When to monitor or control; informed invasive species management using a partially observable Markov decision process (POMDP) framework

Thomas K. Waring[1,2], **Vera L. J. Somers**[1,3], **Michael A. McCarthy**[4], **Christopher M. Baker**[1,2,5*]

[1]School of Mathematics and Statistics, University of Melbourne, Victoria, Australia

[2]Melbourne Centre for Data Science (MCDS), University of Melbourne, Victoria, Australia

[3]ARC Training Centre in Optimisation Technologies, Integrated Methodologies, and Applications (OPTIMA), Victoria, Australia

[4]School of Agriculture, Food and Ecosystem Sciences, University of Melbourne, Victoria, Australia

[5]Centre of Excellence for Biosecurity Risk Analysis, School of BioSciences, University of Melbourne, Melbourne, Victoria, Australia

June 29, 2024

Abstract

1: Resource allocation for invasive species management requires information about the size of the invasive population, which may be expensive and time-consuming to obtain. The trade-off between investment in monitoring and control efforts is a challenging decision problem, and existing mathematical tools are often difficult to interpret, and / or limited to a specific case study.

2: We propose a partially observable Markov decision process (POMDP) framework to help decision makers understand effective monitoring and control policy making. POMDPs can deal with uncertainty in both the model and state of the system, but are more challenging to solve due to

*__Corresponding author.__ *E-mail address:* cbaker1@unimelb.edu.au

the continuous and high-dimensional state space. Rather than limiting the possible states of the system, as do most previously proposed methods, we work through the development of a *density projection* approach, which reduces the dimensionality of the space of beliefs by restricting them to a parametrised family of probability distributions. This serves to align the mathematical representation of the problem with the real-world quantities relevant to human decision making.

3: The result of our model is a sequence of actions which minimises the expected cost incurred in managing the invasive species, where the recommendation depends on an estimate of the species' abundance, and the uncertainty in this estimate. We demonstrate the effectiveness of our proposed framework with a case study on tropical fire ant (*Solenopsis geminata*) control and two generic case studies of varying complexity. Furthermore, we investigate sensitivity of the results to the choices of control cost and efficacy, and monitoring cost and error.

4: The framework proposed by this paper makes the powerful machinery of POMDPs available to environmental managers. It computes the optimal course of action to manage a growing population of an invasive species, incorporating a varying time horizon and multiple control interventions. We sidestep the computational difficulties of general POMDPs to provide a clear, visual overview of decision-making recommendations, and how these decisions change in new situations. Initial results and scenario based analysis show promising results, and the framework could be extended to the related field of disease management.

Keywords: *invasive species, decision making, partially observable Markov decision processes, uncertainty*

---

1. Introduction

An uncontrolled invasive species can have a major impact on the economy and environment (Hoffmann and Broadhurst, 2016; Alvarez and Solís, 2018). Therefore, if damage is to be avoided, conservation authorities must allocate resources effectively to control or eradicate invasive populations. However, making informed decisions relies on knowledge about the severity of the infestation, and obtaining such information can be time-consuming and expensive. Therefore, the decision problem facing managers involves a challenging combination of both monitoring (information gathering) and control efforts. More specifically, two main challenges arise with invasive species management. First, monitoring and control efforts have an influence on one another and can therefore, not easily be determined separately. And second, full knowledge of the system and the dynamics of its evolution is often unknown or unavailable (Büyüktahtakın and Haight, 2018).

Approaches to the problem vary widely: the underlying model might be deterministic or stochastic, the solution method exact or heuristic, and the potential aims include decision making and policy evaluation. For example: Epanchin-Niell et al. (2012) use mathematical programming, optimal control methods appear in Rout et al. (2014); Mehta et al. (2007); Hauser and McCarthy (2009) and Mbah and Gilligan (2010), and heuristic genetic algorithms are used by Carrasco et al. (2010).

If the problem is modelled as a stochastic, sequential decision problem, a natural mathematical framework is that of a Markov Decision Process (MDP). MDPs are conventionally solved by Stochastic Dynamic Programming (SDP) (Bellman, 2003), which computes a plan of action in the present which is optimal for the uncertain future — examples in the management of invasive species include Williams and Brown (2022); Marescot et al. (2013); Hyytiäinen et al. (2013); Polasky (2010); Rout et al. (2011) and Moore et al. (2010). A shortcoming of this approach, however, is the fact that MDPs assume perfect knowledge of the present state of the system.

If instead the manager must make their decision based only on a (possibly imperfect) observation of the environment, the proper extension to our MDP is a partially observable Markov decision process (POMDP) (Monahan, 1982; Littman, 2009). POMDPs are, however, far more difficult to solve than MDPs, and this difficulty becomes more pronounced as the state space of the system grows. Therefore, POMDP approaches to invasive species management tend to model the system as occupying one of a small number of levels of infestation (Chadès et al., 2008; McDonald-Madden et al., 2011; Haight and Polasky, 2010; Regan et al., 2011; Chadès et al., 2011; Rout et al., 2014). In addition, the realism of their models often relies on the specifics of the situation at hand (Williams and Brown, 2022, Section 10) — for example (Regan et al., 2011) model a situation where a paddock may contain no weeds, only seeds, or seeds and adult weeds: such an approach does not translate directly to, say, animal invasions. Furthermore, as discussed by Williams and Brown (2022), representing the results and recommendations of a POMDP model in an intuitive and tractable way remains a challenge (see Regan et al. (2011) for one approach).

Arguably the most natural way to represent the state of a pest invasion is as an unbounded and possibly continuous population variable: the state of the system is the number or amount present of the invasive species in question. However, as indicated by Williams and Brown (2022), continuous state space POMDPs are particularly mathematically challenging, and so far have found little application in the ecology literature.

In this work, we address this issue by working through the construction of a POMDP model based on *density projection* (Williams and Brown, 2022, Section 8). In our model, the uncertain state of the ecological system is represented by a continuous probability distribution, parametrised by a small number of meaningful quantities.

By way of some simplifying assumptions, we develop an instance of the density projection method which admits an analytic solution. This allows us to compute optimal policies for a wide range of states and parameter values, and we investigate the boundaries where the optimal action changes. We represent our results in an intuitive visual way, including optimal policies, the cut-off points where the optimal policy changes, and the sensitivity of our results to parameter values.

Our method sidesteps the computational difficulties of POMDPs while maintaining a generic setup: as such, we hope it will demonstrate the power of POMDP methods and provide an accessible starting point for ecologists incorporating partial observation into their models.

2. Materials and Methods

We begin this section by setting up the problem we aim to solve, and by motivating the introduction of POMDPs and the density projection approach. Our precise mathematical formulation is deferred to Section 2.3.1, after which we derive an analytic solution to a particular instance of this setup.

2.1. Problem statement

We consider the problem of surveillance (i.e., monitoring) and control of an invasive species population with some unknown abundance, and which is growing at a (possibly) uncertain rate. In response a decision maker or manager has some control actions available, each of which eliminates some fraction of the species, at some cost. Alternatively, the manager may opt to monitor the species, that is, to spend some resources to receive some estimate of the species' abundance.

We state this problem as a *Partially Observable Markov Decision Process* (POMDP) (Åström, 1965; Monahan, 1982). A POMDP is specified by sets of states, actions and observations, an objective function to be minimised (or maximised), and the probabilities which govern the state transitions and the observations.

The case at hand runs as follows. The state at time t is the abundance N_t of the species, and we assume each unit of abundance incurs one unit of cost (that is, cost is measured in units of species abundance). The transitions are governed by a model for the growth of the species. For simplicity, we restrict ourselves to the case of exponential growth, so that

$$N_{t+1} = e^{r_t} \cdot N_t,$$

where the growth rate r_t may vary in time. To model an uncertain estimate of the value r, we draw each r_t from a normal distribution with mean r and variance Δr^2 (Keeling and Rohani, 2008, Section 6.2).

The actions available to the manager fall into two categories: *control* and surveillance, i.e., *monitoring*. Control actions are denoted a_i for $i = 1, 2, \ldots$, and have respective efficacies ρ_i and costs c_i. The effect of action a_i, taken at time t, is modelled as reducing the next-step abundance by the fraction ρ_i, so we have

$$N_{t+1} = (1 - \rho_i)e^{r_t} \cdot N_t, \tag{1}$$

with immediate cost $N_t + c_i$. For simplicity, we make the assumption that, in this case, the manager receives no observation of the state. This models, for example, a program of laying baits, as in our case study (described in Section 4).

If the manager were certain of the species' abundance, the decision is between doing nothing (if the abundance N_t is small) and controlling the species (if N_t is large). The threshold abundance where the decision changes is determined by the cost and efficacy of control actions, and how quickly the species' population is growing — this calculation is carried out at the beginning of Section 2.3.2.

Monitoring actions, denoted A_i, allow the manager to observe the state at some cost C_i. We model this as a draw from a lognormal distribution around the true value, with some standard error e_i. In a fully general circumstances, monitoring actions might also have some control effects.

2.2. Belief Markov Decision Process

The conventional approach to solving a POMDP is to consider it as a conventional MDP, where each state is a probability distribution over the original states of the system (Åström, 1965). Therefore, the new 'belief state' space has one dimension for each original state, representing the probability that the system is in the given state, given the past history of actions and observations. Since higher dimensions are computationally more difficult, the belief state approach is forced to work with a small number of states (Chadès et al., 2021; Rout et al., 2014; Chadès et al., 2008), and certainly to avoid an unbounded state space.

Beyond the limits imposed by computational concerns, the results of the conventional belief state MDP can be difficult to interpret and communicate. For example, in a model which allows for low, medium and high abundance, the following unlikely belief is a point in the considered state space:

$$P(N = \text{low}) = 0.5$$
$$P(N = \text{medium}) = 0$$
$$P(N = \text{high}) = 0.5.$$

Beliefs such as these would only occur in very specific circumstances, and as such including such examples in the state space complicates any interpretation of the model. In (Rout et al., 2014), for example, this problem is avoided by only giving scenarios for distributions where $P(N = \text{high}) = 0$, at the cost of a loss in generality. The interpretability problem is exacerbated by the difficulty of representing the value of a given action or state in more than three dimensions, as discussed in (Chadès et al., 2021, Section 5.2).

It is clear, then, that it is valuable from a computational point of view to reduce the dimensionality of our state space. Furthermore, these dimensions should reflect quantities that have real-world meaning, in order to better reflect the actual 'belief state' of a manager.

2.3. Proposed solution

We propose that the 'belief state' of a manager should be restricted to a parametrised family of probability distributions (Zhou et al., 2010). Reflecting the goals stated above, we choose as parameters the abundance of the controlled species, and the uncertainty in this value. For our development, we choose a lognormal distribution, parametrised by $n, \Delta n$ such that $\log N$ follows a normal distribution with mean n and standard deviation Δn. This choice is made for convenience: exponential growth and multiplicative control correspond to translations of the distribution. Such a simplification fulfills the desiderata of the previous section, as demonstrated in Section 3. In short, given an estimate with error of $\log N$ our model computes the optimum policy amongst options of monitoring or control.

2.3.1. Problem setup

We now re-state our decision problem with our suggested simplification. As is conventional for POMDPs, the model is specified by a tuple consisting of: state space, observation space, transition function, observation function, reward function, time horizon and initial belief (Chadès et al., 2021). In our case the state space is simply the abundance $N \in \mathbb{R}_{\geq 0}$, and observations lie in the set $\mathbb{R}_{\geq 0} \cup \{\text{null}\}$. The transition function is the growth model specified in Equation (1): that is, exponential growth with rate e^{r_t} and multiplicative control. The observation is 'null' if the action is not a monitoring action A_i, otherwise it is lognormally distributed around the true abundance with (log-scale) standard error e_i. The reward function is the abundance of the species, plus the cost of any action taken, and the time horizon, discount factor γ, and lognormal initial belief are specified by the user. A full account of the necessary parameters is given in Table 1.

To solve this POMDP, we reduce it to a parametrised belief-state MDP as follows. The states are $(n, \Delta n) \in \mathbb{R} \times \mathbb{R}_{\geq 0}$, corresponding to normally-distributed log-abundance with mean n and standard deviation Δn. The initial values and time horizon are set by the user, and the reward is now cost plus *expected abundance*, given by:

$$\mathbb{E}[N] = \exp(n + (\Delta n)^2/2).$$

We define the transition function in two cases, depending on whether the manager takes a control or a monitoring action. Since growth and control are multiplicative (as in as in Equation (1)), under a control action a_i the estimate n of $\log N$ is translated, viz:

$$n_{t+1} = n_t + r + \log(1 - \rho_i).$$

Accounting for uncertainty in r, we have $\Delta n_{t+1} = \Delta n_t + \Delta r$.

In the case of a monitoring action, the new abundance estimate n_{t+1} is set to the result of the observation made. Since the true abundance is not available, we draw the observation n_{t+1} from the belief distribution $\text{Norm}(n + r, \Delta n)$. If the monitoring error is e, then the new error estimate is given by a Bayesian update formula, with normal likelihood and normal prior (Gelman et al., 2013, Section 2.2), with a term accounting for the uncertain growth rate:

$$\Delta n_{t+1} = \frac{\Delta n \cdot e}{\Delta n + e} + \Delta r. \tag{2}$$

In the case that the monitoring action has some control effect ρ, n_{t+1} would be drawn instead from $\text{Norm}(n + r + \log(1 - \rho), \Delta n)$.

2.3.2. Solution: action value functions

Following a standard method for solving Markov Decision Problems (Bellman, 2003), we aim to calculate the *action-value function* $Q_T(n, \Delta n \mid a)$, defined as the minimal (expected) cost of taking action a in state $(n, \Delta n)$, over some finite time horizon T. To do so, we follow Bellman's principle of optimality, which reads:

$$Q_T(n, \Delta n \mid a) = \text{cost}(n, \Delta n, a) + \gamma \cdot \mathbb{E}\big[\min_{a'} Q_{T-1}(n', \Delta n' \mid a')\big]. \tag{3}$$

Here, the expectation is over the subsequent state $n', \Delta n'$, which (depending on the action a) may be uncertain, and γ is a *discount factor*, which de-emphasises costs incurred further in the future. (Note that the subsequent state has subscript $T-1$, being one step *closer* to the chosen horizon.) Having calculated this value for each action, the best action in any given state is simply the one with the smallest cost function.

For $T = 0$ the second term in Equation (3) is zero, so that $Q_0(n, \Delta n \mid a) = \exp(n+(\Delta n)^2/2)+\text{cost}(a)$. For $T > 0$, we begin by evaluating the expected cost incurred if only control actions are taken. Then, while the evolution of the system is still random, the manager's belief evolves deterministically — with abundance growing or shrinking accordingwhich control actions are taken. As such, rather than evaluating the 'min' in Equation (3) directly, we compute the value of a *sequence* $\mathbf{a}$ of actions, with length T, by the same process (noting that the optimal action with $T = 0$ will always be to do nothing). We claim that, in this case, the expected cost is always a linear function of $N = \exp(n)$, with coefficients (i.e., slope and intercept) depending on Δn. This proceeds by induction on T, the length of the action sequence $\mathbf{a}$. If $T = 0$, then as above

$$Q_0(n, \Delta n \mid -) = \exp(n + (\Delta n)^2/2) = e^{(\Delta n)^2/2} N.$$

Then, suppose that we have computed the slope A and intercept B associated to δn and some sequence of actions $\mathbf{a}$, viz:

$$Q_T(n, \Delta n \mid \mathbf{a}) = A_T(\Delta n, \mathbf{a}) \cdot N + B_T(\mathbf{a}). \tag{4}$$

To extend this expression to length $T+1$, we add an extra action a_i. Then, the value $Q_{T+1}(n, \Delta n \mid a_i\mathbf{a})$

consists of the immediate reward, plus the discounted reward for the rest of the sequence, evaluated at the subsequent state. We compute:

$$\begin{aligned}
Q_{T+1}(n, \Delta n \mid a_i \mathbf{a}) &= \exp(n + (\Delta n)^2/2) + c_i + \gamma Q_T(n + r + \log(1-\rho_i), \Delta n + \Delta r \mid \mathbf{a}) \\
&= \exp(n + (\Delta n)^2/2) + c_i + \gamma \left[A_T(\Delta n + \Delta r, \mathbf{a}) \cdot N + B_T(\mathbf{a}) \right] \\
&= \left[e^{(\Delta n)^2/2} + \gamma(1-\rho_i) e^r A_T(\Delta n + \Delta r, \mathbf{a}) \right] \cdot N + [c_i + \gamma B_T(\mathbf{a})].
\end{aligned}$$

For a fixed Δn, T and list of M possible actions, the full library of these coefficients can be calculated recursively in $O(MT)$ time. For fixed Δn and T, the *state-value function* $V_T(n, \Delta n) = \min_{\mathbf{a}} Q_T(n, \Delta n \mid \mathbf{a})$ is piecewise-linear and convex (in N), making it again straightforward to compute recursively.

To add in monitoring actions, we make one further simplifying assumption: that in any sequence of actions, monitoring only occurs once. That is, the error e in monitoring and rate Δr at which uncertainty grows are small enough that monitoring once ensures it is not necessary again. The density projection method does not rely on this assumption, but the analytic solution we present does. In general, the analogous form of the integrals presented below could be evaluated numerically to provide similar results. With this assumption however, the term inside the expectation of Equation (3) is a piecewise linear function of e^n, which makes it analytically calculable.

Substituting the state value function into Equation (3), we have:

$$Q_T(n, \Delta n \mid \text{monitor}) = \exp(n + (\Delta n)^2/2) + \text{cost}(\text{monitor}) + \gamma \int_{-\infty}^{\infty} V_{T-1}(x + r, \Delta n') \varphi_{n,\Delta n}(x) dx,$$

where $\varphi_{n,\Delta n}(x)$ is the probability distribution function for $x \sim \text{Norm}(n, \Delta n)$. The integration interval splits into pieces where the minimum in the definition of V_{T-1} is achieved by a particular sequence $\mathbf{a}$. For such an interval $l \le x \le u$, the function V_{T-1} takes the form of Equation (4), and as such can be calculated using the following identities (Jawitz, 2004, Section 2.2):

$$\int_l^u \varphi_{n,\Delta n}(x) dx = \frac{1}{2} \left(\text{erf}\left(\frac{u-n}{\Delta n \sqrt{2}} \right) - \text{erf}\left(\frac{l-n}{\Delta n \sqrt{2}} \right) \right)$$

Notation	Variable name in code	Definition
r	`r`	Disease growth rate
Δr	`dr`	Uncertainty in growth rate
ρ_i	`rho`	Control efficacy
c_i	`ce`	Control cost
—	`interventions`	A dictionary containing pairs `(ce,rho)` of cost & efficacy
c_m	`cm`	Cost of monitoring
e	`err_mon`	Error in monitoring
—	`mon_dict`	A dictionary of triples `(cm,err_mon,rho_mon)` of cost, error & control efficacy
T	`num_steps`	Decision horizon
γ	`gamma`	Discount factor

Table 1: Model parameters.

and

$$\int_l^u e^x \varphi_{n,\Delta n}(x)dx = \frac{\exp(n+\Delta n^2/2)}{2}\left(\operatorname{erf}\left(\frac{u-n-\Delta n^2}{\Delta n\sqrt{2}}\right) - \operatorname{erf}\left(\frac{l-n-\Delta n^2}{\Delta n\sqrt{2}}\right)\right).$$

These permit the calculation of the value of monitoring (noting that we can no longer predict what future actions will be, since the outcome of monitoring is stochastic):

$$Q_T(n, \Delta n \mid \text{monitor}) = \exp(n + (\Delta n)^2/2) + \text{cost}(\text{monitor}) + \gamma \int_{-\infty}^{\infty} V_{T-1}(x+r, \Delta n')\varphi_{n,\Delta n}(x)dx, \quad (5)$$

where $\Delta n'$ is given by Equation (2).

For (control) actions *before* a monitoring step, the expression Equation (5) is substituted into Equation (3). For instance, the value of the sequence a_i (control) then monitor is given by:

$$Q_{T+1}(n, \Delta n \mid a_i\,\text{monitor}) = \exp(n + (\Delta n)^2/2) + c_i + \gamma Q_T(n + r + \log(1-\rho_i), \Delta n + \Delta r \mid \text{monitor}).$$

2.4. Model Parameters

The parameters specifying the model are collected in Table 1. Fixing these based on the scenario in question, the optimum action depends on the abundance of the species, and the uncertainty in that value. If we denote the abundance by N, then given an estimate n of $\log N$ with approximate error Δn, our model computes the optimum action. In fact, the model computes the expected reward for a *sequence* of actions which ends either at the time horizon T, or with some monitoring action. We provide code to compute these values, and a vignette explaining how to set parameters.

3. Results

To illustrate our results, we plot the optimum sequence of actions for pairs $(\Delta n, \overline{N})$ of uncertainty and expected abundance. For the following examples, we take $r = 1.2 \pm 0.05$ and $\gamma = 0.9$. Control is assumed to cost $c_e = 4$ units with effectiveness $\rho = 0.6$, and monitoring costs $c_m = 1$ unit, with error $e = 0.1$.

The simplest interesting case is over $T = 2$ time steps: at the first step a decision must be made between suppressing the outbreak (a control intervention), collecting information to inform the second decision (a monitoring intervention), or doing nothing. The outcome of this analysis is given in Figure 1.

Qualitatively, the displayed results match intuition. For small values of the uncertainty (in this case $\Delta n < 0.5$), monitoring is never optimal, as enough is known about the problem already. In this case, the decision between control and doing nothing is made based on the expected future abundance of the species. As uncertainty increases (rightwards on the graph), monitoring actions become more valuable, and as such optimal over a larger range of parameter values, matching the orange region on the right of the graph. Furthermore, the marginal benefit of monitoring is greater when the decision is less "obvious" — that is, when the estimated abundance is not very high or very low. In these extreme cases, even with a large "factual uncertainty" in the prevalence of the species, the "decision uncertainty" is small, so monitoring interventions are not worthwhile.

The simple case of Figure 1 illustrates the qualitative success of our proposed method. For practical applications, however, it could be useful to implement more complexity. In Figure 2 we illustrate some of the ways this is possible. To clarify the entries in the legend of Figure 2a, that is, the actions available to the manager, we include a tree representation of the same, in Figure 2b.

Figure 2 has distinct regions corresponding to '`monitor`' and '`control, monitor`', which arise as we have added one extra time step ($T = 3$). This is informed by the fact that uncertainty increases over time, but also by the benefits of suppressing an exponentially growing outbreak earlier rather than later.

In addition, in Figure 2 we have added a second control action. Specifically, this is modelled as the same intervention performed twice, with independent results, with a cost of $2c_e$ and an efficacy of $1-(1-\rho)^2$. As the abundance increases (upwards on the plot), these more intense control actions become optimal.

In Figure 3, we test the effect that variation of the parameters c_e and ρ has on the optimal control actions — ignoring for the moment any uncertainty, and therefore any monitoring actions. For the

Note that, with different parameter values, the action `ignore,monitor` appears on the plot.

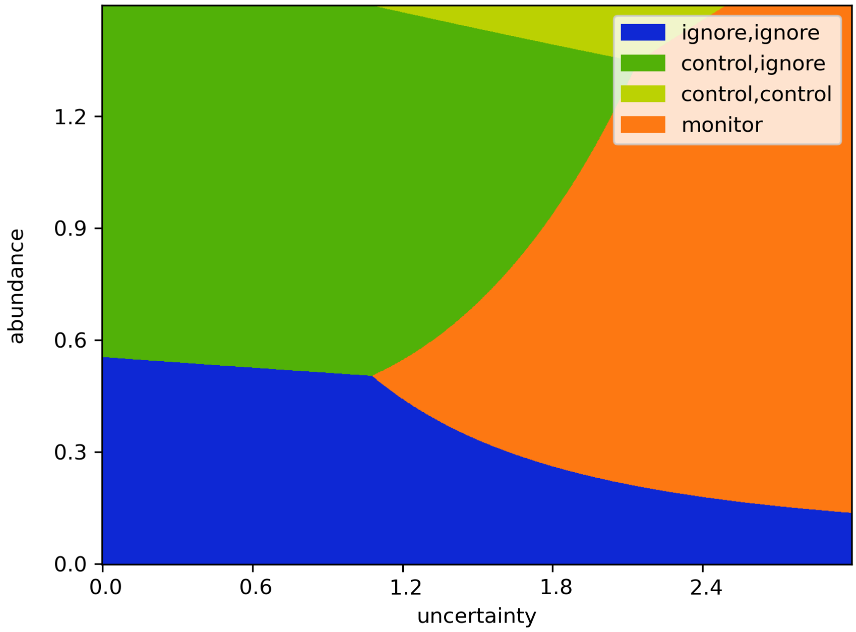

Figure 1: Optimal management actions for a sequence of $T = 2$ time steps for given uncertainty and abundance. Management options consist of 'ignore', 'monitor' or 'control'.

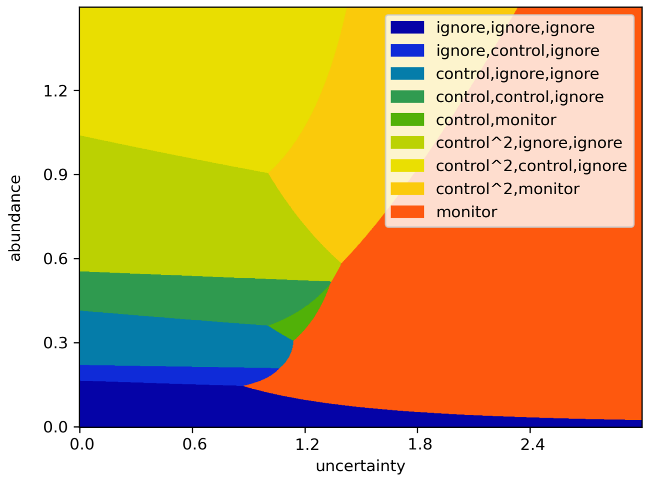

(a) Optimal management actions over three time steps

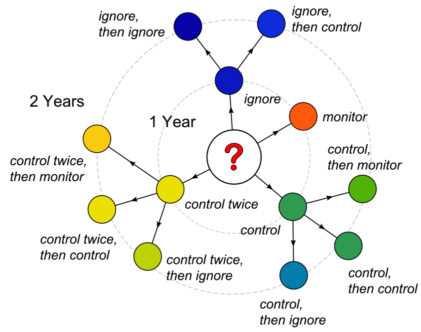

(b) A tree, representing the choices available to the manager.

Figure 2: Optimal management actions for further complexity, here $T = 3$, illustrating the quantitative power of our model.

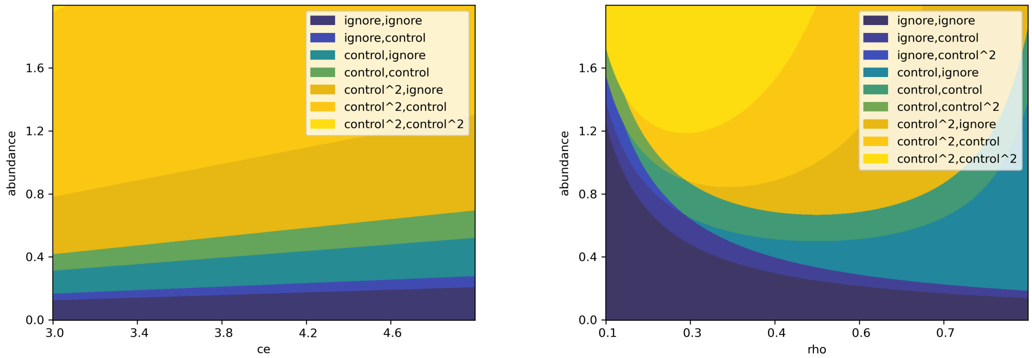

(a) Optimal actions, for varying c_e.

(b) Optimal actions, for varying ρ.

Figure 3: Parameter sensitivity analysis for control parameters (the cost of controlling c_e and control efficacy ρ) with no initial uncertainty ($\Delta n = 0$), for varying abundance $\overline{N}$ for $T = 3$. Apart from c_e (left) and ρ (right), parameters are kept the same to the example in Figure 2.

cost parameter, Figure 3a demonstrates that as the cost increases (moving rightwards on the graph), the species must be more abundant to justify a given control action. Figure 3b analyses the efficacy parameter ρ in an analogous way. For a small ρ, the abundance must be very high to justify any control action. As ρ increases, it becomes worthwhile to control the species at smaller abundances. For ρ close to 1, controlling once tends to be sufficient — as such, the regions corresponding to more expensive control actions (yellow and gold) become smaller.

To examine the effect of varying parameters related to monitoring, we can plot the border of the region where monitoring is optimal, as abundance and uncertainty vary. This corresponds, for example, to the edge of the orange region in Figure 2. Figure 4 plots these frontiers for varying monitoring cost c_m and error e. As either of these parameters increase, monitoring becomes less worthwhile, and as such needs more uncertainty to justify it. This is demonstrated by the fact that the frontiers corresponding to larger cost and error (coloured yellow) are further to the right.

4. Simulation Study

We develop a simulation model to demonstrate how our decision-framework could be used to support decisions. The structure of the simulation is that there is a true underlying abundance of the species, and that gets updated in each time step, deepening on the management decision. At each time step, the decision is made by choosing the optimal decision, which is calculated always assuming a set time window (i.e. we do not update the mapping of abundance and uncertainty to decision through time).

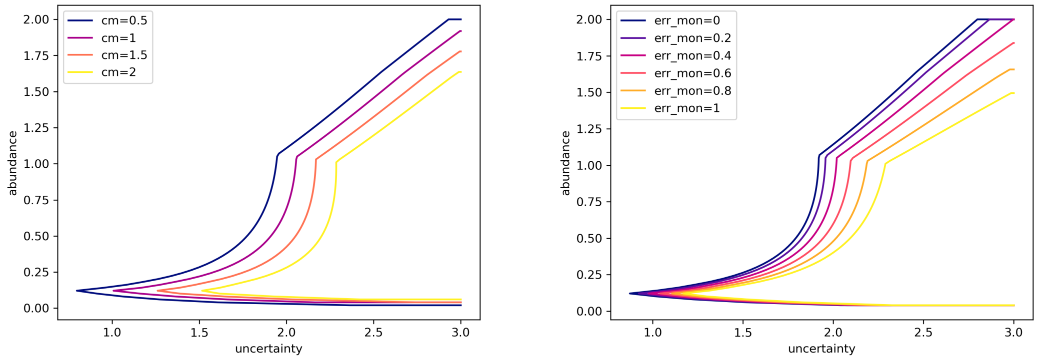

(a) Monitoring frontiers, for varying c_m.

(b) Monitoring frontiers, for varying e.

Figure 4: Parameter sensitivity analysis for monitoring parameters (the cost of monitoring c_m and the monitoring error e) showing the frontier where monitoring becomes optimal, plotted on the same axes as Figures 1 and 2. Other parameters are kept at the values used in the example in Figure 2.

The simulation begins with a belief of the manager, for both abundance and uncertainty, alongside a hidden true abundance. At each time step, the abundance estimate and uncertainty changes depending on the action chosen. If the action is ‘do nothing’ or ‘monitor’ then the increase is multiplicative, drawn from a normal distribution with mean r and standard deviation dr. The difference between ‘do nothing’ and ‘monitor’, is that the subsequent uncertainty estimate for ‘monitor’ is greatly reduced. If ‘control’ is chosen, then the abundance is reduced by a fixed percentage.

To validate the effectiveness of our algorithm, we investigate how our solution performs across many such simulations and compared to a simpler ‘naive’ algorithm. The naive algorithm solves the same problem except that there is no option to monitor, so there is no integration between the control decisions and monitoring decisions. We explore the performance of the algorithms, depending on how accurate the initial abundance estimate (Figure 5). We find the full solution largely outperforms the naive solution. The only exception is if there is a large initial underestimate of the invasive species abundance, and in these cases both solutions perform similarly.

To demonstrate how our method might be applied in practice, we run this simulation using parameter estimates for tropical fire ants (*Solenopsis geminata*) on Ashmore Reef, Australia. The resulting timeseries simulation is shown in Figure 6, and the sequence of actions alternates through ‘control’, ‘do nothing’ and ‘monitor’. The parameters for the case study are listed in Table 2. The growth rate, its uncertainty and control efficacy are set following Baker et al. (2017), while parameter estimates for costs of control effort and monitoring are taken from Walshe (2017). The optimal solution is calculated using a discount

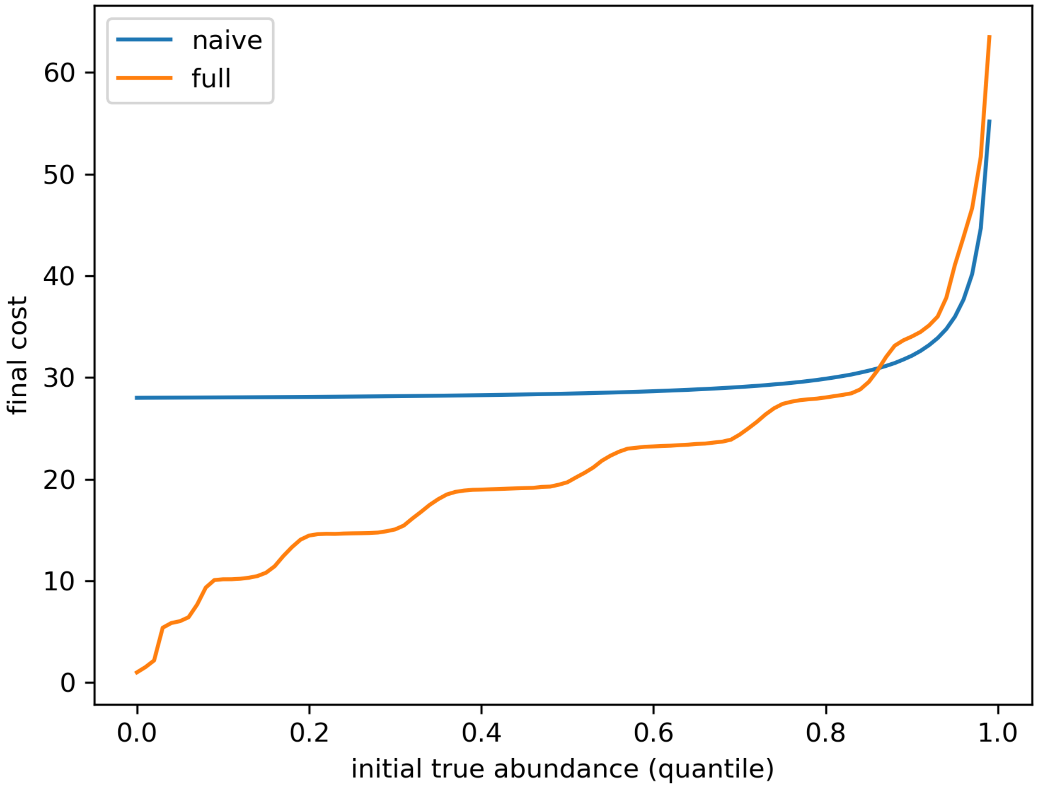

Figure 5: The average final cost for the full solution and the naive solution, depending on initial abundance. The parameters are as in Figure 2, except that the time horizon is extended to five steps. For all simulations, the estimate of the initial abundance is held constant, while the true initial abundance is varied. Hence, initial true abundance values less than 0.5 mean that there is an *overestimate* of the abundance, while if the initial true abundance is greater than 0.5, then there is an underestimate of the true abundance. Low values of the final cost correspond to better performance, compared to high values of final cost. The 'bumps' in the full solution likely correspond to an interaction between the time-window and when monitoring is used.

Table 2: Parameter values for the tropical fire ant cast study simulation.

Parameter	Value
r	2.82
dr	0.015
control cost	31.8
control efficacy	99%
monitoring cost	6.2

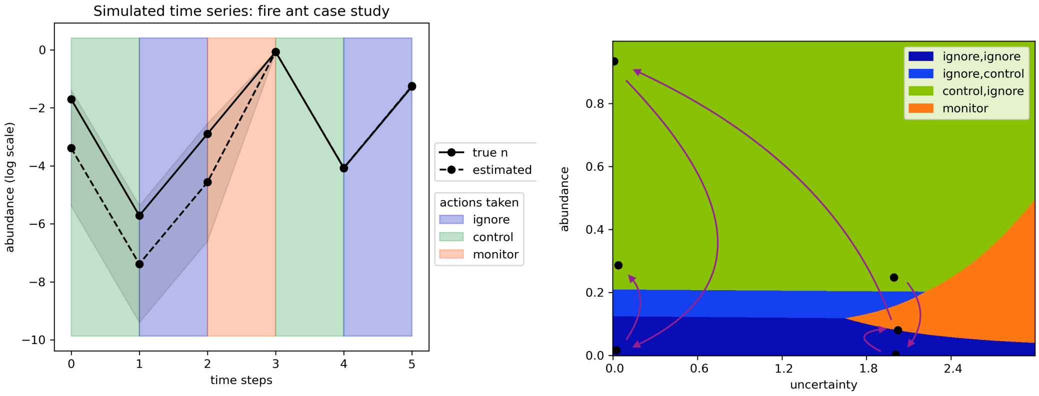

Figure 6: An example time-series for the fire ant case study (left). The high estimated abundance in the first time-step means the optimal decision is to treat and then do nothing. Then the large uncertainty means the optimal decision is to monitor in the third year. The sequence of decisions are mapped out on the phase plane diagram for these parameters (right).

rate of $\gamma = 0.9$ and over a 2-year time-horizon.

5. Discussion

In this paper, we have presented a partially observable Markov decision process (POMDP) framework for informed invasive species management. We believe that the density projection method that we advocate has the potential to make POMDP methods accessible to a broad audience in invasive species management, and ecology more broadly. We work through the development of a model in this framework, and provide an intuitive, visual tool that can help to decide between when to monitor and when to control given an estimate (with uncertainty) of the species abundance.

It is natural to parametrise an ecological problem, such as an invasive species, by its growth rate and the relative costs of control or monitoring compared to doing nothing at all. While our assumptions of exponential dynamics limit the quantitative realism of our model, we believe that its simplicity and interpretability make it a broadly applicable starting point for decision making in the presence of uncertainty about the state of the system being managed.

In aiming for an intuitive and generic model of uncertain abundance, we make different simplifications to previous work. In the past, POMDP approaches have simplified the ecological system dynamics to make the problem computationally tractable (Williams and Brown, 2022; Haight and Polasky, 2010; Regan et al., 2011; Chadès et al., 2011; Rout et al., 2014). Instead, we simplify the computational problem by assuming a parametric distribution for the population size, as recommended in Section 8 of Williams and Brown (2022), which gives the dimensions of our model intuitive content. Therefore, a major advantage of our POMDP approach is that it can deal with uncertainty in the state of the system while providing easily interpretable results.

Similar problems and discussions can be found in the related field of disease management, where an analogous trade-off exists between monitoring, i.e. testing members of the population to gauge the spread of disease, and control, say by deploying medication, vaccination or implementing quarantine strategies. Although a variety of strategies have been proposed, POMDP solutions are still limited (Nowzari et al., 2016; Zino and Cao, 2021). Diseases whose management involves both surveillance (i.e. testing) and control (vaccination, medications and other control strategies) elements, such as HIV and COVID, could be particularly suitable for our method. One initial interesting proposal in this field is due to Hauskrecht and Fraser (2000), who apply a POMDP framework for diagnosis and treatment of ischemic heart disease.

Our framework is currently mainly limited by the assumptions of our model — for example, we assume that the population of the species grows exponentially, and ignore interactions with the rest of the environment. Our model could be extended to more general models of species infestation, possibly at the cost of a numerical solution. However, for small outbreaks and short time-frames exponential growth is often applicable — since these situations frequently coincide with high uncertainty, we believe that our model has value despite its simplicity.

Our model also contains the assumption that certain parameters, such as the control efficacy ρ, are known exactly, which could be relaxed with further work. One way to include parameter uncertainty would be to link our work to an adaptive management framework, such as described by Moore et al. (2017). In such a framework, r or ρ would be added as dimensions to the model, so that the belief state of the manager includes a distribution over possible values of the parameters (Chadès et al., 2021, Section 4.3). In this situation, monitoring actions have the additional benefit of reducing parameter uncertainty, such as about the efficacy of management.

In this paper, we formulate a POMDP model which is applicable to a wide range of invasive species scenarios. In contrast with traditional POMDP approaches (see Williams and Brown (2022)), we limit our model to dimensions of abundance and its uncertainty, retaining sufficient complexity to model the trade-off between monitoring and control, while remaining analytically solvable. This idea is not limited to the field of invasive species management, and has the potential to be applied to the 'reverse' problem of endangered species conservation, and to management of disease. POMDP approaches to the conservation of Sumatran tigers are proposed by Chadès et al. (2008) and McDonald-Madden et al. (2011), who limit the state of the system to 'present' or 'absent'. The density-projection framework proposed in this paper could be altered to investigate these scenarios with a more general model of abundance.

Author Contributions

Thomas Waring developed the model based on preliminary work by Michael McCarthy in consultation with all authors. Vera Somers and Thomas Waring drafted the manuscript. Thomas Waring ran the simulations and generated the results with input from Vera Somers. All authors edited and improved the manuscript.

Acknowledgements

We thank the Associate Editor and the two reviewers for their comments, which helped us improve the paper.

Conflict of Interest Statment

No authors have conflicts of interest to declare.

Data Availability Statement

Python code implementing the model described in this paper is available at `https://github.com/thomaskwaring/to-monitor-or-control`, including scripts to reproduce the figures. See also Zenodo archive at DOI 10.5281/zenodo.11560537 Waring et al. (2024). No data were used.

References

S. Alvarez and D. Solís. Rapid response lowers eradication costs of invasive species. *Choices*, 33(4):1–9, 2018.

K. J. Åström. Optimal Control of Markov Processes with Incomplete State Information I. *Journal of mathematical analysis and applications*, 10:174–205, 1965. ISSN 0022-247X. doi: {10.1016/0022-247X(65) 90154-X}.

C. M. Baker, J. C. Hodgson, E. Tartaglia, and R. H. Clarke. Modelling tropical fire ant (solenopsis geminata) dynamics and detection to inform an eradication project. *Biological Invasions*, 19:2959–2970, 10 2017. ISSN 15731464. doi: 10.1007/S10530-017-1499-9/FIGURES/7.

R. Bellman. *Dynamic Programming*. Dover Books on Computer Science Series. Dover Publications, 2003. ISBN 9780486428093.

I. E. Büyüktahtakın and R. G. Haight. A review of operations research models in invasive species management: state of the art, challenges, and future directions. *Annals of Operations Research*, 271: 357–403, 12 2018. ISSN 15729338. doi: 10.1007/S10479-017-2670-5/TABLES/2.

L. R. Carrasco, R. Baker, A. Macleod, J. D. Knight, and J. D. Mumford. Optimal and robust control of invasive alien species spreading in homogeneous landscapes. *Journal of the Royal Society Interface*, 7: 529–540, 2010. doi: 10.1098/rsif.2009.0266.

I. Chadès, E. McDonald-Madden, M. A. McCarthy, B. Wintle, M. Linkie, and H. P. Possingham. When to stop managing or surveying cryptic threatened species. *Proceedings of the National Academy of Sciences*, 105(37):13936–13940, 2008.

I. Chadès, T. G. Martin, S. Nicol, M. A. Burgman, H. P. Possingham, and Y. M. Buckley. General rules for managing and surveying networks of pests, diseases, and endangered species. *Proceedings of the National Academy of Sciences*, 108(20):8323–8328, 2011.

I. Chadès, E. McDonald-Madden, M. A. McCarthy, B. Wintle, M. Linkie, and H. P. Possingham. When to stop managing or surveying cryptic threatened species. *Proceedings of the National Academy of Sciences of the United States of America*, 105:13936–13940, 9 2008. ISSN 00278424. doi: 10.1073/PNAS.0805265105/SUPPL_FILE/0805265105SI.PDF.

I. Chadès, L. V. Pascal, S. Nicol, C. S. Fletcher, and J. Ferrer-Mestres. A primer on partially observable Markov decision processes (POMDPs). *Methods in Ecology and Evolution*, 12(11):2058–2072, 2021. ISSN 2041-210X. doi: 10.1111/2041-210X.13692.

R. S. Epanchin-Niell, R. G. Haight, L. Berec, J. M. Kean, and A. M. Liebhold. Optimal surveillance and eradication of invasive species in heterogeneous landscapes. *Ecology Letters*, 15:803–812, 8 2012. ISSN 1461023X. doi: 10.1111/J.1461-0248.2012.01800.X.

A. Gelman, J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin. *Bayesian Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis, 2013. ISBN 9781439840955.

R. G. Haight and S. Polasky. Optimal control of an invasive species with imperfect information about the level of infestation. *Resource and Energy Economics*, 32(4):519–533, 2010.

C. E. Hauser and M. A. McCarthy. Streamlining 'search and destroy': cost-effective surveillance for invasive species management. *Ecology Letters*, 12:683–692, 7 2009. ISSN 1461-0248. doi: 10.1111/J.1461-0248.2009.01323.X.

M. Hauskrecht and H. Fraser. Planning treatment of ischemic heart disease with partially observable markov decision processes. *Artificial intelligence in medicine*, 18(3):221–244, 2000.

B. D. Hoffmann and L. M. Broadhurst. The economic cost of managing invasive species in australia. *NeoBiota*, 31:1–18, 2016.

K. Hyytiäinen, M. Lehtiniemi, J. K. Niemi, and K. Tikka. An optimization framework for addressing aquatic invasive species. *Ecological Economics*, 91:69–79, 2013.

J. Jawitz. Moments of truncated continuous univariate distributions. *Advances in Water Resources - ADV WATER RESOUR*, 27:269–281, 03 2004. doi: 10.1016/j.advwatres.2003.12.002.

M. J. Keeling and P. Rohani. *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press, 2008. ISBN 978-0-691-11617-4. doi: 10.2307/j.ctvcm4gk0.

M. L. Littman. A tutorial on partially observable markov decision processes. *Journal of Mathematical Psychology*, 53(3):119–125, 2009.

L. Marescot, G. Chapron, I. Chadès, P. L. Fackler, C. Duchamp, E. Marboutin, and O. Gimenez. Complex decisions made simple: a primer on stochastic dynamic programming. *Methods in Ecology and Evolution*, 4(9):872–884, 2013.

M. L. N. Mbah and C. A. Gilligan. Balancing detection and eradication for control of epidemics: Sudden oak death in mixed-species stands. *PLOS ONE*, 5:e12317, 2010. ISSN 1932-6203. doi: 10.1371/JOURNAL.PONE.0012317.

E. McDonald-Madden, I. Chadès, M. A. McCarthy, M. Linkie, and H. P. Possingham. Allocating conservation resources between areas where persistence of a species is uncertain. *Ecological Applications*, 21(3):844–858, 2011.

S. V. Mehta, R. G. Haight, F. R. Homans, S. Polasky, and R. C. Venette. Optimal detection and control strategies for invasive species management. *Ecological Economics*, 61(2-3):237–245, 2007.

G. E. Monahan. State of the art—a survey of partially observable markov decision processes: theory, models, and algorithms. *Management science*, 28(1):1–16, 1982.

A. L. Moore, L. Walker, M. C. Runge, E. McDonald-Madden, and M. A. McCarthy. Two-step adaptive management for choosing between two management actions. *Ecological Applications*, 27(4):1210–1222, 2017. doi: https://doi.org/10.1002/eap.1515.

J. L. Moore, T. M. Rout, C. E. Hauser, D. Moro, M. Jones, C. Wilcox, and H. P. Possingham. Protecting islands from pest invasion: optimal allocation of biosecurity resources between quarantine and surveillance. *Biological Conservation*, 143:1068–1078, 2010. doi: 10.1016/j.biocon.2010.01.019.

C. Nowzari, V. M. Preciado, and G. J. Pappas. Analysis and Control of Epidemics: A Survey of Spreading Processes on Complex Networks. *IEEE Control Systems Magazine*, 36(1):26–46, 2016. ISSN 1066033X. doi: 10.1109/MCS.2015.2495000.

S. Polasky. A model of prevention, detection, and control for invasive species. In *Bioinvasions and Globalization: Ecology, Economics, Management, and Policy*. Oxford University Press, 2010.

T. J. Regan, I. Chadès, and H. P. Possingham. Optimally managing under imperfect detection: a method for plant invasions. *Journal of Applied Ecology*, 48(1):76–85, 2011.

T. M. Rout, J. L. Moore, H. P. Possingham, and M. A. Mccarthy. Allocating biosecurity resources between preventing, detecting, and eradicating island invasions. *Ecological Economics*, 71:54–62, 2011. doi: 10.1016/j.ecolecon.2011.09.009.

T. M. Rout, J. L. Moore, and M. A. Mccarthy. Prevent, search or destroy? a partially observable model for invasive species management. *Journal of Applied Ecology*, 51:804–813, 6 2014. ISSN 1365-2664. doi: 10.1111/1365-2664.12234.

T. Walshe. An assessment of alternative management interventions for treatment of tropical fire ants on ashmore reef. Technical report, Australian Institute of Marine Science, 2017.

T. K. Waring, V. L. J. Somers, M. A. McCarthy, and C. M. Baker. Code for: When to monitor or control; informed invasive species management using a partially observable Markov decision process (POMDP) framework, June 2024.

B. K. Williams and E. D. Brown. Partial observability and management of ecological systems. *Ecology and Evolution*, 12(9):e9197, 2022.

E. Zhou, M. Fu, and S. Marcus. Solving continuous-state pomdps via density projection. *Automatic Control, IEEE Transactions on*, 55:1101 – 1116, 06 2010. doi: 10.1109/TAC.2010.2042005.

L. Zino and M. Cao. Analysis, prediction, and control of epidemics: A survey from scalar to dynamic network models. *IEEE Circuits and Systems Magazine*, 21(4):4–23, 2021.