

RESEARCH ARTICLE

A general optimal adaptive framework for managing a threatened species

Sam Nicol¹  | Jonathan Ferrer-Mestres¹ | Hannah Lloyd² | James Brazill-Boast³ |
Emma Gorrod^{2,4} | Iadine Chades¹ 

¹CSIRO, Land and Water, Dutton Park, Queensland, Australia

²Saving Our Species Program, Department of Planning Industry and Environment, Sydney, New South Wales, Australia

³Biodiversity Conservation Trust, Sydney, New South Wales, Australia

⁴Centre for Ecosystem Science, University of New South Wales, Kensington, New South Wales, Australia

Correspondence

Sam Nicol, CSIRO Land and Water, Ecosciences Precinct, 41 Boggo Road, Dutton Park QLD 4102, Australia.
Email: sam.nicol@csiro.au

Handling Editor: Punyasloke Bhadury

Funding information

NSW Department of Planning, Industry and Environment

Abstract

1. Managers must determine which interventions best protect threatened species when the outcomes of interventions are uncertain. Adaptive management is a dynamic optimization approach that generates optimal management actions based on current knowledge while learning to improve future management outcomes. Although adaptive management theory is well-developed, uptake has been impeded by its complexity and a tendency to develop bespoke solutions with high implementation costs for problem-specific returns.
2. To increase uptake of adaptive management and improve threat management for species recovery, we developed a general adaptive management decision model, framed as a Mixed Observability Markov Decision process. We embraced principles of generality, simplicity and interpretability to overcome previous implementation challenges. We created a general model structure that is applicable to any species–threat combination, thus avoiding the need to develop customized models for every species. Simplicity was achieved by minimizing states to reduce the information requirements for parameterization. To improve interpretability, we implemented our method as a Shiny application and employed a recent artificial intelligence approach to simplify the optimal strategy. We applied our approach to a case study of fox impacts on a threatened marsupial.
3. Our case study shows that when one management action is robust to uncertainty, the value of information of optimal adaptive management may be low. Cases like these highlight species–threat combinations where investment in adaptive management is not required.
4. Our tool provides a rapid prototype adaptive management approach with minimal cost to management agencies. Our simple yet general model structure improves efficiency for implementing adaptive management for large numbers of threatened species, improving the effectiveness of conservation investments.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Ecological Solutions and Evidence* published by John Wiley & Sons Ltd on behalf of British Ecological Society.

KEYWORDS

adaptive management, artificial intelligence, conservation, decision theory, Markov decision process, MOMDP, SARSOP, threatened species

1 | INTRODUCTION

In the face of accelerating global species decline, conservation biologists must decide how and when to act despite having limited information about effectiveness (Williams & Johnson, 2013). This uncertainty impacts the success of threat management and species recovery (Nicol et al., 2019), so there is an incentive to apply actions which best protect species. A solution to this problem is adaptive management (Chadès et al., 2017; Walters, 1986; Williams, 2011), a dynamic ‘learning by doing’ approach that learns the most effective actions from implementation and monitoring of management outcomes.

Optimal adaptive management solutions find the best sequential decisions where the outcome of actions can be represented probabilistically and the current and/or future system states may be uncertain. The most common way of solving adaptive management problems uses stochastic dynamic programming and Bayesian inference (see Chadès et al., 2017 for a review). Adaptive management has been developed to manage natural resource management (Johnson et al., 1997; Memarzadeh et al., 2019; Nicol et al., 2014), species conservation (Fackler et al., 2014; Rout et al., 2009) and epidemiology (Atkins et al., 2020; Shea et al., 2014). Although adaptive management is conceptually embedded in conservation and land management agencies (Williams et al., 2009), the optimal formulation has had limited uptake beyond theoretical papers (Runge, 2011; Westgate et al., 2013; but see Williams & Johnson [1995] for one long-running example of application).

The reasons why conservation studies are not implemented include poor engagement, failure to embed studies within a broader decision-making framework, and failure to acknowledge the social dimensions of conservation actions (Knight et al., 2008). These issues apply to adaptive management; however, based on the authors’ experience of more than a decade designing optimal adaptive management studies, we assert that the ‘Decision-Theoretic’ (Runge, 2011) optimization approach to adaptive management has additional implementation challenges that may further explain the low uptake. Specifically:

- (i) The complex solutions produced by optimal adaptive management methods can be difficult to interpret.
- (ii) Most published studies are customized, standalone cases that have steep development costs but limited applicability outside their specific domain.
- (iii) The modelling process is time-consuming and a specialist skill.
- (iv) The modelling processes often lack simple user interfaces and can be difficult to install and run (but see MDPSOLVE [Fackler, 2011], the MDPToolbox [Chadès et al., 2014] and an implementation of SARSOP in R [Boettiger et al., 2018] for tools that partially overcome this issue).

In this study, we developed a tool that aims to overcome these barriers to implementation. We collaborated with the Saving our Species program from the Australian state of New South Wales (NSW) to develop an adaptive management approach to guide the implementation of actions for any species–threat combination. Like many other jurisdictions globally, NSW has a large threatened species list (>1000 listed species and communities) and 39 key threatening processes. Designing custom adaptive management approaches for all listed species and threats is not cost-effective or feasible, yet there is a strong push to document the effectiveness and improve the outcomes of investments in threatened species management (Brazill-Boast, 2018).

The adaptive management approach that we propose here is cost-effective for management agencies because it saves on the significant time and expense involved with designing and implementing custom adaptive management studies for every listed threatened species (Stem et al., 2005). Many jurisdictions globally have hundreds or thousands of listed species, and management uncertainty is ubiquitous across threatened species management. The rapid prototyping approach that we propose allows managers to quickly and cheaply generate optimal policies for many species, identifying which actions to implement and which management questions have significant value of information for further modelling.

We built on previous studies (Chadès et al., 2012; Nicol & Chadès, 2012; Nicol et al., 2015) to propose a simple yet general problem formulation that can be applied to any species–threat combination. Previous adaptive management research has shown that similar performance can often be achieved with smaller state spaces (Ferrer-Mestres et al., 2021; Nicol & Chadès, 2012; Pascal et al., 2021), which have the advantage of being more interpretable and easier to collect data for. We used a minimal set of categorical states to reduce the number of model parameters and the information required for parameterization. Further, we implemented our approach as a Shiny application in R (available from <https://conservation-decisions-lw.shinyapps.io/SpeciesThreatAM/>). These modifications allow users to experiment and solve optimal adaptive management problems with minimal input, allowing for rapid experimentation to explore optimal policies for different threat–species case studies.

2 | MATERIALS AND METHODS**2.1 | Method overview**

Our contribution focuses on developing a general approach for managing a species and its main threat. The goal is to cost-effectively minimize the probability of extinction of a population of a threatened

Managers need to cost-effectively protect a species from a threatening process, e.g. protect potoroos from predatory red foxes.



Managers can **choose from 3 management actions**. Here they can control the threat by baiting foxes, using either:




They can also take no action (A0). Each action has a different cost.

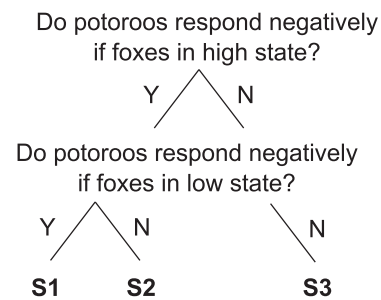
We don't know whether each action reduces foxes nor whether potoroos recover if foxes decrease, but we can elicit all possible combinations ("models").

There are 5 possible threat models (F1-F5) and 3 species models (S1-S3):

Threat models



Species models



Given a starting state and an estimate of the uncertainty about the true model ("belief"), the adaptive management determines the optimal action to protect the potoroo given the uncertainty about which model is correct. Over time, taking actions and observing outcomes improves the belief over time, leading to better predictions as the model uncertainty decreases.

FIGURE 1 Schematic of the threat–species adaptive management problem applied to our case study example of the fox (threat) and potoroo (species)

species over time by implementing actions to reduce threat intensity (see Figure 1). However, it is unclear to managers which management actions are most likely to be effective over time. This type of problem is common in conservation (Salafsky et al., 2002) and is faced by many jurisdictions that manage portfolios of threatened species using limited resources.

Our approach tracks two entities, threat and species, which are modelled using a Markovian state-transition framework. The value of the species or threat at a point in time is called a state. States are described using categorical descriptors (e.g. locally extinct, low or high species state), and transitions describe the probabilities of changing

state after taking an action. The effects of an action may differ with the state of the system, for example the most effective action for managing a low threat level may differ from the most effective action to manage a high threat level. Threats can be managed through threat reduction actions (e.g. active control of invasive predators), and threat reduction affects the species state. The effectiveness of the action in reducing the threat and the impact of threat reduction on the species are both uncertain. We consider all binary combinations of threat reduction (i.e. actions are effective/ineffective at reducing threats) and species response (i.e. threat states impact/have no impact on species). We call each of these combinations 'models'. Each model is a possible

hypothesis about how threat reduction and species response interact. Given any observed state and a set of beliefs in each model (i.e. probabilities that each model is true), the optimization determines the most effective action given the uncertainty about both threat reduction effectiveness and species response.

In the following sections, we describe the adaptive management approach before presenting an illustrative case study.

2.2 | Model definition

Our approach models the expected outcomes of discrete management interventions for a species that is impacted by a threat. As in previous adaptive management studies (Chadès et al., 2012), we model the system with a Mixed Observability Markov Decision Process (MOMDP; Ong et al., 2010). MOMDPs scale up to larger size problems more efficiently than traditional grid-based belief MDP approaches (Chadès et al., 2012; Ong et al., 2010).

The MOMDP is defined by a tuple $\langle A, X, Y, O, T_x, T_y, Z, \theta, \gamma \rangle$:

- A is the set of finite discrete candidate management actions $a_i \in A$, $i \in \{1, \dots, |A|\}$. We assume that this set includes a do-nothing action.
- The state space $X \times Y$ is composed of discrete fully observable (X) and unobservable components (Y). Fully observable components can be measured without uncertainty. Unobservable components cannot be measured, so they must be inferred based on predictions about the expected observed dynamics.
- $O = O_X \times O_Y$ is the finite observation space.
- T_x and T_y are the transition matrices that describe the dynamics of the observable and unobservable states, that is they specify the (Markov) probabilities of transitioning from the current state to any other state when action a_i is taken.
- Z describes the probability of observing a state if an action is taken.
- $\theta(x, a)$ is the reward associated with state x after taking action a .
- γ is a discount factor which determines the convergence rate of the dynamic optimization algorithm, which we solve for the infinite time horizon ($0 < \gamma < 1$).

For our species–threat adaptive management formulation, we define the MOMDP tuple as follows:

- States $X \times Y$ are a combination of fully observable species (s) and threatening process (p) states¹ ($x = (s, p)$; $s \in S$, $p \in P$, $x \in X$) and unobservable ‘models’ of threat reduction and species response ($y = (y_s, y_p)$; $y_s \in Y_s$, $y_p \in Y_p$, $y \in Y$). Specifically,
 - Species components are defined by three states representing qualitative measures of abundance: $S = \{\text{Locally Extinct, Low, High}\}$. Threatening process states are defined by two states $P = \{\text{Low, High}\}$. State definitions should be agreed with domain experts on a case-by-case basis. Experts may refer to their experience or published literature to help them define the state thresholds.

- We allow two responses of the threatening process to management, which we denote using a binary variable r_a . Actions can be either ineffective or effective at reducing the threat, denoted $r_a \in \{0, 1\}$, respectively. A threat model, $y_p \in Y_p$; $y_p = (r_{a_0}, r_{a_1}, \dots, r_{a_{|A|}})$, describes whether the threat will be effectively managed for each action. If we enumerated all possible models, the number of threat models would be $2^{|A|-1}$ (assuming that the do-nothing action cannot be effective). Threat models are defined according to a matrix $T(r_a|y_p)$, which has dimension $|Y_p| \times |A|$. The elements $\Pr(r_a|y_p) \in T(r_a|y_p)$ are encoded such that effective actions have value 1 and ineffective actions have value 0. To simplify notation, the model corresponding to row i of $T(r_a|y_p)$ is referred to as F_i . Without loss of generality, but to include the effect of applying multiple actions, we here solve a special case where $|A| = 4$, but assume this consists of a do-nothing action a_0 , two distinct actions a_1 and a_2 , and an action a_3 that implements both a_1 and a_2 . By assuming that a_3 has this known structure, some combinations of effectiveness are not possible (e.g. if a_1 or a_2 is effective, then a_3 cannot be ineffective) and the number of models Y_p is reduced to $|Y_p| = 5$. Formally,

$$T(r_a|y_p) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 \end{bmatrix},$$

where rows represent the models y_p and columns represent the effectiveness of actions a_0, \dots, a_3 , respectively.

- We allow three models of species response to the threat. The response of species to a threat state is represented as either impacting the species or not, denoted using $r_p \in \{0, 1\}$, respectively. A species model, $y_s \in Y_s$; $y_s = (r_{p_1}, r_{p_2}, r_{p_3})$, describes whether the species will be impacted given the threat state. Species can respond (1) negatively to any level of threat presence (Low or High); (2) negatively to high threat presences (no impact of low threat); or (3) no response to threat presence (no impact of either high or low threat). Species response is independent of the action and depends only on the threat state, that is $|Y_s| = 3$. The three models are defined with a matrix where response to threat is encoded as 1 and no response is coded as 0. Formally,

$$T(r_p|y_s) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \\ 0 & 0 \end{bmatrix},$$

where rows represent the models y_s and columns represent effectiveness of threat states $r_p = \{\text{Low, High}\}$, respectively. Elements of $T(r_p|y_s)$ are denoted as $\Pr(r_p|y_s)$. To simplify notation, the model corresponding to row i of $T(r_p|y_s)$ is referred to as S_i .

- $O = O_X \times O_Y$ is the finite observation space. In this problem, $O = X$, that is X is observable and the models Y are unobservable.
- $T_x(x, y, a, x') = \Pr(x'|x, y, a)$ is the probability of transitioning to state x' from state (x, y) , given that action a is implemented. Similarly,

¹ We use the notation $x = (s, p)$ and $y = (y_s, y_p)$ for simplicity and consistency with previous MOMDP presentations.

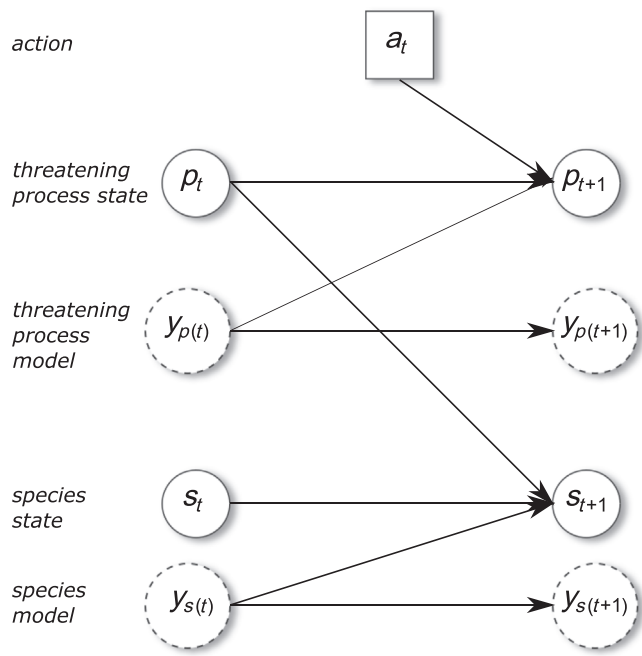


FIGURE 2 Influence diagram for the species–threat adaptive management problem showing the conditional dependence of variables. Arrows show the direction of influence between variables.

$T_y(x, y, a, x', y') = \Pr(y'|x, y, a, x')$ is the probability of transitioning from state y to y' when action a is implemented and the observed state changes from x to x' . We assume that the model dynamics are stationary (i.e. that the reactions of species and threats to management actions do not change over time due to climate change or other factors), so that $T_y(x, y, a, x', y') = 1_{y=y'}$ and 0 otherwise.

- $Z(a, x', y', o') = \Pr(o'|a, x', y')$ is the probability of observing $o' \in O$ if the state is (x', y') after taking action a . In our problem, Z is simplified because state components are either perfectly observed (probability of correct observation = 1) or unobservable (probability of observation = 0). Formally, since Y is unobservable, $Z(a, x', y', o') = 1_{x'=o'}$ (0 otherwise).
- $\theta(x, a) = \theta(s, a) = v(s) - c(a) + \max(c(a))$, where $v(s) \geq 0$ is the benefit of maintaining the species in state s , and $c(a) \geq 0$ is the cost of action a . The term $\max(c(a))$ is a constant that forces the reward function to be positive, which is useful for simplifying the MOMDP policy.
- We set the discount factor controlling optimization convergence to $\gamma = 0.9^2$.

The structure of our MOMDP can be represented using an influence diagram (Figure 2).

Our problem framing is an example of adaptive management under model uncertainty (Chadès et al., 2017), whereby a discrete set of

pre-specified models are evaluated. This approach assumes that the expert-specified transition probabilities in one of the candidate models are similar to the true system dynamics, but this assumption is difficult to test. The number of possible models depends on the state and action spaces, and usually only a subset of plausible models can feasibly be included as candidates to represent the true dynamics. However, because our formulation uses a small state space, we can evaluate all possible models that allow for the threat to have negative impacts on the species. Our model approach can capture the range of threat responses for any two actions and the interactions between them.

2.3 | Parameterization

We simplify the transition matrices to account for independence between state variables (Figure 2). The full transition matrix is specified by

$$T(x', y'|x, y, a) = T_x(x, y, a, x') T_y(x, y, a, x', y').$$

Recall that $T_y(x, y, a, x', y') = 1_{y=y'}$; 0 otherwise, so parameterization requires specifying the matrices $T_x(x, y, a, x')$ that define the probabilities of transitioning between observable states when an action is implemented. Because the threat is independent of the species response and the species response is independent of the action, the elements of $T(x', y'|x, y, a)$ can be simplified to

$$\Pr(x'|x, y, a) = \Pr(p'|p, a, y_p) \Pr(s'|s, p, y_s).$$

We parameterize the models by eliciting the transition matrices $\Pr(p'|p, a, y_p)$ for each action.

Threat models y_p are defined by specifying which actions are effective. We assume that ineffective actions follow the transition probabilities for the do-nothing action. The transition matrix for any threat model y_p can thus be built as a combination of the transition matrices for each action, that is

$$\begin{aligned} \Pr(p'|p, a, y_p) &= \Pr(r_a|y_p) \Pr(p'|p, a, r_a) \\ &+ (1 - \Pr(r_a|y_p)) \Pr(p'|p, \text{do nothing}, r_a). \end{aligned}$$

A simpler way to express this is

$$\Pr(p'|p, a, y_p) = \begin{cases} \Pr(p'|p, a), & \text{if } a \text{ effective under } y_p; \\ \Pr(p'|p, \text{do nothing}), & \text{if } a \text{ ineffective under } y_p. \end{cases}$$

Our assumption of binary responses to threat means that $\Pr(p'|p, a, y_p)$ depends on $|A|$ matrices of dimension $|P| \times |P| = 2 \times 2$. We need only elicit two probabilities for each action, for example $P(p' = \text{High}|p = \text{High}, a)$ and $P(p' = \text{Low}|p = \text{Low}, a)$, with the other two probabilities inferred since $\sum_{p'} P(p'|p, a) = 1$. Consequently, we require $2|A|$ elicitation questions to parameterize the full set of threat models Y_p . The elicitation burden can be further reduced by

² The discount factor determines the relative contribution of future utility compared to current utility. In most ecological applications, it is set between 0.9 and 1. We used 0.9 in the results reported in this paper but also tested a discount factor of 0.99. The optimal policy was unchanged by the increased discount factor.

interpolation following the methods outlined in Cain (2001), although the gains are minimal for small numbers of actions.

Species models y_s are defined by specifying whether species are impacted by threat states. We assume that if species are not affected by a threat, then the response follows the transition probabilities for species in the absence of the threat. The transition matrix for any species model y_s can thus be built as a combination of the transition matrices for each threat level, that is

$$\Pr(s'|s, p, y_s) = \Pr(r_p|y_s) \Pr(s'|s, p, r_p) + (1 - \Pr(r_p|y_s)) \Pr(s'|s, \text{Not Present}, r_p),$$

where the response of the species in the absence of the threat is denoted using the 'Not Present' notation. The previous expression can be written as

$$\Pr(s'|s, p, y_s) = \begin{cases} \Pr(s'|s, p), & \text{if threat level } p \text{ impacts species under } y_s; \\ \Pr(s'|s, \text{Not Present}), & \text{if threat level } p \text{ has no impact on species under } y_s. \end{cases}$$

Our assumption that species are impacted by threats in a binary way means that $\Pr(s'|s, p, y_s)$ depends only on three species transition matrices. Matrices containing the values of $\Pr(s'|s, p)$ are elicited to cover the range of possible species responses in y_s , that is one to represent the species response for each threat level $p = \{\text{Low, High}\}$ and one for when the threat is not present. Elicited $\Pr(s'|s, p)$ matrices are 3×3 matrices; however, one row is known a priori ($P(s' = \text{Locally Extinct} | s = \text{Locally Extinct}, p) = 1; 0$ otherwise); so we need only elicit four probabilities for each matrix; a total of 12 elicitation questions (see data files at <https://doi.org/10.6084/m9.figshare.16386510> for an example elicitation).

2.4 | Optimization

Any mapping of states to actions is called a *policy*. The objective in adaptive management is to find the policy which yields the highest discounted expected sum of rewards over time. This *optimal policy* depends on the state of the system and a belief $b(y)$, which is a vector of length $|Y|$ whose elements represent the probability that each model is correct. As actions are taken, the belief is updated ('adapted') to summarize observed outcomes. The optimization finds the best action for the belief space, that is for any updated belief, we can query the optimal policy and return the best action.

For our problem, maximizing the reward means selecting management actions to maximize the return on investment, which is determined by the balance between the benefit of maintaining the species in a state ($v(s)$) and the cost of action $c(a)$.

Details of the optimization procedure for MOMDPs are contained in Supporting Information S1. We solved the MOMDP using the SAR-SOP algorithm (Kurniawati et al., 2008). Although SAR-SOP generates an optimal policy, the solutions are difficult to represent due to the

complexity of the solution. We implemented a recent approximation method, alpha-min-fast (Ferrer-Mestres et al., 2021), that allows users to simplify the solution by specifying a maximum number of alpha vectors and a desired precision.

2.5 | Case study

We demonstrated our approach on a case study which sought to determine the optimal management actions to control impacts of feral red fox (*Vulpes vulpes*) predation on an endangered marsupial, the long-footed potoroo (*Potorous longipes*). Red foxes have been implicated in the declines of several ground-dwelling mammals and birds (Mahon, 2009), including the long-footed potoroo (Dexter & Murray, 2009). Foxes are controlled using various methods, including baiting with sodium monofluoroacetate. Although studies have demonstrated the importance of fox removal for recovering some species (Dickman, 1996; Kinnear et al., 1998), there are alternative ways to deliver baits which vary in cost and effectiveness. The duration, frequency and intensity of baiting and whether the baiting is delivered from the ground or aerially all impact the effectiveness of the control programme. The consistency with which baiting programmes are applied is also important, with several studies citing the reinvasion of foxes after intermittent baiting as a key factor in failed control attempts (Gentle et al., 2007; Mahon, 2009). Of particular interest is the effectiveness of aerial baiting compared to ground baiting. Aerial baiting is perceived to be expensive compared to ground baiting but achieves broader landscape coverage since ground baiting tends to cluster baits along existing tracks. We apply our adaptive management framework to investigate the conditions under which aerial and/or ground baiting should be applied.

We set thresholds for the states in consultation with experts in fox control and potoroo monitoring, who referred to previous studies to aid their estimates. The species (potoroo) states were defined as (Arthur et al., 2012; Catling & Burt, 1994; Claridge et al., 2010):

- **Locally Extinct State:** Species detected on 0% of sand plot nights/camera traps.
- **Low State:** Prints detected on 1%-10% of sand plot nights; or detection on >10% of camera traps.
- **High State:** Prints detected on >10% of sand plot nights; or detection on >30% of camera traps.

Threat (fox) states were defined as (Diment, 2010):

- **Low fox density:** <0.3 foxes/km².
- **High fox density:** ≥ 0.3 foxes/km².

Although several factors can be varied to define candidate actions, for both modelling purposes and to reflect operational realities, it is only possible to focus on a subset of fixed action scenarios. We agreed to focus on the effectiveness of aerial compared to ground

baiting. In consultation with experts, we defined scenarios that fixed the frequency, density and coverage of baiting for each action. Specifically, we considered four actions: doing nothing, ground baiting, aerial baiting and combined ground and aerial baiting. Ground baiting was assumed to take place every 4 weeks along tracks, with a bait density of 1 bait/500 m. Aerial baiting was assumed to occur three times per year and achieve a uniform bait density of 4 baits/km². The candidate model set consists of all combinations of hypotheses about whether ground and/or aerial fox baiting will effectively recover potoroos (15 possible models; Figure 1).

Transition probabilities were estimated by species experts. Experts were academics, land managers and pest controllers with experience managing foxes and potoroos. Experts were recruited in consultation with the Saving our Species program. Eleven experts were invited to participate by email, of which 6 provided estimates. Experts were asked questions of the form: (i) given a fox state and action, estimate the probability of remaining in the current fox state in the following year (i.e. $\Pr(p'|p, a)$), and (ii) given a fox state and a potoroo state, estimate the probabilities of potoroo states in the following year (i.e. $\Pr(s'|s, p)$). Estimates were combined by taking the mean of each estimated probability over all experts. Combined estimates were renormalized to ensure that probabilities for each row of the transition matrices summed to 1.

The relative cost of actions was estimated using itemized cost data for ground and aerial baiting of foxes from two previous Saving our Species projects. Ground baiting was assigned a cost of 1. The relative cost of aerial baiting was obtained by dividing the cost for aerial baiting by the cost of ground baiting. Using this method, the relative cost of aerial baiting was set to 1.2. The cost of the combined aerial and ground baiting was assumed to be the sum of its component actions, that is 2.2.

Generating the rewards for the MOMDP requires the benefit of species persistence relative to the cost of action. In this case, the units of benefit are relative values measured against the benchmark of the cost of ground baiting (i.e. $c = 1$), so a benefit of $v = 10$ means that the persistence of the potoroo is valued 10 times as much as the annual cost of ground baiting. Since these data do not exist, we set $v(s) = 0$ if the species is locally extinct and trialled multiple benefit values (10,15,20) otherwise. These benefit values were chosen because they spanned a variety of optimal policies (see Section 3). The benefit value can be varied to determine how the decision maker's perception of extinction risk affects the optimal policy. The effect of the benefit value on the optimal policy can easily be varied by manually changing the 'Benefit of non-extinction' parameter in the Shiny application.

2.6 | Simulating the policy

To test the effectiveness of the optimal policy, we simulated different policies using an assumed ('known') model dynamic. We tested policies that always applied each of the four actions irrespective of the belief ('static' actions) and the MOMDP policy which changes actions based on the observed state component and the belief. Under a known model, we know which actions are effective, so the static policies with

effective actions provide a benchmark to assess the MOMDP policy's performance.

For brevity, we report the results for one known model here (all other models were run; these can be tested using the Shiny application). In the simulations reported below, the simulated threat model assumes that ground baiting and the combined action are effective, but aerial baiting is ineffective (i.e. the 'true' model is F_2). The species model assumes that the species responds negatively to high threat levels but is unaffected by a low threat (i.e. 'true' model is S_2). The performance of ground baiting and the combined action provide an approximate upper limit to the simulation performance. If the MOMDP policy is learning, then it should approximate the performance of the static policies for ground baiting and the combined action.

2.7 | Generating the policy graph

A policy graph is a visual representation of a policy where each node represents the action to apply given the current observable state component and the belief, and edges represent observations. The policy graph for our problem was complex due to the large number of possible models (15), which created large policy tree diagrams that were difficult to plot or interpret. To simplify these, we plotted stationary policy graphs for each possible model dynamic (the stationary belief state was assessed visually from simulations of the belief over time). Transitions for the stationary policy graphs are created from a weighted average of the transition probabilities, where weights are drawn from the stationary belief state. However, the stationary belief may differ depending on the true system dynamics which are unknown until observed. Thus, our simplified policy graphs depict the optimal policy conditional on a 'true' model dynamic. Users can generate the current belief based on existing observations and follow the conditional policy graph for the most likely model given the belief.

3 | RESULTS

Simulating the performance of the optimal MOMDP policy for our case study against static action strategies demonstrated that the MOMDP solution produces good policies despite not knowing the underlying dynamics. The MOMDP policy outperforms all but the combined action in both reducing the threat and preventing species extinction, and has a cumulative reward close to those of the 'effective' ground baiting and combined action policies (Figure 3). The combined action is the optimal action under most simulations for this problem (Figure 4c). The policy graph (Figure 5) shows that the combined action is optimal in the (Low Fox, Low Species) state, but doing nothing is optimal in all other states. Since the system is almost always in the (Low Fox, Low Species) state (Figure 3), the combined action is frequently selected in practice.

Our simulations showed that despite expert optimism that the control actions would reduce fox numbers (Figure 3a), declining trends meant that none of the policies were likely to maintain potoroos in the long term regardless of the action taken (Figure 3b), suggesting




FIGURE 3 Fifteen-year simulations of threat (fox predation) and species (potoroo) under different management strategies with an assumed model dynamic (model F2 S2; see Section 2). Panels (a) and (b) show the simulated threat and species states, respectively; panel (c) shows cumulative rewards. Solid lines represent the mean of 100 simulations; dashed lines show the standard error. The performance of ground baiting and the combined action provide an approximate upper limit for performance since these actions are pre-defined to be effective in the simulation. In panel (a), the performance of aerial baiting is obscured because it is identical to the performance of the ‘Do-nothing’ policy.

pessimism about the effectiveness of fox control as a sole means to recover potoroos. During the elicitation, experts verbally indicated that fox control can be highly effective (regardless of whether aerial or ground control is used), but that it requires sustained intense control to prevent immigration of foxes to the control area, which is rarely achieved in practice (Gentle et al., 2007; Mahon, 2009).

The belief simulations (Figure 4a,b) showed that the correct species response model could be identified rapidly (the correct model emerged after 3 years given initial uniform belief; this was typical when testing across different assumed model dynamics). However, threat models could not be distinguished, other than ruling out model F1 (all actions ineffective) within the first year. This inability to identify the correct threat model from the belief state was typical when testing different assumed model dynamics except when F1 was the true threat dynamic, in which case it was identified after a single year. The optimal policies under different model dynamics almost always selected the combined action, with the exception of model F1 which alternated between trying the combined action and doing nothing.

We tested our case study with different values for the benefit of species persistence (we tested benefit $v(s) = 10; 15; 20$) to identify the decision space where other actions may become optimal. However, for this problem this decision space is small: while we could find some decision uncertainty between doing nothing or the combined action, we did not find the space where ground or aerial actions competed with the combined action in the optimal policies due to a trade-off in cost–benefit. For our case study, the most effective policy is to do nothing when the benefit of species persistence is <10 ; and to always

take the combined action when the benefit of species persistence is 20. In between these values, there is uncertainty about when to manage, but not about which action to take. This result is particular to the reward values used in our case study and may not occur with other parameterizations.

4 | DISCUSSION

Our optimal adaptive management formulation can be applied to any single threat and species combination to discover the optimal management policy from a discrete set of alternative actions. This approach can reduce the cost of implementation for large lists of threatened species because it reuses a common framing and has low information needs. We minimized the number of observable states to reduce the high information requirements needed to parameterize Markov decision processes and make expert elicitation feasible. With two threat states and three species states, the number of elicitation questions is only $2|A| + 12$. These simplifications allowed us to explore all combinations of unobserved system dynamics, ensuring that the true dynamics are in the model set and enabling us to identify the most likely model of how species respond to management.

Our case study showed the promising result that the true species response dynamics model (i.e. species response to threat reduction) can be rapidly identified, and this finding was robust across different response dynamics. While this is encouraging, there are some caveats. For example, we made the simplifying assumptions that threat

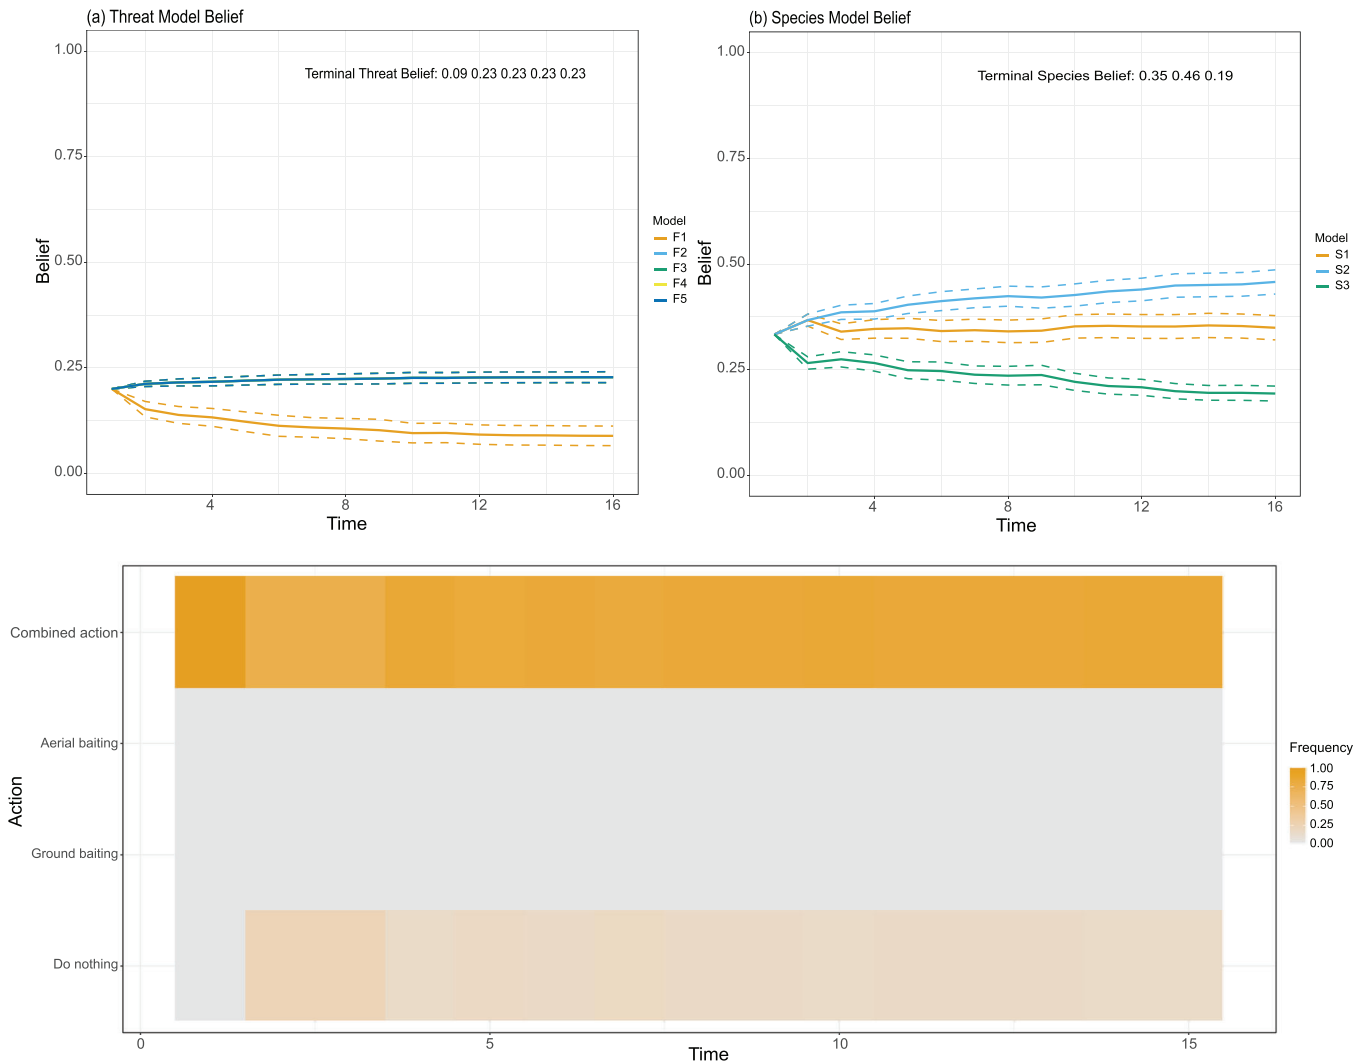


FIGURE 4 Fifteen-year simulated MOMDP policy with an assumed model dynamic. Panels (a) and (b) show the simulated belief for the threat (fox) and species (potoroo), respectively. Solid lines represent the mean of 100 simulations; dashed lines show the standard error. Panel (c) shows the frequency that actions were applied by the optimal policy in the 100 simulations. The simulated beliefs in panel (a) are identical for all threat models except F1; we annotated the terminal threat beliefs to help readers identify which lines overlap in the plot.

and species responses are binary. Environmental, demographic and antecedent conditions, variability in management approaches and interacting threats mean that responses may differ from predictions. When the states are simplified and the goal of management is to avoid extinction (a binary outcome), these simplifying assumptions are likely to be appropriate; however, more complex models would likely be necessary if the management goals were more specific, for example identifying thresholds for threat reduction.

Our case study could not uniquely identify the true model for threat response to management actions. This was likely caused by the assumptions in the model set rather than being a general result. In our threat model set, if any baiting action was effective, then the combined baiting action was also effective. The combined action was also the optimal policy for every species and threat state likely to occur, so each model always chose the same action and observed the same ‘effective’ response regardless of the state visited. Consequently, for this prob-

lem there was never any optimal action that could distinguish between threat models for which at least one baiting action was ineffective (F2–F5). We could not identify the threat model, but in practice we do not really care since the optimal policy (combined action) is the same for all threat models. For the potoroo problem that we modelled, for all states likely to be visited, the best action was either to always apply both ground and aerial baiting or do nothing.

Although the combined action was the most expensive action, there was a low cost differential between the actions compared to the benefits of non-extinction (costs of action ranged from 1 to 2.2; but the benefit of non-extinction was between 10 and 20) and the combined action was more likely to avoid species extinction than the other actions. With such a high value on the benefit of species persistence and little difference in cost, there was no incentive to choose actions with higher probabilities of extinction (i.e. ground or aerial baiting only) that would help to distinguish models. Put differently, there was no

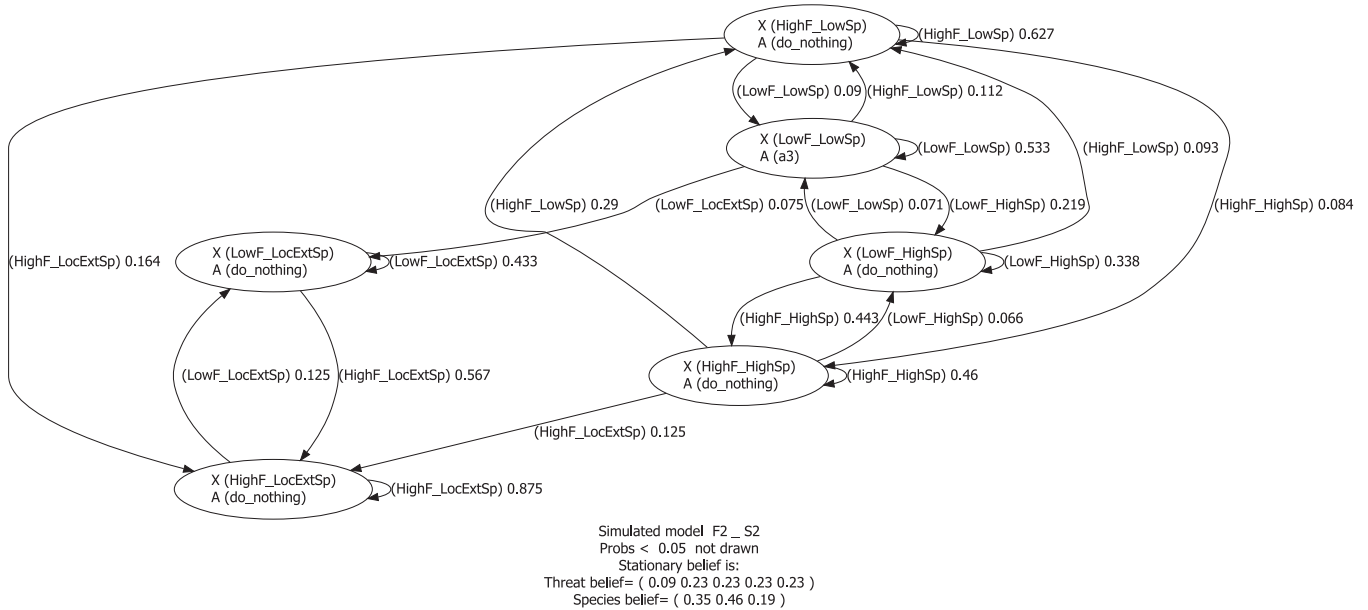


FIGURE 5 Simplified policy graph showing the optimal stationary policy assuming model F2_S2 is the ‘true’ model. Nodes depict the observable states and corresponding optimal actions. Edges show the transition probabilities between nodes. Edges with probability < 0.05 are not shown. Note that the optimal action in all states (depicted inside each node) is to do nothing, with the exception of the (Low Fox, Low Species) state, for which the optimal action is the combined action (a3).

value of information in resolving the adaptive management problem for our case study parameters because the benefits of taking the most effective action were never influenced by the cost (Maxwell et al., 2015). The value of information would be higher where (1) the cost of action was high relative to benefit of non-extinction; (2) the reward function was more nuanced (e.g. different benefits for high and low species, and/or benefits associated with low threat state); or (3) the combined action was not much better than (or equivalent to) one of the other actions. Nonetheless, this finding is instructive: it tells us that there is no decision uncertainty about the most effective action, so resources would be better spent on implementation rather than trying to determine the most effective actions.

This study sought to overcome the challenge of developing custom adaptive management studies for long threatened species lists. Our findings highlight the crucial importance of the reward function, which requires decision-makers to be explicit about their values. In practice, the reward function is rarely expressed, as the existence value of a species is a taboo trade-off (Walshe et al., 2015). Despite this, we postulate that our case study may be illustrative of many threatened species, where there is a narrow decision space in which there is any real uncertainty about the optimal action so that the ‘existence value’ is of little interest since we know the most effective action to take a priori (Maxwell et al., 2015; Moore & Runge, 2012; Nicol et al., 2018). For these species, adaptive management studies about effectiveness are not required, reducing the burden of adaptively managing long threatened species lists. Without this type of rapid modelling approach, the traditional approach is to establish controlled experimental designs, often at significant expense. As well as being useful for selecting optimal actions where decision uncertainty is important,

our approach could be used to determine whether adaptive management is necessary and to prioritize limited funds to where they are most needed.

We applied our approach to a case study where there was a single main threat; however, most species are affected by multiple threats. Although our modelling approach can model multiple threats by adding additional threat states, doing so would increase the complexity of both the elicitation and the solution. We recommend applying our approach to cases where a dominant threat can be identified, for example using threat assessment approaches (IUCN, 2012; Salafsky et al., 2008).

Provided that a single threat can be identified, our approach is not restricted to any specific ecosystem. Most threat types can be modelled; however, our approach assumes stationary dynamics and so may not be suitable for systems that are changing over time, including climate change. Adaptive management approaches for non-stationary dynamics have been developed (Martin et al., 2011; Nicol et al., 2015; Nicol et al., 2014). With some additional assumptions and several extra questions during expert elicitation, most notably regarding the expected rate of change, these approaches could be used with the approach we present here.

Our approach demonstrates how to formulate a general archetype for species–threat adaptive management problems that can rapidly inform decision-making for threatened species, reducing the costs of uptake. General formulations could be created for other classic classes of adaptive management problems, such as optimal harvesting (Memarzadeh et al., 2019) or epidemiology (Shea et al., 2014). As well as simplifying the cost of creating adaptive management solutions for individual applications, creating a library of such general

formulations may present opportunities for further generalizations or theory to further advance the application of adaptive management.

AUTHOR CONTRIBUTIONS

Sam Nicol, James Brazill-Boast, Emma Gorrod and Iadine Chadès conceived the ideas for the study. Sam Nicol and Iadine Chadès designed the adaptive management approach. Sam Nicol, James Brazill-Boast, Emma Gorrod, Hannah Lloyd and Iadine Chadès designed the case study. Sam Nicol collected and analysed the data and led the writing of the manuscript. Sam Nicol and Jonathan Ferrer-Mestres wrote the code and the Shiny application. All authors contributed critically to the drafts and gave final approval for publication.

ACKNOWLEDGEMENT

Human research ethics clearance for the elicitation was granted by CSIRO's Human Research Ethics Framework (project reference 088/19).

CONFLICT OF INTEREST

This work was co-funded by the New South Wales Government's Saving our Species program and CSIRO.

DATA AVAILABILITY STATEMENT

Averaged expert-elicited data files and the cost estimates used in the case study are available from the Figshare data repository (<https://doi.org/10.6084/m9.figshare.16386510>).

ORCID

Sam Nicol  <https://orcid.org/0000-0002-1160-7444>

Iadine Chadès  <https://orcid.org/0000-0002-7442-2850>

PEER REVIEW

The peer review history for this article is available at: <https://publons.com/publon/10.1002/2688-8319.12186>.

REFERENCES

- Arthur, A. D., Catling, P. C., & Reid, A. (2012). Relative influence of habitat structure, species interactions and rainfall on the post-fire population dynamics of ground-dwelling vertebrates. *Austral Ecology*, 37(8), 958–970.
- Atkins, B. D., Jewell, C. P., Runge, M. C., Ferrari, M. J., Shea, K., Probert, W. J. M., & Tildesley, M. J. (2020). Anticipating future learning affects current control decisions: A comparison between passive and active adaptive management in an epidemiological setting. *Journal of Theoretical Biology*, 506, 110380.
- Boettiger, C., Ooms, J., & Mermarzadeh, M. (2018). SARSOP: Approximate POMDP planning software in R. R package.
- Brazill-Boast, J. (2018). Saving our species: A cost-effective, large scale monitoring and evaluation program for threatened species. In S. Legge, N. Robinson, D. Lindenmayer, B. Scheele, D. Southwell, & B. Wintle (Eds.), *Monitoring threatened species and ecological communities* (pp. 225–238). CSIRO Publishing.
- Cain, J. (2001). *Planning improvements in natural resource management: Guidelines for using Bayesian networks to support the planning and management of development programmes in the water sector and beyond*. Centre for Ecology and Hydrology.
- Catling, P., & Burt, R. (1994). Studies of the ground-dwelling mammals of eucalypt forests in south-eastern New South Wales: The species, their abundance and distribution. *Wildlife Research*, 21(2), 219–239.
- Chadès, I., Carwardine, J., Martin, T., Nicol, S., & Buffet, O. (2012). MOMDP: A solution for modelling adaptive management problems. Proceedings of the 26th AAAI Conference on Artificial Intelligence, pp. 267–273.
- Chadès, I., Chapron, G., Cros, M.-J., Garcia, F., & Sabbadin, R. (2014). MDPtoolbox: A multi-platform toolbox to solve stochastic dynamic programming problems. *Ecography*, 37(9), 916–920.
- Chadès, I., Nicol, S., Rout, T. M., Péron, M., Dujardin, Y., Pichancourt, J.-B., Hastings, A., & Hauser, C. E. (2017). Optimization methods to solve adaptive management problems. *Theoretical Ecology*, 10(1), 1–20.
- Claridge, A. W., Cunningham, R. B., Catling, P. C., & Reid, A. M. (2010). Trends in the activity levels of forest-dwelling vertebrate fauna against a background of intensive baiting for foxes. *Forest Ecology and Management*, 260(5), 822–832.
- Dexter, N., & Murray, A. (2009). The impact of fox control on the relative abundance of forest mammals in East Gippsland, Victoria. *Wildlife Research*, 36(3), 252–261.
- Dickman, C. R. (1996). Impact of exotic generalist predators on the native fauna of Australia. *Wildlife Biology*, 2(3), 185–195.
- Diment, A. (2010). *Monitoring the ecological impacts of invasive predator control*. University of Sydney.
- Fackler, P. (2011). *MDPSolve: A MATLAB toolbox for solving markov decision problems with dynamic programming: User's guide*. North Carolina State University.
- Fackler, P. L., Pacifici, K., Martin, J., & McIntyre, C. (2014). Efficient use of information in adaptive management with an application to managing recreation near golden eagle nesting sites. *PLoS ONE*, 9(8), e102434.
- Ferrer-Mestres, J., Dietterich, T. G., Buffet, O., & Chadès, I. (2021). K-N-MOMDPs: Towards interpretable solutions for adaptive management. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 14775–14784.
- Gentle, M. N., Saunders, G. R., & Dickman, C. R. (2007). Poisoning for production: How effective is fox baiting in south-eastern Australia? *Mammal Review*, 37(3), 177–190.
- IUCN. (2012). *Threats classification scheme (Version 3.2)*. <http://www.iucnredlist.org/technical-documents/classification-schemes/threats-classification-scheme>
- Johnson, F. A., Moore, C. T., Kendall, W. L., Dubovsky, J. A., Caithamer, D. F., Kelley, J. R., & Williams, B. K. (1997). Uncertainty and the management of mallard harvests. *The Journal of Wildlife Management*, 61(1), 202–216.
- Kinney, J. E., Onus, M. L., & Sumner, N. R. (1998). Fox control and rock-wallaby population dynamics - II. An update. *Wildlife Research*, 25(1), 81–88.
- Knight, A., Cowling, R. M., Rouget, M., Balmford, A., Lombard, A. T., & Campbell, B. M. (2008). Knowing but not doing: Selecting priority conservation areas and the research-implementation gap. *Conservation Biology*, 22(3), 610–617.
- Kurniawati, H., Hsu, D., & Lee, W. (2008). SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. Robotics: Science and Systems IV.
- Mahon, P. S. (2009). Targeted control of widespread exotic species for biodiversity conservation: The red fox (*Vulpes vulpes*) in New South Wales, Australia. *Ecological Management & Restoration*, 10(s1), S59–S69.
- Martin, J., Fackler, P. L., Nichols, J. D., Lubow, B. C., Eaton, M. J., Runge, M. C., Stith, B. M., & Langtimm, C. A. (2011). Structured decision making as a proactive approach to dealing with sea level rise in Florida. *Climatic Change*, 107(1), 185–202.
- Maxwell, S. L., Rhodes, J. R., Runge, M. C., Possingham, H. P., Ng, C. F., & McDonald-Madden, E. (2015). How much is new information worth? Evaluating the financial benefit of resolving management uncertainty. *Journal of Applied Ecology*, 52(1), 12–20.
- Mermarzadeh, M., Britten, G. L., Worm, B., & Boettiger, C. (2019). Rebuilding global fisheries under uncertainty. *Proceedings of the National Academy of Sciences of the United States of America*, 116(32), 15985.

- Moore, J. L., & Runge, M. C. (2012). Combining structured decision making and Value-of-Information analyses to identify robust management strategies. *Conservation Biology*, 26(5), 810–820.
- Nicol, S., Brazill-Boast, J., Gorrod, E., McSorley, A., Peyrard, N., & Chadès, I. (2019). Quantifying the impact of uncertainty on threat management for biodiversity. *Nature Communications*, 10(1), 3570.
- Nicol, S., & Chadès, I. (2012). Which states matter? An application of an intelligent discretization method to solve a continuous POMDP in conservation biology. *PLoS ONE*, 7(2), e28993.
- Nicol, S., Fuller, R. A., Iwamura, T., & Chadès, I. (2015). Adapting environmental management to uncertain but inevitable change. *Proceedings of the Royal Society B: Biological Sciences*, 282(1808), 20142984.
- Nicol, S., Griffith, B., Austin, J., & Hunter, C. M. (2014). Optimal water depth management on river-fed national wildlife refuges in a changing climate. *Climatic Change*, 124(1), 271–284.
- Nicol, S., Ward, K., Stratford, D., Joehnk, K. D., & Chadès, I. (2018). Making the best use of experts' estimates to prioritise monitoring and management actions: A freshwater case study. *Journal of Environmental Management*, 215, 294–304.
- Ong, S. C. W., Shao Wei, P., Hsu, D., & Wee Sun, L. (2010). Planning under uncertainty for robotic tasks with mixed observability. *The International Journal of Robotics Research*, 29(8), 1053–1068.
- Pascal, L. V., Akian, M., Nicol, S., & Chades, I. (2021). A universal 2-state n-action adaptive management solver. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17), 14884–14892.
- Rout, T. M., Hauser, C. E., & Possingham, H. P. (2009). Optimal adaptive management for the translocation of a threatened species. *Ecological Applications*, 19(2), 515–526.
- Runge, M. C. (2011). An introduction to adaptive management for threatened and endangered species. *Journal of Fish and Wildlife Management*, 2(2), 220–233.
- Salafsky, N., Margoluis, R., Redford, K. H., & Robinson, J. G. (2002). Improving the practice of conservation: A conceptual framework and research agenda for conservation science. *Conservation Biology*, 16(6), 1469–1479.
- Salafsky, N., Salzer, D., Stattersfield, A. J., Hilton-Taylor, C., Neugarten, R., Butchart, S. H. M., Collen, B., Cox, N., Master, L. L., O'Connor, S., & Wilkie, D. (2008). A standard lexicon for biodiversity conservation: Unified classifications of threats and actions. *Conservation Biology*, 22(4), 897–911.
- Shea, K., Tildesley, M. J., Runge, M. C., Fonnesebeck, C. J., & Ferrari, M. J. (2014). Adaptive management and the value of information: Learning via intervention in epidemiology. *PLoS Biology*, 12(10), e1001970.
- Stem, C., Margoluis, R., Salafsky, N., & Brown, M. (2005). Monitoring and evaluation in conservation: A review of trends and approaches. *Conservation Biology*, 19(2), 295–309.
- Walshe, T., Estevez, R., Knapp, Z., & Fonzo, M. D. (2015). Sacred values and filthy lucre-taboo trade-offs in conservation. A perspective on the triage debate. *Australasian Plant Conservation*, 23(4), 20–21.
- Walters, C. (1986). *Adaptive management of renewable resources*. The Blackburn Press.
- Westgate, M. J., Likens, G. E., & Lindenmayer, D. B. (2013). Adaptive management of biological systems: A review. *Biological Conservation*, 158, 128–139.
- Williams, B. (2011). Adaptive management of natural resources—Framework and issues. *Journal of Environmental Management*, 92(5), 1346–1353.
- Williams, B., Szaro, R. C., & Shapiro, C. D. (2009). *Adaptive management: The U.S. Department of the Interior technical guide*. U.S. Department of the Interior, Adaptive Management Working Group.
- Williams, B. K., & Johnson, F. A. (1995). Adaptive management and the regulation of waterfowl harvests. *Wildlife Society Bulletin*, 23(3), 430–436.
- Williams, B. K., & Johnson, F. A. (2013). Confronting dynamics and uncertainty in optimal decision making for conservation. *Environmental Research Letters*, 8(2), 025004.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

The Shiny application can be run from <https://conservation-decisions-lw.shinyapps.io/SpeciesThreatAM/>. A readme user manual for the app, as well as code for the analysis in this paper is available from <https://github.com/nicols02/SpeciesThreatAM>. We recommend running the app from the link above (any OS), however instructions for installation (Windows) and interpretation of the app are available the readme file included in the github repository.

Supporting Information S1: MOMDP optimization

How to cite this article: Nicol, S., Ferrer-Mestres, J., Lloyd, H., Brazill-Boast, J., Gorrod, E., & Chades, I. (2022). A general optimal adaptive framework for managing a threatened species. *Ecological Solutions and Evidence*, 3, e12186. <https://doi.org/10.1002/2688-8319.12186>