{N - M'} \right) \left(\frac{\beta \tilde{C} \tilde{\mu} (1 - \tilde{\mu})}{\tilde{C} + 1} + 1 \right) \\ C_t &= \frac{\mu_t'(1-\mu_t')}{V(p'_t)} - 1 \\ M_t &= (1-\gamma)(M_{t-1} + K_t), \end{aligned}$$

where ρ is an exponential discount factor and $\tilde{\mu}(\mu_{t-1}, C_{t-1} \mid n_{t-1}, K_{t-1})$ and $\tilde{C}_t(\mu_{t-1}, C_{t-1} \mid n_{t-1})$ are given in Equations (4) and (5), respectively. This formulation will be used to guide our choice of an optimal policy, which will be used to construct simulations that produce estimates of the true value of our described strategy.

IV. Methodology: Estimation Approach

This section describes how we incorporate the function forms detailed in the "Theoretical Framework" section into the construction estimates. Due to the complexity of the problem at hand, it will be

necessary to approximate the dynamic value of disease control. Our first step is to create reliable estimates of the optimal testing level given a set of belief and state variables, X. This step requires precise estimation of posterior parameters given prior parameters, actions and outcomes. With these estimates we are able to deduce the optimal testing level using a reverse induction methodology: value function iteration.

Given these optimal choices, we will be able to forward simulate decisions and predict the true value trajectories for given scenarios in which the initial disease prevalence and beliefs regarding prevalence are allowed to vary. The specific approach chosen for each of these steps is detailed in the following subsections.

All of our estimation is conducted in MatLab.

Expected payoffs

Real payoffs based on the actual underlying prevalence will be used to compare the desirability of alternative policies, while expected payoffs are used to determine the level of testing chosen in a given period. The AHB will be unaware of real payoffs ex-ante, and will only have partial information regarding these profits ex-post. The AHB will not be able to attribute all variation in this ex-post profits to changes in health and trade status and thereby identify the true prevalence within the system. Health status, and production in turn, may also vary with the level of other diseases within the system, environmental conditions, and production choices[6]. Choices will become more efficient as information is collected over time, but will not maximize the stream of real payoffs.

We estimate the expected payoffs simply by assigning the payoffs associated with each level of prevalence with the appropriate probability weight and estimating the number of identified latently

[6] While we treat farms as homogeneous, there are many factors that make these firms distinct. For example, the genetic composition of the herd may vary based on differences between the derived demand for fluid milk and milk used for other dairy products.

infected facilities given the number of tests executed. This is a straightforward application of the formula given in (15).

Calculate Bayesian updating

Once the number of outcomes has been determined within our simulation, the parameters μ and C may be updated in the way described in equation (4) and (5). The calculations are simple and approximations are unnecessary.

Density Projection

To match the distribution after accounting for physical dynamics, we attempt to satisfy the equations shown in (8). The first step is estimate (9), which complicated by the divergence of these values to negative infinity when the expectation of p' becomes very close to either 0 or 1. We use an enhanced version of the function qnwbeta, which is provided in the CompEcon Toolbox, to construct an approximation for these values.

The estimates of (9) are passed into MatLab's fsolve program to obtain posterior estimates of μ' and C'. More specifically, μ' and C' are estimated to satisfy the conditions defined are (8). Here, we are forced to restrict the possible values of μ to an interval that approximates a $(0,1)$ interval because values of 0 or 1 cause one of the moment conditions to take on a value of negative infinity. Here, we choose an arbitrarily small value of 1e-10.

Value function Iteration

To find the optimal level of testing given an initial set of beliefs, we use a value function iteration. This procedure relies on backward induction to deduce the optimal (testing) policy in the current period by accounting for its effect on the welfare in future periods. Such models are often employed in marcoeconomic theory. Our approach will follow those analyses targeted at describing the relationship between stochastic growth and consumption (Christiano, 1990).

Forward simulation

With results obtained using the procedures detailed in the previous sections, we are able to forward simulate our optimal policy. To do so, we must specify an initial, underlying prevalence, initial beliefs, and a reasonable time horizon. For this analysis we assume that the underlying beliefs are uniform and that allow the underlying prevalence to vary over the unit interval. We explore a time horizon of 5 years.

To avoid the potential problems that may arise when prevalence lies outside the unit interval, we model the number of infected individuals rather than the prevalence. From this group we conduct hypergeometric sampling (i.e. sampling without replacement), excluding potential cases where the number of identified facilities exceeds the number of latently infected facilities. In the theoretical section we assumed that the sampling process was binomial, which is a reasonable approximation of hypergeometic sampling when the total number of herds is large and the number of tests is small.

V. Results

Policy function

The first steps in our approach culminate in the construction of a policy function. For the system in which the AHB can choose to test between 0 and 10 facilities per month, the optimal testing levels for any combination of $\mu \in (0,1)$ and $C \in [2,61]$ are represented in the heat map below. In Figure 4:

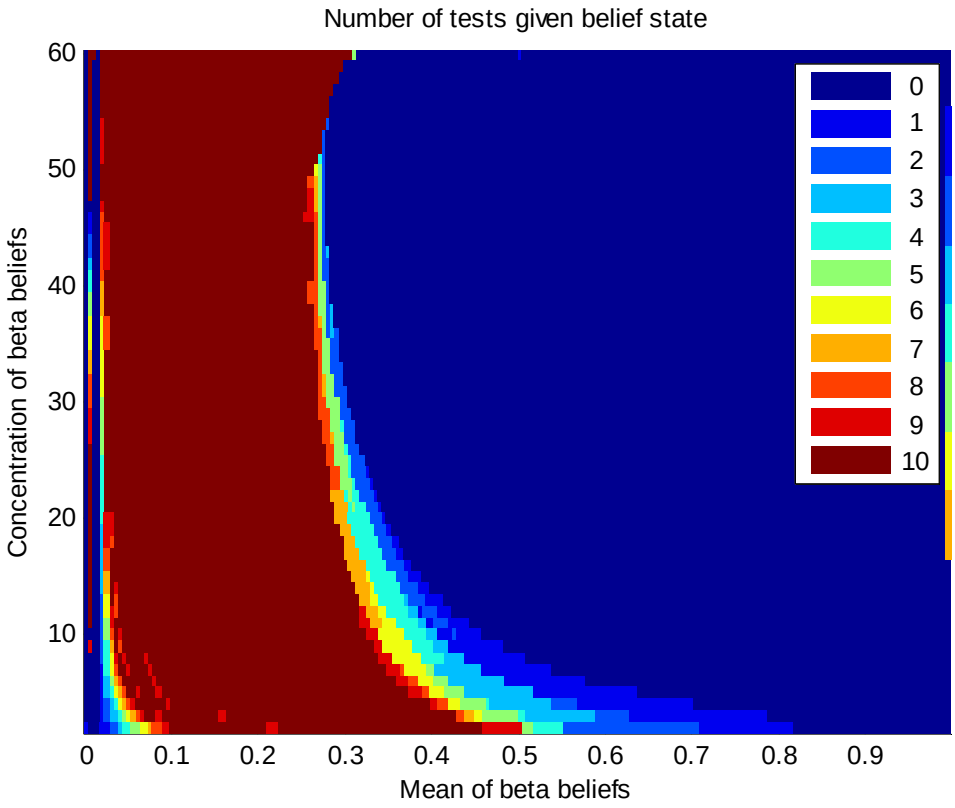

Figure 4. Level of testing given belief state when N=5,990.

Here, we can clearly see that when the decision maker believes that the prevalence is very high (e.g. more than 80% of herd on thought to be infected), the testing level is zero. This relationship between testing in beliefs is likely attributable to the fact that any attempt to decrease prevalence through testing will only lead to an increase in the number of infections. Our specified density dependent transmission function will cause simulated infections to be highest when $p = 0.5$. This result may change if the number of tests is allowed to increase or the state space is increased because the value function iteration step accounts for expectations of changes in the parameters[7].

While we do not increase the maximum value of C in this report, we are able to make inferences regarding the effects of increasing the maximum fraction of herds tested. In the Figure 5, the number of

[7] Currently we are placing severe restrictions on the value of C. This parameter could take on very large values, particularly if the number of tests in a given period is particularly high relative and the effect of physical changes on C is relatively small.

herds is reduced from 5,990 to 100 to represent testing over a much smaller group of interacting facilities. We can see that there is change in the absolute number of tests for any (μ, C) pair, but that the qualitative shape is the same. Testing rates generally decline as the mean of the belief distribution increases, and the concentration of this distribution increases.

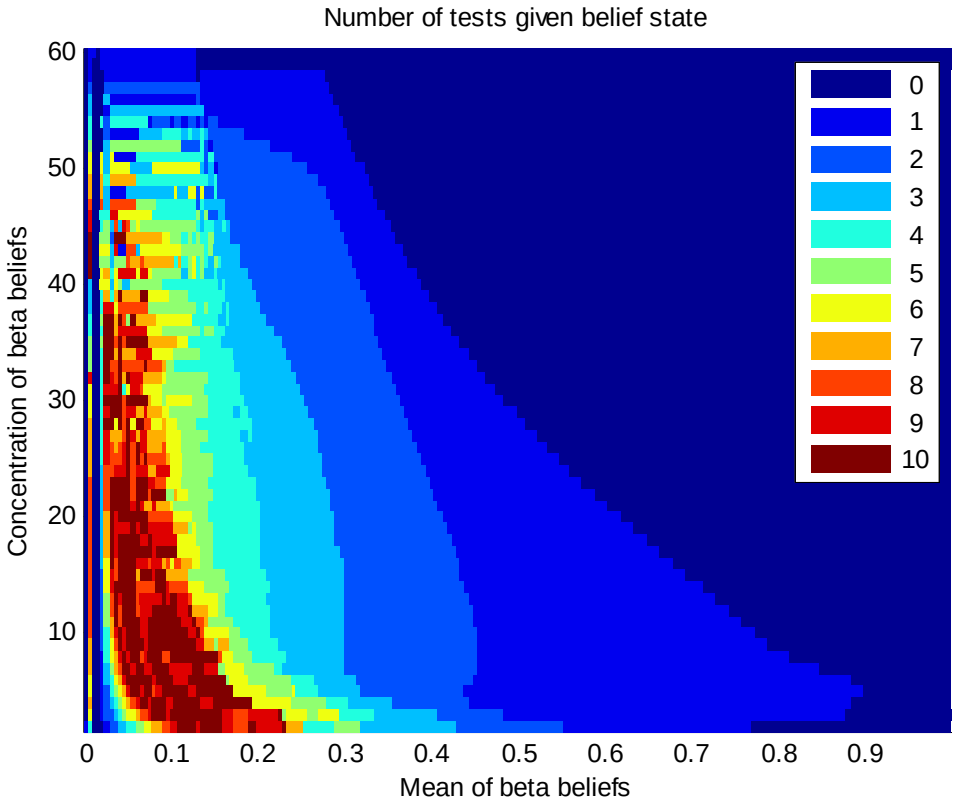

Figure 5. Level of testing given belief state when N=100.

Progression through time

The previous section provided both a policy function and functions to map the parameters that define this periods beliefs into the next. With this information we are able to compare to important scenarios. In the first scenario, a central veterinary authority uses all information available. In the second scenario, the data only informs Bayesian updating. Implicit in this approach is the assumption that the central veterinary authority beliefs that, within a given time horizon, there exist a constant, unknown

prevalence. In the third scenario, the central veterinary authority selects a level of testing and does not change the testing rate in response to test results.

Our main measure of concern is the actual level of profits realized in a given period of 5 years, which is identical to (14) except that the μ is replaced by p. We compare the cumulative stream of discounted value realized by the producers for the three scenarios discussed previously. We allow the actual, underlying probability to take value along the interval $[0.03,0.15]$ and report cumulative value measures for each. These are include in Table 1:

Bayesian and Comp	0		
Bayesian only	-0.12%	5 tests	0.09%
0 tests/mo	-0.12%	6 tests	0.13%
1 tests	-0.08%	7 test	0.17%
2 tests	-0.04%	8 test	0.21%
3 tests	<0.01%	9 tests	0.25%
4 tests	0.05%	10 tests	0.29%

Table 1. Percent differences from the value realized when both Bayesian updating and compartmental dynamics are accounted for.

Here, we can see that the including compartmental dynamics clearly outperforms a simply using Bayesian updating. Additionally, it outperforms naïve guesses that also happen to be bad. In this case, the AHB is better off at higher testing rates. Therefore, if the AHB happens to take an initial guess that is high, its welfare will be greater.

The difference between the streams of value between the models in which compartmental dynamics are included versus not increases as the time horizon becomes longer. For example, for a 10 year time horizon, the percent difference between the two models is 0.4% of total value. This relationship is unsurprising because of small transmission coefficient associated with bTB.

VI. Discussion

Our paper is a first attempt at addressing the problem of disease testing and control when the underlying prevalence is uncertain to the central veterinary authority. Our simulations suggest that accounting for both the physical dynamics will lead to a higher level of social welfare than the cases where the central veterinary authority incorrectly assumes that it is learning about a fixed prevalence or when a level of testing is arbitrarily chosen.

Enhanced realism will be an essential step in future iterations of our modeling approach. Our model lacks some elements such as endogenous responses in trade volumes and biosecurity investment that are essential to accurate estimates. It does already, however, provide evidence that suggests the need to explicitly model the relationships between disease testing and beliefs regarding disease.

This paper makes two contributions. First, we provide an extension to the existing POMDP literature by allowing for a system in which the decision maker learns about a changing underlying probability. Second, the results suggest that data collection *and* rigorous analyses of that data may provide more efficient broad-based controls than otherwise delivered. Such an enhancement in efficiency may not alone constitute sufficient motivation to institute tracing programs, but rather an additional benefit which is currently left unconsidered.

References

1. Bicknell, K.B., Wilen, J.E. and Howitt, R.E. (1999). “Public policy and private incentives for livestock disease control.” *The Australian Journal of Agricultural and Resource Economics* 43, 501–521.

2. Christiano, Lawrence J. "Solving the stochastic growth model by linear-quadratic approximation and by value-function iteration." *Journal of Business & Economic Statistics* 8.1 (1990): 23-26.

3. Cousins DV, Florisson N. A review of tests available for use in the diagnosis of tuberculosis in non-bovine. http://www.cfsph.iastate.edu/Factsheets/pdfs/bovine_tuberculosis.pdf

4. Cryer, J. Sampling With and Without Replacement and the Hypergeometric and Bionomial Distributions. Accessed 2014. http://homepage.divms.uiowa.edu/~jcryer/Hypergeometric.pdf

5. Fackler, Paul L., and Robert G. Haight. “Monitoring as a Partially Observable Decision Problem.” *Resource and Energy Economics* (2014).

6. Gelman, Andrew, et al. *Bayesian data analysis*. CRC press, 2013.

7. Gramig, Benjamin M., and Richard D. Horan. "Jointly determined livestock disease dynamics and decentralised economic behaviour." *Australian Journal of Agricultural and Resource Economics* 55.3 (2011): 393-410.

8. California Department of Food and Agriculturee (CDFA). Fact Sheet: TB Testing Information for Cattle Producers http://www.cdfa.ca.gov/AHFSS/Animal_Health/pdfs/TB/TB_Testing_Fact_Sheet.pdf.

9. Vallat, Bernard. “Feeding the world better by controlling animal diseases.” OIE Website, Editorials. Accessed May 21st, 2014. Available at http://www.oie.int/for-the-media/editorials/detail/article/feeding-the-world-better-by-controlling-animal-diseases/

10. D. L. Whipple, C. A. Bolin, A. J. Davis, J. L. Jarnagin, D. C. Johnson, R. S. Nabors, J. B. Payeur, D. A. Saari, A. J. Wilson, M. M. Wolf. “Comparison of the sensitivity of the caudal fold skin test and a commercial gamma-interferon assay for diagnosis of bovine tuberculosis.” Am J Vet Res. 1995 April; 56(4): 415–419.