Should I Spread My Risk or Concentrate My Efforts: Is Triage of a Subpopulation Ever the Best Decision?

McDonald-Madden E.[1,4], Chades [1,2], McCarthy M. A.[3], Wintle B.[3], and H. P. Possingham[1]

1. The Ecology Centre, The University of Queensland, St Lucia QLD , Australia
2. Unité de Biométrie et Intelligence Artificielle, INRA, BP 27 F-31326 Castanet-Tolosan, France
3. The School of Botany, The University of Melbourne, Parkville VIC 3010, Australia
4. Invasive Animal Cooperative Research Centre, University of Canberra, Australia email: e.mcdonaldmadden@uq.edu.au

Keywords: *Threatened species management, monitoring, partially observable Markov decision process, decision theory*

EXTENDED ABSTRACT

Threatened species often exist in a small number of isolated subpopulations. Given limitations on conservation spending, we must ask the question: should we put all our eggs in one basket and manage the best quality subpopulation or the subpopulation most likely to benefit from management, or should we hedge our bets and manage both subpopulations? A further complexity arises when we consider that most threatened species are cryptic and their presence in an area can be uncertain as a result of the imperfect nature of most detection methods. Managers of cryptic species thus face several dilemmas: if they are unsure whether a species is present in an area or has been extirpated, should they continue to manage the species in that area or instead invest some of their limited resources in surveying to determine if the species is still present (and viable)? How much negative evidence do they need in order to give up and make the decision to cease management? The ecology and conservation literature present little guidance on how to approach such problems, though some analogous problems have been tackled within a decision theory framework (Gerber et al. 2005; Regan et al. 2006; Wilson et al. 2006). Here we build on lessons from these studies and others investigating optimal conservation decision making (Possingham et al. 2001; Dorazio & Johnson 2003) to develop a coherent decision framework for allocating resources between two subpopulations of a threatened species where we are uncertain about the persistence of the species in our management areas. In this problem we must make a decision about how to allocate finite resources to three separate actions in each subpopulation; management, surveying and doing nothing. Management reduces a subpopulation's risk of extinction. Surveying, while not reducing extinction risk, improves our knowledge about whether the species is present, therefore avoiding costly unnecessary expenditure. Both management and surveying cost money and thus the decision to perform either of these actions in a subpopulation will alter the resources available and therefore the success of the action implemented in the other subpopulation. At any point in time managers will have a belief about whether a subpopulation is still extant. In this paper we assess how our optimal decisions change as a function of those beliefs and the time remaining in the management period. The goal of efficient conservation planning and management is to find a resource allocation strategy, or set of actions, that maximises the net expected long-term benefit. Here the optimal strategy involves a trade-off between the persistence of our subpopulations at the end of the management period, and the impact of our decisions on the probability of subpopulation extinction. We pose this problem as a Partially Observable Markov Decision Process (POMDP) and solve a multi time-step scenario using the incremental pruning algorithm (Cassandra et al. 1997). The POMDP algorithm finds an optimal resource allocation each year given the current belief about the state of the species (extant or extinct) in each subpopulation. This paper has two major aims; (i) to extend the framework proposed by Chades et al. (in review) to incorporate two subpopulations of a threatened species, addressing the issue of triage in conservation management, and (ii) to introduce more ecological complexity and realism to the problem by considering subpopulationws of differing habitat quality. We illustrate our findings with a case study using parameters for a Sumatran Tiger (Panthera tigris sumatrae).

1 INTRODUCTION

Worldwide, threatened species have been adversely affected by habitat loss and fragmentation. These changes can be caused by short-term human impact, such as land clearing, or through long-term impacts such as climate change. The resulting habitat fragmentation means that many threatened species tend to exist in a small number of relatively isolated subpopulations (Harrison & Bruna 1999). While a number of ecological theories have provided conservation biologists with general principles for considering the persistence of threatened species remaining in scattered subpopulations, for example island biogeography theory, metapopulation theory and the source-sink paradigm (Andrewartha & Birch 1954; MacArthur & Wilson 1967; Pulliam 1988; Hanski 1999), their use in providing practical and economically astute management plans for conservation management is limited (Possingham *et al.* 2001). Given limitations on conservation spending, we must ask the question should we put all our eggs in one basket and manage the best quality subpopulation or the subpopulation most likely to benefit from management, or should we hedge our bets and manage all subpopulations.

A further complexity arises when we consider that most threatened species are cryptic and their presence in an area can be uncertain as a result of the imperfect nature of most detection methods (MacKenzie & Kendall 2002; Tyre *et al.* 2003). Several enigmatic species have been presumed extinct for long periods before being inadvertently rediscovered (e.g. ivory-billed woodpecker, *Campephilus principalis* (Fitzpatrick *et al.* 2005)). Managers of cryptic threatened species are prone to two sorts of error. It is possible, if not likely, that some populations of a threatened species are being managed even though the species has already disappeared or become functionally extinct in that area, this is the first type of error. What managers need to know is how long they should continue investing in conservation management without strong evidence that the species is still present, and when to shift their resources from saving to surveying for the species? Ultimately, if their belief in the species' existence continues to decline, when should managers surrender resources to another conservation problem? The second possible error is that managers could give up on a species too soon, failing to invest in sufficient surveying to be adequately sure further management is unwarranted.

The problem of how to best allocate conservation resources can be couched in terms of a trade-off between managing and surveying, or doing nothing (surrendering and redistributing resources to other problems). Whether or not to invest scarce management resources and time in surveying may be a difficult decision for managers, though some might argue that expenditure on determining the presence of a potentially viable population is a prerequisite for managing it. Similarly difficult is the decision to give up on the species and stop management, especially if there is possibility that the species may be still extant. We pose this problem as a Partially Observable Markov Decision Process (POMDP) and solve a multi time-step version using the *incremental pruning* algorithm (Cassandra *et al.* 1997). This paper has two major aims: (i) to introduce POMDP as a coherent approach to optimal allocation of resources in a system with partially observable states (e.g., the current status of a cryptic species), and (ii) to extend this framework to a more complex ecological scenario in which decisions need to be made in multiple subpopulations of a threatened species, addressing the issue of triage in conservation management. We illustrate our findings with a case study using parameters for a Sumatran Tiger (*Panthera tigris sumatrae*).

2 METHOD

2.1 The system and ecological complexities

We consider a threatened species that exists in two subpopulations in remnant habitat patches, referred to as population A and population B. In this paper we assume that the subpopulations are isolated from each other and thus that there is no chance of recolonisation once a subpopulation becomes locally extinct. Each subpopulation has associated with it a probability of extinction when it is not managed, p_{0A} and p_{0B}, and a probability of extinction when management is implemented, p_{mA} and p_{mB}. These values are derived from three functional relationships between probabilities of extinction of a subpopulation and resources invested in management (see Figure 1). Each functional relationship represents a different ecological scenario; in this case the level of threat to the subpopulation given it was not managed. This can be interpreted in a number of different ways but in this paper we refer to it as the quality of the habitat in the patch for our threatened species. Two broad questions are assessed: how will we manage if our subpopulations are of equal quality (and thus risk), and how will we act if they differ?

Within the first question we explore the impact of patch quality by investing when both populations have high, medium and low probabilities of extinction when not managed. We then explore the more complex second question comparing patches of high quality to low quality, high to medium, and medium to low quality.

2.2 The decisions

In each subpopulation one of three main decisions can be made (1) to manage, (2) to survey, or (3) to do nothing. Managing a subpopulation reduces the probability of extinction, thus $p_m < p_0$, and surveying allows observations of the system to be obtained with a detection probability (d_s) higher than that of management (d_m), thus $d_s > d_m$. As the budget is traded off between both subpopulations there are six possible overall conservation actions when we consider two subpopulations:

(i) manage both A and B,
(ii) manage A and do nothing in B,
(iii) do nothing in A and manage B,
(iv) manage A and survey B,
(v) survey A and manage B, and
(vi) do nothing in both A and B.

The probability of extinction when we manage a subpopulation and the detection probability when we survey a subpopulation are dependent on the decision implemented in the other subpopulation. That is, if we manage A and B then the budget must be split and thus the probability of extinction in both will be greater than if we were to only manage one subpopulation. However if we were to manage one subpopulation then this would have a low probability of extinction while the second subpopulation would not be managed and have a much higher risk of extinction. Thus there is a clear trade-off between the probability of extinction of an individual subpopulation and our ability to save both subpopulations given a fixed budget, *C*.

2.3 Partially Observable Markov Decision Process (POMDP)

The first step in formulating the conservation resource allocation problem is to define a quantifiable objective. Our objective is to find the optimal allocation of resources given a fixed budget that maximises the expected long term benefits for the conservation of a cryptic threatened species. The final reward associated with a strategy is based on the final state of the system at the end of the management horizon. In this case the state of the system is based on whether subpopulations are *extant* or *extinct*, and thus the system can be characterised by one of four possible states (1) both populations extant, (2) both populations extinct, (3) population A extant and population B extinct, and (4) population A extinct and population B extant. This problem could be solved using traditional stochastic dynamic programming (SDP) algorithms. In doing this, however, we assume that we know exactly what state each subpopulation is in each year. Unfortunately our problem is far more complex as we are concerned with the conservation of cryptic threatened species and are thus likely to be uncertain about the real state of the system. Our optimization method must therefore take into account the incomplete observability of each subpopulation. In other words an optimal decision must depend on the entire history of previous observations and actions rather than on instant observations. We pose this problem as a Partially Observable Markov Decision Process (POMDP) and solve a multi time-step scenario using the incremental pruning algorithm (Cassandra *et al.* 1997).

The POMDP algorithm finds an optimal resource allocation each year given the current belief about the state of the species (extant or extinct) in each subpopulation. Partially Observable Markov Decision Process enhances the SDP model adding a set of observations (absence or presence of the species in each subpopulation) and observation probability matrices that provide the probability of an observation given the current state and the performed action i.e. detection probabilities. Rather than keeping track of the past observation-action history one can use belief states. A belief state is a probability distribution over real states capturing the relative likelihood of being in each of our four overall population states. The computed optimal allocation of resources matches an optimal action to each possible belief state, that is, the policy maps belief states, to actions ($\pi : Beliefs \rightarrow Actions$). The stochastic consequences of reserve-managers actions on the subpopulation are represented with transition probabilities. These transitions are populated with the corresponding probability of persistence/extinction from the appropriate population extinction model (see Figure 1) such that each transition represents the probability of moving from real state s, to real state s' given action a is implemented ($\Pr(s' \mid s, a)$). Probability are also derived for our observations to represent the likelihood of an observation, z, given the real state of the system, s' ($\Pr(z \mid a, s')$. This is based on the detection probability of the species given the resources allocated to that action, surveying (d_s) or management (d_m).

In order to use the optimal solution, decision-makers first need to determine the current belief state of the species. This can be done by answering two simple questions: when is the last time we saw the species in each subpopulation? And what have we decided since? This is the basis on which the POMDP algorithm works, given a belief state ($b(s)$) a decision is selected and the action implemented and an observation, z, obtained. Using this information the previous belief ($b(s)$) is updated to give the current belief state ($b_z^a(s')$) (see Figure 2 for a diagram of this process). Bayes theorem enables us to update the belief state throughout our management time horizon for all possible combinations of actions that could be implemented and observations that may be obtained:

$$b_z^a(s') = \frac{\Pr(z \mid a, s') \cdot \sum_{s \in S} \Pr(s' \mid s, a) \cdot b(s)}{\Pr(z \mid b, a)}.$$

A reward function is specified based on the final state of our system at the end point of our management horizon, N ($R(s,a)$). Here we explore different rewards based on the number of subpopulations remaining extant at the conclusion point of decision-making horizon. The POMDP algorithm iterates backwards through our decision-making horizon calculating at each time step, n, the action, a, that gives the maximum value, $V_n^*(b)$, based on the reward function, $R(s,a)$, the real state transitions, $P(s' \mid a, s)$, the observation probabilities, $P(z \mid a, s')$, and the optimal value from the previous time step, $V_{n-1}^*(b_z^a)$:

$$V_n^*(b) = \max_{a \in A} \sum_{s \in S} b(s) R(s,a) + \gamma \sum_{s \in S} \sum_{s' \in S} \sum_{z \in Z} b(s) P(s' \mid a, s) P(z \mid a, s') V_{n+1}^*(b_z^a)$$

where $n = 1, 2, \ldots, N-1$.

The actions with the maximum value at each time step make up the optimal management policy, π, for a specific ecological scenario.

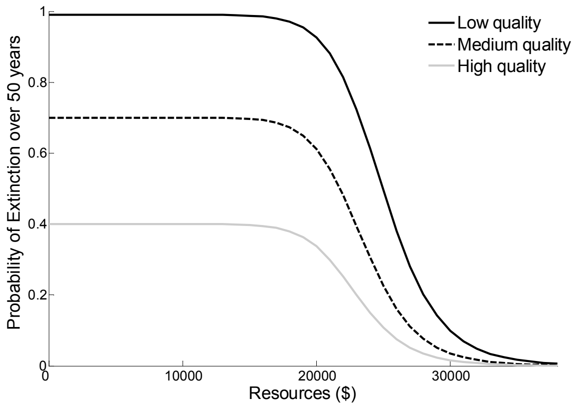

Figure 1. Assumed relationships between probability of extinction in 50 years and management intensity. Each curve represents a habitat quality measure, high, medium, or low. The black curve is derived from Linkie *et al.* (2006) for the Sumatran tiger.

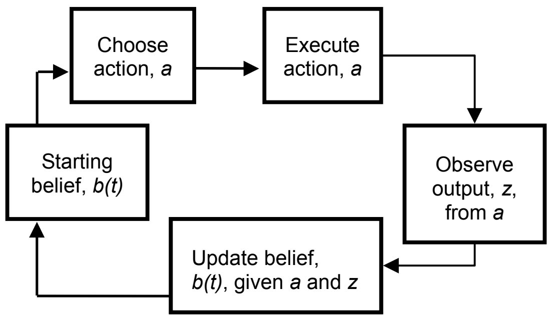

Figure 2. Partially Observable Markov Decision Process iterative belief updating procedure.

3 CASE STUDY – SUMATRAN TIGER

3.1 Problem description

The Sumatran tiger, like most of the tiger species, has suffered from dramatic population decline as a result of reduction in prey abundance, habitat clearance and illegal poaching (Linkie *et al.* 2006). The Kerinci Seblat region of Sumatra is dedicated as a level 1 'tiger conservation unit' (Wikramanayake *et al.* 1998) and significant resources are spent annually to implement management strategies for this population including anti-poaching patrols. Linkie et al. (2006) investigated the effect of resources invested in anti-poaching protection on the probability of losing this population of Sumatran tigers. The current conservation strategy for this species includes reducing the level of poaching by patrolling the population and assessing its status through surveying. Currently about $30,000 are spent annually implementing these two actions with approximately one third of this budget spent on surveying and the remainder on protection measures (Linkie pers. comm.). We interpolated a yearly local extinction probability of 0.058 when the park is managed and 0.100 when it is not (Linkie *et al.* 2006). Similarly, detectability of tigers living in the reserve was estimated at 0.782 when surveyed (d_s) and 0.001 (d_m) when not surveyed. The sensitivity of the optimal strategy to assumed extinction probabilities, detection probabilities and the overall budget was assessed with an extensive sensitivity analysis. We assess the impact of these factors on our initial action in relation to the belief state of each subpopulation.

3.2 Results and comments

Uncertainty is inherent in conservation problems. It is therefore remarkable that this study is one of the first uses of POMDP in a conservation context considering its utility in aiding decision-making under uncertainty. The POMDP framework yields a policy graph describing which action we should implement given our starting belief in the persistence of each subpopulation, the observation gleaned from that action and the time horizon remaining (see Figure 3). We might expect that when our belief in the persistence in both subpopulations is low we should do nothing in both (see Figure 4). Furthermore we may expect that when our belief is high in subpopulation A and low in subpopulation B we should manage A and do nothing in B, and vice versa if our beliefs are reversed. With a large enough budget and high belief in the persistence of both subpopulations we should manage both. All these decisions may be intuitive but what is elusive is the decisions we should make when we have an intermediate belief in the persistence in one subpopulation or both subpopulations. Should we manage one subpopulation and survey the other? If so, which subpopulation should we survey? How will decisions change if our budget increases or decreases, or surveying becomes more efficient, or if our subpopulations have differing habitat quality? Another component that is unclear is how our decisions should change through time. In this work we derive the answers to these questions from the policies obtained using POMDP. Ultimately we address the question of whether to spread our risk or concentrate our effort when managing threatened species where our belief in the persistence of subpopulations guides our actions.

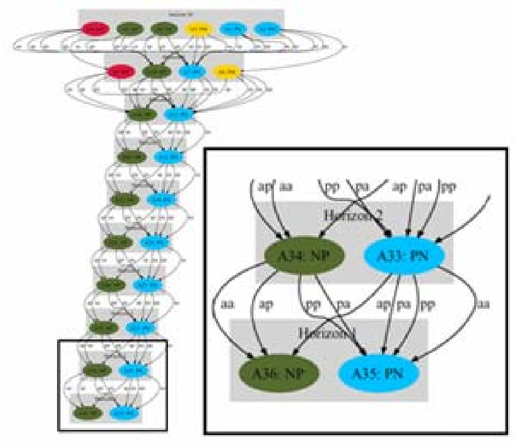

Figure 3: Example policy graph from Partially Observable Markov Decision Process for two subpopulations of Sumatran tiger (*Panthera tigris sumatrae*) with identical extinction risk based on the low quality extinction curve (see Figure 1). Inset shows specific details of graph, nodes represent starting belief states, colours represent different actions, grey blocks represent time steps (descending from top), and arrows show which decision to make next given a particular observation is observed (aa = absent in A and absent in B, pp = present in A and present in B, etc.).

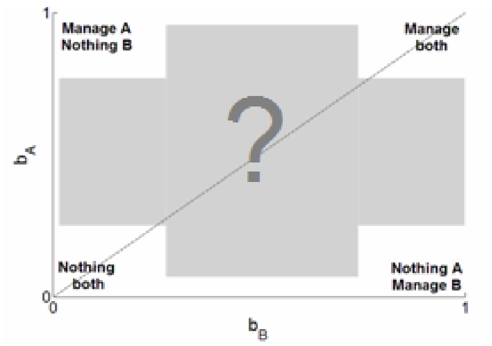

Figure 4. Possible expression of results for an individual time horizon as a relationship between our belief in subpopulation A, b_A, and our belief in subpopulation B, b_B. The shaded area is where we are uncertain of how to act; this will be elucidated by using a POMDP.

4 ACKNOWLEDGEMENTS

This research was funded by an ARC Discovery grant to HPP and MAM, an ARC Federation Fellowship to HPP, and INRA travelling fellowship to IC. EM PhD Studentship was supported by the Invasive Animals CRC within the Detection, Prevention and Mitigation Program and by a MYQRS scholarship from the University of Queensland. We are grateful to E. Game and two anonymous reviewers for helpful comments on this paper and P. Baxter for insightful discussion and advice.

5 REFERENCES

Andrewartha H.G. & Birch L.C. (1954) *The Distribution and Abundance of Animals*. The University of Chicago Press, Chicago, Illinois.

Cassandra A.R., Littman M.L. & Zhang N.L. (1997) Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In: *Proceedings of the International Conference on Uncertainty in Artificial Intelligence*.

Chades I., McDonald-Madden E., McCarthy M.A., Wintle B.A., Linkie M. & Possingham H.P. (in review) Save, survey, or surrender: Optimal management of a cryptic threatened species.

Dorazio R.M. & Johnson F.A. (2003) A comparison of optimal decisions computed using bayesian updating and stochastic dynamic programming. In: *Adaptive Management conference Series*, Laurel, Maryland, USA.

Fitzpatrick J.W., Lammertink M., Luneau Jr M.D., Gallagher T.W., Harrison B.R., Sparling G.M., Rosenberg K.V., Rohrbaugh R.W., Swarthout E.C.H., Wrege P.H., Swarthout S.B., Dantzker M.S., Charif R.A., Barksdale T.R., Remsen Jr J.V., Simon S.D. & Zollner D. (2005) Ivory-billed woodpecker (Campephilus principalis) persists in continental North America. *Science*, 308, 1460-1462.

Gerber L.R., Beger M., McCarthy M.A. & Possingham H.P. (2005) A theory for optimal monitoring of marine reserves. *Ecology Letters*, 8, 829-837.

Hanski I. (1999) *Metapopulation Ecology*. Oxford University Press, New York.

Harrison S. & Bruna E. (1999) Habitat fragmentation and large-scale conservation: what do we know for sure? *Ecography*, 22, 225-232.

Linkie M., Chapron G., Martyr D.J., Holden J. & Leader-Williams N. (2006) Assessing the viability of tiger subpopulations in a fragmented landscape. *Journal of Applied Ecology*, 43, 576-586.

MacArthur R.H. & Wilson E.O. (1967) *The Theory of Island Biogeography*. Princeton University Press, Princeton, N.J.

MacKenzie D.I. & Kendall W.L. (2002) How should detection probability be incorporated into estimates of relative abundance? *Ecology*, 83, 2387-2393.

Possingham H.P., Andelman S.J., Noon B.R., S. T. & Pulliam H.R. (2001) Making smart conservation decisions. In: *Conservation Biology: research priorities for the next decade* (eds. Soule ME & Orians GH). Island Press, Washington.

Pulliam H.R. (1988) Sources, sinks and population regulation. *The American Naturalist*, 132, 652-661.

Regan T.J., McCarthy M.A., Baxter P.W.J., Panetta F.D. & Possingham H.P. (2006) Optimal eradication: when to stop looking for an invasive plant. *Ecology Letters*, 9, 759-766.

Tyre A.J., Tenhumberg B., Field S.A., Niejalke D., Parris K. & Possingham H.P. (2003) Improving precision and reducing bias in biological surveys: Estimating false-negative error rates. *Ecological Applications*, 13, 1790-1801.

Wikramanayake E.D., Dinerstein E., Robinson J.G., Karanth U., Rabinowitz A., Olson D., Mathew T., Hedao P., Conner M., Hemley G. & Bolze D. (1998) An ecology-based method for defining priorities for large mammal conservation: The tiger as case study *Conservation Biology*, 12, 1427-1427.

Wilson K.A., McBride M.F., Bode M. & Possingham H.P. (2006) Prioritizing global conservation efforts. *Nature*, 440, 337-340.