Optimally Managing Oscillating Predator-Prey Systems

[1,2]Michael Bode and [1,2]Hugh Possingham

[1]Department of Mathematics, The University of Queensland, Brisbane, 4079, Australia.
[2]The Ecology Centre, The University of Queensland, Brisbane 4079, Australia.

Keywords: *Predator-prey cycles; optimal management; stochastic dynamic programming*

EXTENDED ABSTRACT

Many interacting predator-prey populations have a natural tendency to exhibit persistent limit-cycle or damped oscillations, especially in the presence of environmental stochasticity. The restriction of habitats, and the resultant compression of ecosystems into small conservation reserves can induce these oscillations in previously stable predator-prey relationships.

During the course of these oscillations, the abundance of both species regularly decreases to low levels. At these times, the inherent stochasticity of the environment will result in a high probability of species extinction. In these situations, actions should be taken by management authorities to stabilise the system and avoid extinction. In many practical cases, funding limitations restrict the actions available to managers, and so we assume that the only options available to a manager are to cull individuals from the two populations. Operating within these constraints, we show in this paper that strategic culling of organisms can greatly reduce the probability of a species becoming extinct.

We use the mathematical technique of Stochastic Dynamic Programming to determine the optimal management strategy for these oscillation-prone populations. Application of the optimal strategy results in a much more persistent system, with only a small number of interventions being necessary. These methods can be applied to many different species, and can incorporate more complex system dynamics without a significant increase in computational time. The optimal strategies that result from our analyses offer general insights into how such systems should be managed.

1 INTRODUCTION

It is well known that predator and prey populations of certain species exhibit natural cyclic behaviour: their abundances do not settle down to particular values, but change in a manner that repeats through time. The most famous example are the Canadian lynx and snowshoe hare populations in North America (Ginsburg & Colyvan 2004). The causes of natural limit-cycles in these populations are contentious, and the problem has received considerable theoretical and empirical attention (see: Turchin 2003, Ginsburg & Colyvan 2004 for reviews).

Population cycles can also be created by human interference in naturally stable systems. When the habitat available to an ecosystem is reduced, its species are "compressed" into a smaller region, and limit cycles can develop (Barnes 1983). Typically, high predator density leads to low prey abundance, followed by a rapid decline in the predator population through starvation. This low predator abundance allows the prey population to increase, at which point the predators also begin to increase, and the cycle begins again. In a large system these cycles are dampened by spatial factors (Jansen & de Roos 2000), but in small systems such as many conservation reserves, they can be very dangerous to species' persistence, especially the predator. When the abundance of a species is very low, a few unfortunate events (e.g. accidental deaths, poaching) can be devastating.

The number of areas designated as conservation reserves is set to increase in the coming decades (James et al. 2001), and these areas will represent islands of safety amid an otherwise hostile landscape. The gradient of habitat quality across the boundary of the reserve will act to funnel unnaturally high numbers of organisms into these reserves. Reserve managers would obviously like to ensure the persistence of the protected species, but while it is possible to manipulate the abundances of the species, options that will have a positive effect on numbers (e.g. captive breeding, translocation) are generally too expensive, as are non-lethal measures that will reduce growth rates (e.g. sterilisation, contraception, isolation). Frequently the managers' options are limited to culling. The idea of killing organisms to prevent their extinction seems quite contradictory – indeed, if such measures are applied at the wrong time, the effects can be negative to species persistence. However, with careful analysis, judicious culling can benefit ecosystem persistence.

Whether or not culling should occur at a particular time will depend heavily on the abundances of both species. The non-linear time-progression of the system means the effects of culling will not be immediately obvious, and will not necessarily be intuitive (e.g. removing individuals of a particular species might lead to that species developing greater abundances than would otherwise have occurred). Taking into account the progression of the oscillatory system dynamics is crucial to developing management strategies that will best ensure the persistence of both predator and prey. This problem is an exercise in dynamic Optimal Control. Although the problem of when to cull individuals has been discussed, the optimal control strategy has not been explicitly formulated. We use dynamic decision theory to formulate this management problem explicitly, and then apply Stochastic Dynamic Programming (SDP) to determine the management strategy that will best ensure the persistence of the predator-prey system. SDP will formulate a strategy that will take into account the system's inherent unpredictability, and the costs and benefits associated with the fluctuating populations and their control.

2 METHODS

The Population Model: To understand the effects of management interventions on the time-evolution of the predator-prey system, we must first devise a quantitative ecosystem model that includes important system dynamics. We set up a general predator-prey model, where the dynamics of the prey population N_t are deterministic, and those of the predator population P_t are stochastic. In reality both species' dynamics will be unpredictable, but stochasticity will be most important to predator dynamics as their population will be much smaller. We keep track of the actual numbers of each species ($N_t \in \{0, 1, 2, \dots\}$, $P_t \in \{0, 1, 2, \dots\}$). To allow the application of SDP, our model is discrete in time ($t = 1, 2, \dots$). The prey population grows logistically, with carrying capacity K and growth rate r.

$$N_{t+1} = N_t \left(1 + r(1 - N_t/K)\right). \tag{1}$$

The predator population grows stochastically. Each timestep the predators seek out prey, a process that becomes more difficult as prey become more scarce. To reflect this, the amount of prey caught by each predator is drawn from a uniform distribution between zero and $C_{\max} N_t/K$. The upper bound of this distribution depends on prey density: if the prey numbers are at their carrying capacity, each predator can catch as many as $C_{\max}$ prey per timestep. As prey become scarce, this maximum decreases. The number of prey caught by the i^{th} individual is thus selected from the distribution:

$$C_t^i \sim \text{Uniform}(0, C_{\max}/N_t K).$$

This catch is then rounded to the nearest whole number of individuals.

The number of prey that each predator catches defines its dynamics. If the predator catches very few prey, it will starve to death. The predator may catch enough prey to survive, but not enough to support reproduction. If the predator is fortunate enough to catch above a certain number of prey, they will have sufficient energy reserves to reproduce, and will give birth to one offspring. The predator dynamics thus follow the equation:

$$P_{t+1} = P_t + \sum_{i=1}^{P_t} b_i - \sum_{i=1}^{P_t} d_i, \tag{2}$$

$$b_i = \begin{cases} 0 & \text{if } C_t^i < C_{\text{birth}}, \\ 1 & \text{if } C_t^i \geq C_{\text{birth}}, \end{cases}$$

$$d_i = \begin{cases} 0 & \text{if } C_t^i > C_{\text{death}}, \\ 1 & \text{if } C_t^i \leq C_{\text{death}}, \end{cases}$$

where b_i represents the reproduction of the i^{th} individual, and d_i indicates whether the i^{th} individual died from starvation. The numbers of predators and

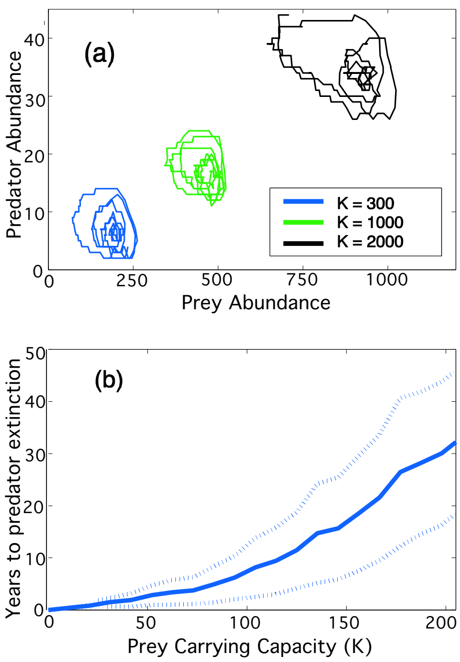

Figure 1. (a) Simulated population cycles in phase space for different prey carrying capacities K. The probability of population extinction decreases as K increases. (b) Expected time to extinction as a function of K. Solid line indicates the mean, dotted lines indicate ± 1 standard deviation.

prey will depend primarily on the habitat's carrying capacity – the system is controlled from the bottom-up. The importance of stochasticity will likewise depend on the carrying capacity. The predator population will increase as the prey population increases, and once the number of predators becomes very large, the net effect of the random fluctuations will become relatively unimportant. The possibility of the fluctuations causing a stochastic extinction in a large population becomes remote as the prey carrying capacity increases, as demonstrated in figure 1b.

Extinction is most likely to occur when the predator population is at a low abundance level. At this point, a series of unfortunate events (e.g. unsuccessful predator foraging, or a large number of random deaths) could result in the predator population declining to zero (see figure 2).

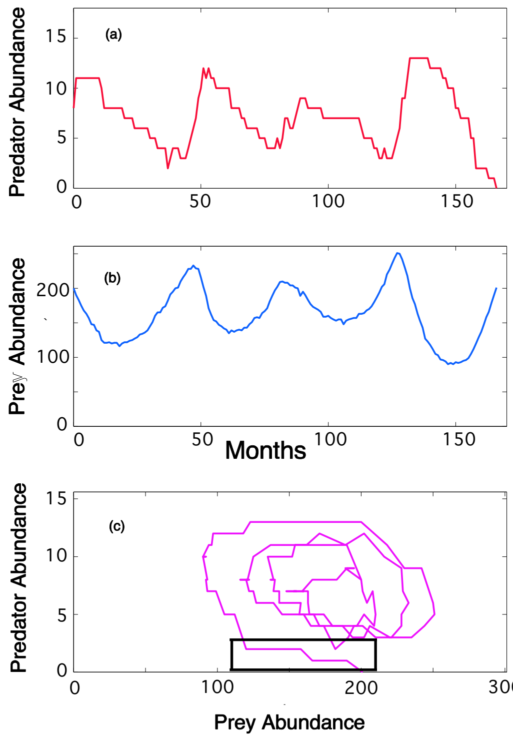

Figure 2. Typical dynamics displayed by a single realisation of a stochastic predator-prey system, in the lead-up to the extinction of the predator population, highlighted by the black box. (a) the predator abundance, (b) the prey abundance, (c) the stochastic cycles displayed in the phase plane.

The System Dynamics as a Markov Process Our optimisation method, Stochastic Dynamic Programming, requires that the system dynamics be expressed as a discrete time Markov process, where the system's evolution is defined by the probability of transition between system "states". In this predator-prey system, a state is defined by the abundance of predators and prey. The changes that the system will undergo – births, deaths, predation – depend only on the two species' abundances, and so we use these two values to define a state. Regardless of the carrying capacity, the potential population of both species is theoretically unlimited, however very large populations of either species are not sustainable for any length of time. We therefore place a "cap" on the abundance of predators (P_c), and on the abundance of prey (N_c). These caps ensure that our state-space is finite, but we will set them at values much higher than the system would naturally encounter, so they do not artificially interfere with the system dynamics.

We use the notation S_i to represent the i^{th} state of the system. State i is defined as:

$$\mathrm{S}_i \equiv \left[N_i = (i-1)(\operatorname{mod} N_c + 1), P_i = \lfloor \frac{i}{N_c+1} \rfloor \right].$$

Where $a(\operatorname{mod} b)$ is the remainder of a divided by b, and $\lfloor c \rfloor$ is the greatest whole number that is less than c (the *floor* of c). There are a total of $S_{max} = (N_c + 1)(P_c + 1)$ states. If $N_c = 10$, then some example states would be:

$$\begin{aligned} \mathrm{S}_1 &\equiv [N(t) = 0, P(t) = 0], \\ \mathrm{S}_2 &\equiv [N(t) = 1, P(t) = 0], \\ \mathrm{S}_{14} &\equiv [N(t) = 2, P(t) = 1], \\ \mathrm{S}_{20} &\equiv [N(t) = 8, P(t) = 1], \\ \mathrm{S}_{40} &\equiv [N(t) = 6, P(t) = 3]. \end{aligned}$$

The predator-prey dynamics are inherently stochastic, as the future evolution of the system is not known with certainty. However, given that the system is in a particular state, its subsequent state will one of a number of possible states (some more likely than others). The dynamics of such a Markov process are governed by a transition matrix $\mathbf{T}$, whose elements $\mathbf{T}_{ij}$ represent the probability that the state S_i will evolve into the state S_j in the next timestep.

Instead of defining $\mathbf{T}$ *a priori*, it can be constructed from repeated simulations of the predator-prey equations and 2. After running a simulation, the observed transitions are stored in a matrix. When the entire simulation has been thus recorded, the matrix is normalised by the sum of its rows. This process is repeated many times, and the mean of all the simulations is calculated. After sufficient repetitions, the mean of these matrices converges to the transition matrix $\mathbf{T}$. The resultant matrix encapsulates the probabilistic time-evolution of the predator-prey system, and is quite simple to construct. Our methods can therefore be easily applied to more complex multi-species ecosystems.

Application of SDP: With this transition matrix, we can obtain the optimal management scheme by the application of a technique from Operations Research called Stochastic Dynamic Programming. SDP has been applied to many problems in dynamic decision theory (e.g. behavioural ecology, Mangel & Clark 1988; management of ecological communities, Richards et al. 1999). Given a range of decisions, SDP will yield an optimal management strategy for our predator-prey system in the form of an state-dependent "optimal decision space", which identifies the best decision for the manager to take in every state. SDP calculates the strategy that will best achieve the stated objective in a set timespan (T years), which must first be defined. Our objective cannot be simply: "we would like to ensure the persistence of the predator population". We must also specify over what timespan e.g. "we would like to ensure that the predator species persists for T years". Theoretically this means that the resulting optimal solution is also time-dependent, and that the optimal decision space to ensure persistence over T years will not necessarily be optimal to ensure persistence over a different time interval. In practice however, as T is allowed to increase the solutions converge rapidly to a time-independent solution (Richards et al. 1999).

Of course, the best decision for a manager to make will depend on what needs to be achieved. Different objectives will result in different optimal strategies. The manager's primary objective is that the predator species persists over some time period. This objective implicitly ensures the prey's persistence, as the prey population must be available for the predator to persist. Each state of the system will have a value assigned to it that reflects how well it fulfills the objective. If we were interested only in the *persistence* of the predator population, every state with a positive number of predators would have the same value assigned to it. However if each predator is deemed to be valuable, we would make the value of each state dependent in some way on the quantity of predators. For example, if our predator population were elephants and attracted tourism income, we might consider higher densities of elephants (and thus a higher likelihood of tourists observing them) more valuable. We will accordingly define the value of the i^{th} state as:

$$\psi(i) = \mathrm{V}_P P(i) + \mathrm{V}_N N(i),$$

the value of each predator (V_P), multiplied by the predator abundance ($P(i)$), added to the value of each prey individual (V_N), multiplied by the prey abundance in that state ($N(i)$). In SDP, we can assess value at the terminal time alone, or throughout the entire timeseries. In the former case, we do not particularly desire large numbers of predators throughout the timeseries, only at the terminal time T. In the latter case, a high predator abundance throughout the timeseries is desirable. Continuing with our elephant example, as tourism is occurring continually, we must add value each year.

To best achieve the management objective, at each timestep the manager must make a decision. In SDP we must outline the full set of decisions available to the manager, and the costs that will be incurred by taking each one. All of our decisions involve either the culling of prey individuals, or the culling of predator individuals. We will further limit the number to be culled. We label our decisions as follows:

1. Do nothing
2. Cull a single predator
3. Cull two predators
4. Cull five prey
5. Cull ten prey

Of course, many more options could be defined. The manager might be able to cull more than just two predators, for example, or might be able to cull both predators and prey simultaneously. These actions will each have an associated cost that will act to dissuade managers from applying them unless they are necessary. These costs will reduce the value of the system by a constant amount each time they are applied. We define the cost associated with the k^{th} management decision as $c(k)$. The effect of each decision will be to alter the transition matrix $\mathbf{T}$, yielding a new matrix $\mathbf{T}(k)$ – after intervention, the predator-prey system will not evolve in the same way. At each timestep, one of the five decisions must be taken (although decision 1 represents taking no action).

The optimal strategy is determined by application of the the "dynamic programming equation" (DPE), defined according to the system value $\psi(i)$, the intervention costs $c(k)$, and the transition matrices $\mathbf{T}(k)$:

$$\mathrm{V}(i,t,T) = \max_{k\in[1,5]}[\psi(i)-c(k)+\sum_{j=1}^{S_{\max}}\mathbf{T}_{ji}(k)V(j,t+1,T)]. \quad (3)$$

The dynamic programming equation is applied as an iterative optimisation algorithm. Following Bellman's principle of optimality, the optimal action to take at time t, when the system is in state i, is the decision k that maximises the net value of the system in the next time step $(t+1)$, *given that all subsequent decisions are also optimal*. This net value incorporates the value of the state $\psi(i)$, minus the cost of the decision k. It then adds the value of those states that the current state i will evolve into, given that the intervention k has occurred, weighted by their probability of occurrence through the transition matrix $\mathbf{T}(k)$. The DPE is repeatedly applied from the terminal time T, backwards to the initial time 0. We must therefore define the value of the system at the terminal time (where our algorithm begins), T.

$$\mathrm{V}(i,T,T) = \psi(i).$$

Problems typically only arise when the state-space is particularly large, as it then becomes computationally difficult to store the necessary transition matrices $\mathbf{T}(k)$. To avoid this, we have binned the prey abundance into multiples of 5, markedly decreasing the optimisation runtime. Comparisons with optimal solutions without this binning revealed negligible differences in the optimal strategy. For a more thorough explanation of SDP, see Mangel & Clark (1988) or Intriligator (1971)

Example Parameters: To illustrate the methods outlined above, we apply SDP to a predator-prey system with a given parameter set. The carrying capacity of the habitat is set as $K = 250$, with a growth rate of $r = 0.15$, or 15 percent each year. Each predator can eat a maximum of $C_{\max} = 10$ prey each year; if the predator does not catch two prey, it will die ($C_{\text{death}} = 1$); if the predator catches 4 or more prey, it will reproduce ($C_{\text{birth}} = 4$). We limit the transition matrix to $N_c = 250$ prey individuals, and $P_c = 20$ predator individuals. Each predator individual produces an annual revenue of 1, prey individuals provide no revenue. The cost of culling a predator organism is set at 50, higher than the cost of culling a prey organism, 5.

3 RESULTS

Figure 2 is a single simulation of a predator-prey system with the above parameter values. The application of SDP yields an optimal decision space: a set of state-dependent decisions that will result in a timeseries with the maximal expected value. Figure 3 shows the optimal decisions for the example system. The resultant decision space can be better interpreted if an example oscillation of the system is superimposed upon it, as shown by the thin black line. The population dynamics represented by this simulation have not been affected by management, and the predator population quickly goes extinct. We see immediately that in many of the states the

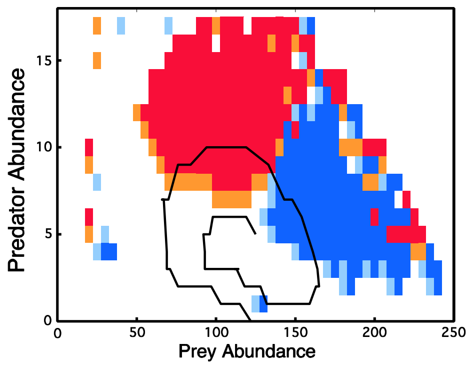

Figure 3. The state-dependent optimal decisions for the example system defined. Lighter red indicates culling a single predator, darker red indicate culling two predators. Lighter blue indicates culling five prey, darker blue indicates culling ten prey. The black line is a realisation of an unmanaged system, superimposed upon the decision space.

optimal decision is *not* to intervene, but to let the system evolve unaltered. Some of these states are in the interior of the simulated population cycle, and thus represent quite stable predator-prey abundances.

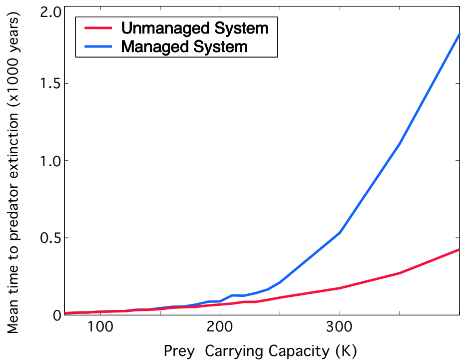

Figure 4. Expected time to predator extinction for SDP-managed and unmanaged systems, as a function of the prey carrying capacity. The SDP-managed system becomes notably more persistent than the unmanaged system as K increases.

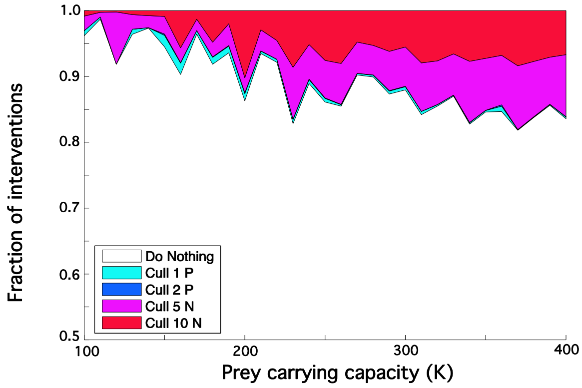

Figure 5. The simulated frequencies of a set of possible management decisions as a function of the carrying capacity. Frequencies are the average of decisions taken in 5000 independent simulations. In most years, management intervention is not required, and culling 2 predator organisms almost never occurs

Taking active management decisions when the system is in one of these states would not be particularly beneficial, as the populations are not in any immediate danger of extinction. The cost of management actions would act to further dissuade interference. In states where the abundance of both species is particularly low, the optimal management action is also to do nothing. This reflects the restrictive range of decision options available to the manager. Although there is some danger of extinction, none of the available management options would be beneficial.

In regions of the state-space where the predator abundance is particularly high, and the prey abundance is quite low, the optimal decision is still to do nothing. Although this may seem counter-intuitive, it again reflects the restrictive management options. Although in these states the predator population is in imminent danger of extinction through starvation, it is not optimal to cull one or two predators. Although taking such action may decrease the probability of the predator population becoming extinct, the system state that would result from taking action is still a high probability of extinction, and this small reduction is not sufficient to outweigh the costs involved. SDP takes into account the costs of the available actions, and weighs them against the potential benefits of action. In these cases the probability of extinction is so high that inaction is optimal.

For the chosen parameter values, there is a large region where it is optimal to remove individuals from the predator population, and another region where it is optimal to remove individuals from the prey population. The superimposition of a simulated population trajectory can help us understand the placement of these regions. The state of the system naturally cycles in an anti-clockwise direction. The actions of the manager are limited to "pushing" the oscillations to the left, by removing prey individuals, or "pushing" them downward, by culling predators. The persistence of the populations will best be achieved by containing the population cycles – i.e., by reducing their radius. It is when the radius of these cycles becomes too large that there is a danger of species extinction, when the cycle intersects with one of the axes. The decision to cull individuals is always taken in a region of the state space where the resultant "push" will reduce the amplitude of the population cycle. If these decisions were taken at any other time, they would result instead in larger amplitude oscillations, and a higher probability of extinction. There are very few places where it is optimal to cull only a single individual; most decisions are either to do nothing, or to cull the maximum number of individuals possible. This implies that the oscillations very quickly become dangerous if they begin to build up a larger amplitude.

4 MANAGEMENT CONCLUSIONS

To determine the effectiveness of the SDP-created optimal management strategy, we ran parallel simulations where one system was managed according to the SDP strategy, and the other was unmanaged, and compared the expected time until the predator became extinct. We repeated this for increasing values of the prey carrying capacity K, to determine the importance of the habitat quality. The results are shown in figure 4. Managing the system becomes increasingly beneficial as the prey's carrying capacity K increases. For low K values, the predator-prey systems are so unstable that the limited management options available cannot prevent the inevitable rapid extinction of the predators. For larger K systems, the management options increasingly allow longer-term persistence. This increased persistence has a positive effect on the costs associated with management.

Although taking action costs money, these predator-prey systems are providing ongoing revenue, and can thus conceivably generate profit. The costs necessary to ensure a persistent system are far less than the profit it creates, and so the SDP-managed system is far more profitable, as well as more persistent. During the course of the managed system simulations, we recorded the interventions made by the SDP-trained managers. The proportion of timesteps that each decision was taken is shown as a function of the prey carrying capacity K in figure 5. The other parameters in these systems were the same as our previous system. We can see that interventions occur very infrequently – no management intervention took place in over 80% of the years, for all of the different ecosystem sizes.

As K increases, the culling of large numbers of prey becomes more prevalent in the manager's interventions. This reflects the increasing population size, rather than the need to remove a higher percentage of the individuals. Most of the interventions called for involve culling prey species, which is a fortunate coincidence for the managers. Our representation of the effects of culling on each species is simplistic, in particular for the predator species. Prey species have evolved to being harvested to some extent, by the predators. Predator species, on the other hand, may not have been historically subject to culling, being situated near, or at the apex of their food web. Elephants, for example, have a social structure that can be significantly disrupted by culling. The low prevalence of predator culling is therefore beneficial. Ideally however, the effects of culling a specific predator species would be included in decision-augmented transition matrices $\mathbf{T}(k)$.

In states where an active decision was necessary, intervention mainly involved culling the minimum number of individuals possible, rather than the higher option. This is despite the relatively small number of states where culling the lower number of individuals was the optimal decision. The more severe culling states were typically avoided with the early application of judicious culling. There are two conclusions we should draw from this. First, the prevalence of particular interventions in decision-space does not necessarily reflect how frequently they will be implemented by a manager. For example, although in much of the state space, two predator organisms should be culled, this action is infrequent in practice, as previous interventions keep the system away from this state. This leads to the second point: in dynamic systems, taking pre-emptive action that acknowledges possible future developments can help prevent their occurrence. Management that does not quantitatively include the system dynamics would find this difficult to achieve. As Barnes (1983, pp. 139) notes, "by the time [the problem of dangerously high levels of predators] is recognised, the time of maximum cost-effectiveness of culling has passed." It is therefore crucial that the future evolution of the system dynamics are considered when management strategies are formulated.

Optimal control of dynamic systems is far more complicated than static control, especially when the options available to management are heavily constrained. Nevertheless, we have shown that if the objectives of management, and the dynamics of the stochastically oscillating predator-prey system are correctly formulated in a dynamic decision theory framework, an optimal management strategy can be determined upon application of Stochastic Dynamic Programming. Our model was devised to reflect the dynamics of a general predator-prey system, and the methods outlined can be applied to any particular system. The general conclusions of the model will be applicable to many similar systems, and can be used in the absence of specific modelling, to guide management actions.

5 BIBLIOGRAPHY

Barnes, R.F.W. (1983) The elephant problem in Ruaha National Park, Tanzania, *Biological Conservation* 26, 127–148.

Ginsburg, L., and M. Colyvan (2004) *Ecological orbits: how planets move and populations grow*, Oxford University Press, 168 pp., New York.

Intriligator, M.D. (1971) *Mathematical optimisation and economic theory*, Prentice-Hall, 508 pp., New Jersey.

James, A., K.J. Gaston, and A. Balmford (2001) Can we afford to conserve biodiversity? *Bioscience*, 51, 43–52.

Jansen, V.A.A, and A.M. de Roos (2000) The role of space in reducing predator-prey cycles. In: *The Geometry of Ecological Interactions*, (Eds. U. Dieckmann, R. Law, J.A.J. Metz). Cambridge University Press, 183–201, Cambridge.

Mangel, M., and C.W. Clark (1988) *Dynamic modeling in behavioural ecology*, Princeton University Press, 320 pp., Princeton.

Richards, S.A., H.P. Possingham, and J. Tizard (1999) Optimal fire management for maintaining community diversity, *Ecological Applications*, 9, 880–892.

Turchin, P. (2003) *Complex population dynamics: a theoretical/empirical synthesis*, Princeton University Press, 264 pp., Princeton.