

THE USE OF STOCHASTIC DYNAMIC PROGRAMMING IN OPTIMAL LANDSCAPE RECONSTRUCTION FOR METAPOPOPULATIONS

MICHAEL I. WESTPHAL,^{1,2,5} MARCUS PICKETT,³ WAYNE M. GETZ,² AND HUGH P. POSSINGHAM^{1,4}

¹The Ecology Centre, Department of Zoology and Entomology, University of Queensland,
St. Lucia, Queensland 4072, Australia

²Department of Environmental Science, Policy, and Management, University of California, Berkeley, California 94720 USA

³Conservation Council of South Australia, Adelaide, South Australia, Australia

⁴Department of Mathematics, University of Queensland, Queensland, Australia

Abstract. A decision theory framework can be a powerful technique to derive optimal management decisions for endangered species. We built a spatially realistic stochastic metapopulation model for the Mount Lofty Ranges Southern Emu-wren (*Stipiturus malachurus intermedius*), a critically endangered Australian bird. Using discrete-time Markov chains to describe the dynamics of a metapopulation and stochastic dynamic programming (SDP) to find optimal solutions, we evaluated the following different management decisions: enlarging existing patches, linking patches via corridors, and creating a new patch. This is the first application of SDP to optimal landscape reconstruction and one of the few times that landscape reconstruction dynamics have been integrated with population dynamics. SDP is a powerful tool that has advantages over standard Monte Carlo simulation methods because it can give the exact optimal strategy for every landscape configuration (combination of patch areas and presence of corridors) and pattern of metapopulation occupancy, as well as a trajectory of strategies. It is useful when a sequence of management actions can be performed over a given time horizon, as is the case for many endangered species recovery programs, where only fixed amounts of resources are available in each time step. However, it is generally limited by computational constraints to rather small networks of patches. The model shows that optimal metapopulation management decisions depend greatly on the current state of the metapopulation, and there is no strategy that is universally the best. The extinction probability over 30 yr for the optimal state-dependent management actions is 50–80% better than no management, whereas the best fixed state-independent sets of strategies are only 30% better than no management. This highlights the advantages of using a decision theory tool to investigate conservation strategies for metapopulations. It is clear from these results that the sequence of management actions is critical, and this can only be effectively derived from stochastic dynamic programming. The model illustrates the underlying difficulty in determining simple rules of thumb for the sequence of management actions for a metapopulation. This use of a decision theory framework extends the capacity of population viability analysis (PVA) to manage threatened species.

Key words: Australia; conservation; decision theory; metapopulation; optimal landscape reconstruction; Southern Emu-wren; *Stipiturus malachurus intermedius*; stochastic dynamic programming; threatened species.

INTRODUCTION

Decision theory has made important contributions to resource management through the use of more qualitative techniques like hierarchical ranking methods (Ralls and Starfield 1995) and quantitative techniques such as stochastic dynamic programming, SDP (Mangel and Clark 1988, Clark and Mangel 2000). Stochastic dynamic programming has been used to find the best ways of harvesting populations (Johnson et al. 1997, Milner-Gulland 1997, Spencer 1997, Hanson and Ryan 1998), releasing a biological control agent (Shea

and Possingham 2000), maintaining ecosystem diversity (Richards et al. 1999), translocating individuals between two populations to ensure persistence (Lubow 1996), and conserving a spatially structured population (Possingham 1996).

In this paper, we use SDP and a presence–absence metapopulation model to explore the landscape reconstruction strategies for the Southern Emu-wren (*Stipiturus malachurus intermedius*), a critically endangered Australian bird. We consider and rank three management options: creating corridors, enlarging patches, and creating new patches. This is the first application of SDP to explicitly look at landscape reconstruction for a real metapopulation.

The Southern Emu-wren is made up of eight subspecies, one of which is confined to the Mount Lofty Ranges of South Australia (Schodde and Mason 1999).

Manuscript received 26 June 2001; revised 21 March 2002; accepted 26 July 2002. Corresponding Editor: J. A. Wiens.

⁵ Address correspondence to M. I. Westphal, The Ecology Centre, University of Queensland, St. Lucia, Queensland 4072, Australia. E-mail: mwestphal@zen.uq.edu.au

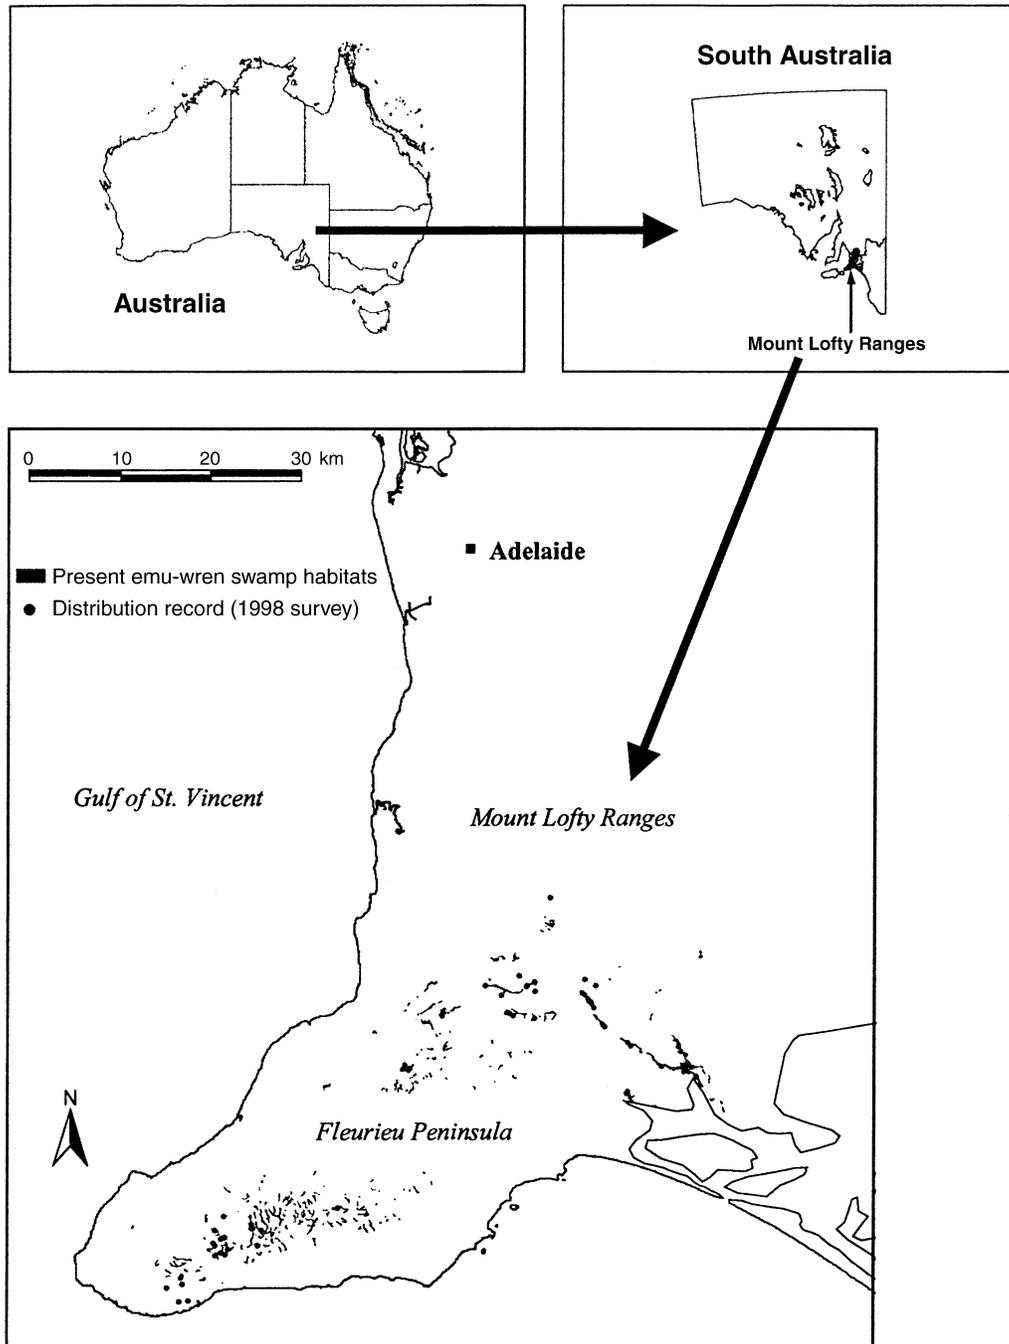


FIG. 1. The location of Southern Emu-wren swamp habitat and distribution. The southernmost points that do not correspond with swamps are heath habitat (Deep Creek Conservation Park).

The critically endangered Mount Lofty Ranges Southern Emu-wren (IUCN Species Survival Commission 1994, Garnett and Crowley 2000) is restricted to swamp and heath habitats of the Fleurieu Peninsula in the southern Mount Lofty Ranges (Fig. 1). About 25% of the original swamp habitat remains, most on private land. Most swamps (~75%) are <5 ha. These swamps are characterized by dense vegetation in the 1–2 m zone

above the ground, composed of tea-tree (*Leptospermum* spp.), sedges (*Lepidosperma longitudinale*), rushes (*Baumea rubiginosa*, *Baumea teragona*, *Juncus* spp.), and ferns (*Blechnum minus*, *Gleichenia microphylla*, *Pteridium esculentum*), often above a layer of peat (Litley and Cutten 1994). The swamps are found along low-lying creeks, perched on hillsides and spring fed, or at the bottom of gullies.

The emu-wrens exclusively use the dense swamp vegetation to feed and nest in, not utilizing the surrounding pasture or eucalypt woodland matrix. Emu-wrens are poor fliers, and only engage in short bursts of flight. They have been called “button grass moths,” and their flight has been compared to that of a dragonfly (Schodde 1982). Thus, the emu-wren is thought to be a very poor disperser (Little and Cutten 1994).

The metapopulation ecology framework (Nicholson 1933, Andrewartha and Birch 1954, den Boer 1968, Levins 1969, 1971, Hanski 1998, 1999), in which individual patches “wink” in and out, but the spatially disjunct network persists as long as the patches are recolonized as quickly as they experience local extinctions, has developed into an important lens through which to understand spatially structured populations. However, it cannot help us to make sound management decisions without being embedded in a decision-making tool. For example, although metapopulation theory tells us that increasing the area and connectedness of habitat patches can attenuate the risk of extinction, it does not give us an ability to rank our decisions.

With a finite pool of resources and a certain exigency to immediately execute the most beneficial course of strategies for an endangered species, managers need decision theory to adjudicate among potential options. For conservation biology to mature as a “triage” discipline, a quantitative framework must be developed that can explicitly incorporate costs and benefits of management options and weigh them accordingly, particularly with lacunas in data and a highly uncertain, stochastic world. This represents the integration of population viability analysis (PVA) with optimization tools (Beissinger and Westphal 1998, Possingham et al. 2002)

In this paper, we first construct a discrete-time Markov chain model for the presence-absence dynamics of the metapopulation. For many endangered species, there is often a paucity of empirical data on density, and management decisions need to be made expediently without the necessary resources to gather detailed population demographic information. Stochastic patch occupancy models (SPOMS), which model only extinction and recolonization events in patches, ignoring population dynamics, can be formulated from snapshot data of patch occupancy (Hanski 1998, 1999, Moilanen et al. 1998). Second, we use stochastic dynamic programming to evaluate various conservation strategies, such as the enlargement of patches, the creation of new patches, and the construction of corridors between patches. This enables us to explore the optimal solution sets for different population states and patterns of metapopulation occupancy.

METHODS

Formulation of the stochastic patch occupancy model (SPOM)

We illustrate our methods using a Southern Emu-wren metapopulation in the Fleurieu Peninsula of South

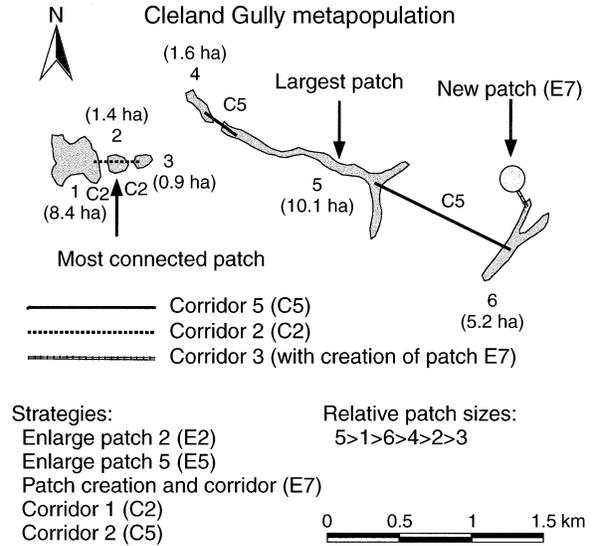


FIG. 2. The Cleland Gully Southern Emu-wren metapopulation, showing size and location of patches and corridors.

Australia (the Cleland Gully metapopulation; Fig. 2). We have partial survey information on patch occupancy from 1993 and 1998.

Our formulation of the SPOM follows that of Day and Possingham (1995). Let the population state of the metapopulation at any time, t , be the set of n patches that are occupied at time t ($\leq n$ total patches). We can represent this as an n -dimensional vector,

$$\mathbf{u}(t) = [u_1(t), u_2(t), \dots, u_n(t)]$$

where the elements are binary variables, $u_i(t) \in \{0, 1\}$. If $u_i(t) = 0$, then patch i is empty; if $u_i(t) = 1$, then patch i is occupied.

Now, we can construct an extinction matrix, \mathbf{X}_n of dimension $2^n \times 2^n$, which represents the probabilities of transitions from any population state through extinction alone. The probability of going from any population state \mathbf{u} to \mathbf{v} by extinction in one time step is given by the product

$$x_{\mathbf{u}\mathbf{v}} = \prod_{i=1}^n I_1(u_i, v_i) \tag{1}$$

where the function $I_1(u_i, v_i)$, the probability that patch i goes extinct in a transition from \mathbf{u} to \mathbf{v} , is

$$I_1(u_i, v_i) = \begin{cases} 1 - E_i & \text{if patch } i \text{ remains occupied} \\ E_i & \text{if patch } i \text{ goes extinct} \\ 1 & \text{if patch } i \text{ remains empty} \end{cases} \tag{2}$$

and E_i is the patch-specific extinction probability. For instance, for a three-patch system, if \mathbf{u} is $[1, 1, 0]$ and \mathbf{v} is $[1, 0, 0]$, then, $x_{\mathbf{u}\mathbf{v}} = (1 - E_1)E_2$. The full matrix, \mathbf{X}_3 , is given in Table 1.

Similarly, we can construct a recolonization matrix, \mathbf{R}_n . The probability of going from any population state

TABLE 1. Extinction matrix for a three-patch system, where E_i is the patch-specific extinction probability. The order of the population states is: [0,0,0], [0,0,1], [0,1,0], [0,1,1], [1,0,0], [1,0,1], [1,1,0], [1,1,1].

\mathbf{I}	0	0	0	0	0	0
E_3	$1 - E_3$	0	0	0	0	0
E_2	0	$1 - E_2$	0	0	0	0
E_2E_3	$E_2(1 - E_3)$	$(1 - E_2)E_3$	$(1 - E_2)(1 - E_3)$	0	0	0
E_1	0	0	0	$1 - E_1$	0	0
E_1E_3	$E_1(1 - E_3)$	0	0	$(1 - E_1)E_3$	$(1 - E_1)(1 - E_3)$	0
E_1E_2	0	$E_1(1 - E_2)$	0	$(1 - E_1)E_2$	0	0
$E_1E_2E_3$	$E_1E_2(1 - E_3)$	$E_1(1 - E_2)E_3$	$E_1(1 - E_2)(1 - E_3)$	$(1 - E_1)E_2E_3$	$(1 - E_1)E_2(1 - E_3)$	0

\mathbf{u} to \mathbf{v} by recolonization alone in one time step is given by the product

$$r_{\mathbf{uv}} = \prod_{i=1}^n I_2(u_i, v_i) \tag{3}$$

where the function $I_2(u_i, v_i)$, the probability that patch i becomes recolonized in a transition from \mathbf{u} to \mathbf{v} , is

$$I_2(u_i, v_i) = \begin{cases} 1 - C_i & \text{if patch } i \text{ remains unoccupied} \\ C_i & \text{if patch } i \text{ becomes recolonized} \\ 1 & \text{if patch } i \text{ remains occupied} \end{cases} \tag{4}$$

and C_i is the patch-specific colonization probability.

Assuming that the order of events is first extinction and then recolonization, the elements of the transition matrix, \mathbf{A}_n , represent the probability of going from any population state \mathbf{u} to \mathbf{v} through both extinction and recolonization:

$$a_{\mathbf{uv}} = \sum_{\mathbf{w}} x_{\mathbf{uw}} r_{\mathbf{wv}} \tag{5}$$

where \mathbf{w} is an intermediate state after extinction.

Parameterization of the stochastic patch occupancy model (SPOM)

To parameterize the SPOM, we use the functions for patch recolonization and extinction from the incidence function model, IFM (Hanski et al. 1996, Hanski 1998, 1999, Moilanen and Hanski 1998, Moilanen et al. 1998). Assume that the patch-specific extinction probabilities are inversely related to patch area:

$$E_i = \min \left[\frac{e}{A_i^p}, 1.0 \right] \tag{6}$$

where A_i is the patch area (in square meters) for patch i , and e and p are two parameters that scale area to extinction. We ignore the possibility of a rescue effect (Brown and Kodric-Brown 1977), in which immigration from neighboring patches leads to a decreasing risk of extinction. To reduce the number of parameters, we assume that the extinction probability is 1.0 for patches $< 2000 \text{ m}^2$ (0.2 ha), which is reasonable, considering that estimates of the size of a pair's territory vary from several hectares to < 0.5 ha (MLR Southern Emu-Wren Recovery Team 1998).

We build a function for the probability of patch i becoming recolonized by first developing an expression for the relative number of dispersers reaching the patch, M_i . If we assume that individuals have an exponentially declining probability of reaching a patch with distance, then the relative number of dispersers reaching a patch is

$$M_i = \beta \sum_{j \neq i}^n \exp(-\alpha d_{ij}) p_j A_j \tag{7}$$

where p_i equals 1 for occupied patches and 0 for empty patches, d_{ij} is the distance (total gap distance, using "stepping stone" dispersal) between patches i and j , α is a dispersal mortality parameter, β is a parameter that scales connectivity to the number of migrants reaching a patch, and n is the number of patches in the metapopulation. The area term effectively weights the number of dispersers originating from a patch by its size.

Let the per patch recolonization probability be sigmoidally related to the relative number of migrants reaching a patch:

$$C_i = \frac{M_i^2}{M_i^2 + \gamma^2} \tag{8}$$

where γ is a location parameter. The squared term gives the curve a sigmoidal shape (Hanski et al. 1996, Hanski 1998, 1999, Moilanen and Hanski 1998, Moilanen et al. 1998). The parameter β effectively is subsumed within γ^2 . One can envisage other functional forms, but it is reasonable to assume that the recolonization probability will be related to the relative number of dispersers reaching a patch in a manner more complex than linearly.

We evaluated three different dispersal scenarios (Fig. 3). Under low dispersal conditions, we assumed α to be 0.015 m^{-1} , which for a distance of 200 m gives a relative dispersal probability of only 0.1. This comes from the fact that emu-wrens apparently are unable to cross a matrix of pasture > 200 m, although their movement would be facilitated by other vegetation such as blackberry, *Rubus* spp. (Little and Cutten 1994). Little and Cutten (1994) considered patches to be isolated if they were > 100 m apart, noting that emu-wrens were never seen flying > 20 m at a time, even when disturbed. Combined with a value of γ set to 5×10^5 , this makes

TABLE 1. Extended.

0	0	}
0	0	
0	0	
0	0	
0	0	
0	0	
$(1 - E_1)(1 - E_2)$	0	
$(1 - E_1)(1 - E_2)E_3$	$(1 - E_1)(1 - E_2)(1 - E_3)$	

small patches (~1 ha) essentially unable to recolonize neighboring patches (Table 2).

Under the baseline (medium dispersal) scenario, α is 0.0046 m^{-1} , which gives a relative probability of dispersal of 0.1 for 500 m. All patches within 100 m of each other have at least a 10% chance of recolonizing each other, and the first cluster of patches (1,2,3) has a roughly 5% chance of recolonizing the second cluster (4,5) of patches when no corridors are present and at their initial sizes. Finally, under the high-dispersal scenario, α is 0.0023, giving a relative dispersal probability of 0.1 for 1000 m. Each patch has at least a 10% chance of recolonization, as long as one patch is occupied. The value of γ is the same for all three scenarios. Although there is uncertainty in the emu-wren recolonization ability, it would fall with this range of possibilities.

To get an estimate for α when corridors are present, we fitted a negative exponential distribution to dispersal data from banding studies of emu-wrens in the nearby Nangkita metapopulation in the Fleurieu Peninsula (Pickett 2000). In total, 19 interpatch movements were recorded among seven patches, which are effectively connected via vegetative corridors, giving a value of $\alpha = 0.0015 \text{ m}^{-1}$.

Although we have survey data from 1993 and 1998, some patches were not surveyed in both years. Assuming the most optimistic extinction scenario, we used maximum likelihood techniques (Moilanen 1999) to derive ρ , the parameter scaling patch area to extinction rate, giving a value of 0.6. We varied this parameter by $\pm 50\%$. The extinction and recolonization parameters represent plausible scenarios.

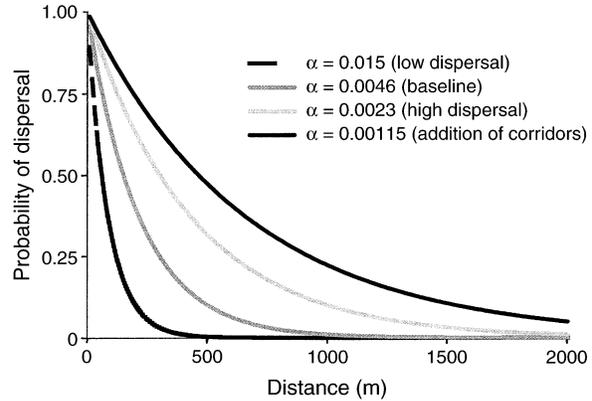


FIG. 3. The dispersal curves for the three different scenarios considered and with the addition of corridors.

The objective and stochastic dynamic programming equations

Stochastic dynamic programming is a state-based, backwards iteration method that explicitly computes the optimal strategy at any time step, based on the state of the system and the time horizon, assuming that optimal strategies are subsequently made (Mangel and Clark 1988, Clark and Mangel 2000). We assume that the goal of the optimization is to maximize the probability that the metapopulation will remain extant at the end of a time period, T . From the SPOM, the total number of population states (patterns of patch occupancy) is $S = 2^n$ (Day and Possingham 1995). For this emu-wren metapopulation, there are seven patches, including the newly created patch, so $S = 128$.

To implement the SDP, we must assign a value to the state of the system when the strategy is evaluated. We let the value of the metapopulation be 1.0, if at least one patch is occupied at the end of the time period (the metapopulation is extant), and 0 otherwise. Suppose there exist k conservation strategies that can be implemented at each time step for the emu-wren, which include the enlargement of patches, the creation of new swamp patches, and connections between patches via corridors. Finding the optimal strategy involves backwards iteration. If we know the value of the metapopulation state at the terminal time, T , then we can cal-

TABLE 2. Extinction (Extinc.) and recolonization (Recolon.) rates under the various parameter scenarios considered.

Patch no.	Baseline		High dispersal		Low dispersal		High extinction		Low extinction	
	Extinc.	Recolon.	Extinc.	Recolon.	Extinc.	Recolon.	Extinc.	Recolon.	Extinc.	Recolon.
1	0.11	0.17	0.11	0.49	0.11	0.03	0.33	0.17	0.03	0.17
2	0.31	0.71	0.31	0.84	0.31	0.45	0.56	0.71	0.18	0.71
3	0.41	0.69	0.41	0.85	0.41	0.32	0.64	0.69	0.27	0.69
4	0.29	0.73	0.29	0.86	0.29	0.42	0.54	0.73	0.16	0.73
5	0.10	0.13	0.10	0.49	0.10	0.02	0.31	0.13	0.03	0.13
6	0.14	0.02	0.14	0.24	0.14	0.00	0.38	0.02	0.05	0.02

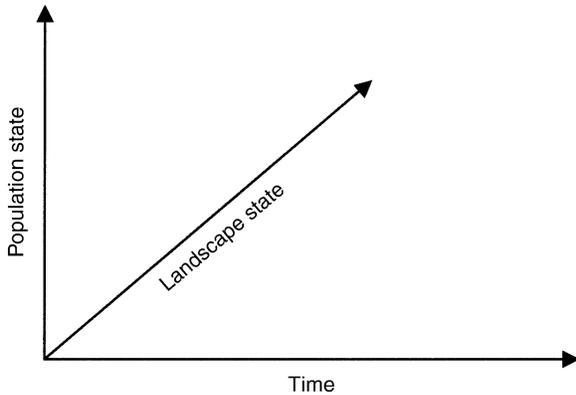


FIG. 4. The output of the SDP (stochastic dynamic programming) model. For every landscape state, population state, and time, an optimal strategy, k_b , is calculated.

culate the value of any state in the previous time step, which is given by the dynamic programming equation (Mangel and Clark 1988, Clark and Mangel 2000):

$$V(\mathbf{u}, \mathbf{z}, t, T) = \max_k \left[\sum_{\mathbf{v}=1}^S a_{\mathbf{u}\mathbf{v}}(\mathbf{y}) V(\mathbf{v}, \mathbf{y}, t + 1, T) \right] \quad (9)$$

where $V(\mathbf{u}, \mathbf{z}, t, T)$ is the value of population state \mathbf{u} and landscape state \mathbf{z} at time t ; T is the terminal time; $a_{\mathbf{u}\mathbf{v}}(\mathbf{y})$ is an element of the transition matrix $\mathbf{A}_n(\mathbf{y})$, the probability of going from population state \mathbf{u} to \mathbf{v} , given that we implement the k th management strategy with the metapopulation initially in the landscape state \mathbf{z} . The landscape state is the combination of possible patch areas and the presence-absence of corridors, and \mathbf{y} represents the new landscape state after the k th strategy has been implemented, $\mathbf{y} = F(\mathbf{z}, k)$. Because each management strategy changes the patch areas and/or connection of patches, each combination of landscape state and population state has a different transition matrix for each management strategy. Therefore, the value term must be a function of both the population and landscape states. Because we set the value to 1 if the metapopulation remains extant and 0 if it goes extinct, $V(\mathbf{u}, \mathbf{z}, 0, T)$ gives the exact extinction probability of the metapopulation for the \mathbf{u} population state and \mathbf{z} landscape state at the initial time. The SDP equation makes no assumptions about the functional forms of extinction and recolonization and simply gives the strategy that will maximize the persistence of the metapopulation at any time and for any state, given that optimal decision are made in the future. In most cases for metapopulations, we are interested in modeling only the occupancy, and the main currency of interest is the probability of extinction for the entire metapopulation. One can think of the result of the SDP analysis as a three-dimensional matrix (Fig. 4): at any point in time, landscape state, and population state, the model derives an optimal strategy.

One can use this technique to determine an optimal trajectory of strategies. Because we are interested in the long-term optimal strategy, i.e., what management action we should perform now to minimize the long-term extinction, we have back-stepped to stationarity, that is, the point where the decision matrix, whose elements are the optimal strategy for each population and landscape state, no longer changes with time. This provides the strategy that maximizes long-term persistence of the metapopulation. For this metapopulation, it was found at 90 yr.

The management strategies for the Southern Emu-wren

We evaluate six landscape reconstruction strategies for the metapopulation (Fig. 2): increasing the area of the largest patch; increasing the area of the most connected (and smaller) patch; creating a corridor from the largest patch to its two closest neighboring patches; creating a corridor from the most connected patch to its neighboring patches; creating (and subsequently enlarging) a new habitat patch with a corridor connecting it to the nearest patch; and executing no strategy.

An existing habitat patch, which is degraded habitat and unsuitable, has been chosen as the site for the patch creation. Although the number of possible points across the landscape where one could model creating new habitat is infinite, sites of degraded patches are the most likely places with suitable habitat to be restored. Habitat for the emu-wren only occurs in topographic depressions, which would hydrologically facilitate the creation of swamps.

In each time step, the manager can execute one strategy; therefore, we have standardized the "cost" of each strategy to be equivalent to ~ 0.9 ha of revegetated area. This is the amount of area needed to create the largest corridor, assuming the requisite corridor width of 10 m (MLR Southern Emu-Wren Recovery Team 1998). For the strategy of creating a new patch, this amount is added, minus the amount of area needed for the construction of a corridor to the nearest patch. We also assume that once a patch is enlarged, the new area will be suitable habit within five years, the time step of the model; hence, there is no time delay. For the patch enlargements, each patch can be enlarged six times. This is justified by the fact that there is only so much available land for habitat reconstruction. So, including the corridors, there is a total of 1372 possible landscape states, or configurations ($7 \times 7 \times 7 \times 2 \times 2$), for which the best strategy is found at each time step and each population state.

Moreover, we have added to the model the chance of catastrophic fire, which is an important ecological process in most Australian ecosystems, although the fire frequency in the area is poorly understood. We have assumed that fire is an independent event in each patch, results in local extinction of the subpopulation, and does not spread to other patches. It occurs just after

TABLE 3. The initial optimal management decision for various population states as calculated from stochastic dynamic programming.

Metapopulation patch occupation	Baseline		High dispersal		Low dispersal		High extinction		Low extinction	
	Low fire	High fire	Low fire	High fire	Low fire	High fire	Low fire	High fire	Low fire	High fire
At penultimate time ($t = T - 1$)										
Two smallest patches	E2	E2	E2	E2	E2	E2	E2	E2	E2	E2
Largest patch	E5	E5	E5	E5	E5	E5	E5	E5	E5	E5
All patches	E2	E2	E2	E2	E2	E2	E2	E2	E2	E2
Two largest patches	E5	E5	E5	E5	E5	E5	E5	E5	E5	E5
At stationary ($t = 0$)										
Two smallest patches	E2	E2	E2	E2	E2	E2	E2	E2	E2	E2
Largest patch	C5	C5	E5	E5	C5	C5	C5	C5	C5	C5
All patches	E2	E2	E2	E2	C2	C2	E2	E2	E2	E2
Two largest patches	C5	E2	C5	E2	C5	C2	E2	E2	C5	C5

Note: Strategies are abbreviated as: E2, enlarge largest patch; E5, enlarge the most connected (smaller) patch; C2, create a corridor connecting the most connected (smaller) patch to neighbors; C5, create a corridor connecting the largest patch to neighbors.

demographic extinction in the model. Fire occurrence is assumed to vary from once every 100 yr to once every 20 yr (extreme case). Once a subpopulation is extirpated by fire, the regeneration time is considered to be five years.

We have compared the 30-yr extinction probability of the optimal strategy from the SDP analysis to that of some fixed-population state-independent strategies, which have the same financial costs. By this we mean that the sequence of strategies is executed regardless of the population and landscape state, something that the stochastic dynamic programming approach explicitly considers. These strategies disregard what patches are occupied at any given time. The “big patch” strategy is simply to expand the largest patch six times. The actions for the “corridor-focused” strategy are in order: construct a corridor connecting the largest patch to its neighbors, construct a corridor connecting the most connected patch to its neighbors, enlarge the largest patch, enlarge the most connected patch, and repeat the last two actions again. The “risk spread” strategy involves first creating a new patch, followed by constructing corridors in the same order as the “corridor-focused” strategy, and then subsequently enlarging the patches in order of decreasing size. To calculate the extinction probabilities for the non-optimal, state-independent sets of strategies, we simply multiply out the transition matrices. Other possible fixed sets of actions can be envisaged, but these represent typical habitat reconstruction paradigms that a manager would be expected to use in the absence of applying SDP. It is worth noting that typical analyses of extinction risk involve Monte Carlo simulations to get an estimate of the probability distribution; however, this is unnecessary, as the Markov chain metapopulation model gives the exact extinction probability.

RESULTS

First we will look at the time and state dependence of the optimal strategies and the trajectories of long-

term optimal decisions, then we will compare fixed, state-independent strategies with the optimal one over a 30-year period. It is cumbersome to show the optimal strategy for each population state, so we have presented four possible benchmark states: only the largest patch occupied, only the two largest patches occupied, all the patches occupied, and only the two smallest patches occupied.

One can see that the optimal solution at the initial time ($t = 0$) varies depending on the population state of the metapopulation and the extinction and recolonization probabilities (Table 3). Some generalities can be gleaned for the extinction and recolonization parameters and the particular patch geometry of the model. When only small patches are occupied, the best initial strategy is to enlarge those patches. However, when a more extinction-resistant, large patch is occupied, connecting it to neighbors via corridors is the optimal initial strategy rather than enlarging it, except when dispersal is quite high. As more patches are occupied, the initial strategy varies with the extinction and recolonization rates. The higher the extinction rate, the more augmenting the size of the smaller patch is the favored strategy, whereas the lower the dispersal, the more building corridors is favored. Under high fire conditions, the extinction probability of all patches is greater. As the number of occupied patches increases, the optimal strategy is either to increase the size of vulnerable patches or to construct corridors, depending on the dispersal ability of the organism. Under very poor dispersal conditions, the shorter corridor is more beneficial. With only a short time frame, patch enlargement is always favored over connecting patches, regardless of the extinction and recolonization parameters. The penultimate time ($T - 1$) is one step before the end of the time horizon, and the optimal strategy at this point is the strategy that minimizes the extinction probability over one time step. These results illustrate the time-dependent nature of the optimal solution.

It must be noted that the optimal initial strategies ($t = 0$) in Table 3 assume that the optimal strategy is executed at every time step in the future (85 yr into the future in this model). This initial strategy is not the same as the best strategy if only one strategy could be performed over the entire time horizon of the model. The essence of SDP is the backwards iteration, which allows one to calculate the optimal strategy now, taking into consideration all possible future states of the metapopulation. The optimal initial strategy does not tell the whole story, nor does it indicate what strategies should be performed in the future; this is determined by the future occupancy of the metapopulation.

The optimal trajectories at stationarity for the baseline model also can be evaluated (Fig. 5A–D). In reality, after each strategy is implemented, there is a chance that the metapopulation will transition to any other population state, as given by the elements of the matrix, A_n . Because it is a Markov process, one cannot speak of there being only *one* optimal trajectory of strategies, because the transitions to new population states are probabilistic. Which management strategy one should execute in the next time depends on which patches are occupied in that time step. Here we show the trajectories of optimal solutions if and only if the metapopulation remains in that population state over time. If an occupied patch is relatively large, it is important to first connect it to corridors and then enlarge it. The smaller a patch, the more critical it is to enlarge it at an earlier stage. From the trajectory, it is clear that there is a very complex interplay of strategies, and the course through time is not easily summarized. It is important to note that the optimal solutions may be quite different for alternative extinction and recolonization parameters, patch geometry, and distribution of patch areas. This is underscored by Fig. 6A and B, which show the optimal trajectory through time if all of the patches remain occupied under the scenarios of low extinction and high dispersal.

For the particular spatial location of the new patch and the parameters of the model, creating a new patch is not a very desirable option, although this would be expected to vary with the spatial arrangement of the metapopulation and the potential location of the new patch. We have not sought exhaustively to vary the location of the new patch around the landscape, but have only used the location of a degraded patch as the focal site. This result makes intuitive sense. Because of the nonlinear form of the relationship between area and patch extinction, the creation of a new patch generally would not be a good strategy unless the magnitude of an area-independent extinction (e.g., catastrophe) were great. In such a case, “spreading the risk” by creating a new patch may be more favorable than augmenting the size of an existing patch, if the patch is close enough to receive dispersers. The occurrence of “do nothing” before all of the active strategies have been exhausted indicates that the difference between

the value of doing nothing and any active strategy is smaller than the numerical precision of our calculations.

Compared to fixed, population state-independent strategies, one can see that the optimal state-dependent set of strategies is superior (Fig. 7), particularly with the population state most vulnerable to extinction. Under low fire conditions, the optimal management set of strategies is 50–80% better than no management, whereas the best state-independent set of strategies is only 10–30% better than no management. Simply increasing the size of the largest patch has the highest extinction probability of the strategy sets considered, although the generality of this result is open to speculation. This is strong evidence for the idea that the sequence of actions, not just the particular management actions, is paramount. It is worth emphasizing that the extinction probabilities between the unmanaged scenario and the scenarios with continued management at each time step would diverge over time (Fig. 8).

Table 4 shows the best strategy to perform if only one strategy could be executed over the 90 years under the baseline scenario, determined by multiplying out the transition matrices. As expected, this is also highly state dependent. Connecting the largest patch to its neighbors is the most optimal strategy when the largest patch is occupied. When only enlargements are considered, enlarging the smallest patch is favored if it is occupied. These are not general results, but are highly contingent on the model parameters and spatial configuration of the metapopulation. Because the extinction vs. area curve is nonlinear, the patch with the greatest marginal value for enlargement (greatest reduction in extinction per unit of area) will have the most benefit in terms of minimizing the metapopulation extinction. It is interesting to note that the best one-time strategy over a time horizon will not necessarily be the same as the initial optimal strategy or the optimal strategy over one time step (the penultimate strategy).

DISCUSSION

We have shown how stochastic dynamic programming can be used to evaluate the optimal landscape reconstruction strategies for a real metapopulation. The technique has advantages over Monte Carlo simulations because it would require an inordinate number of simulations to explore all of the complex combinations of population and landscape states and find the optimal strategy for each time step. If one were simply interested in finding the best strategy to execute in order to minimize the extinction probability over one time step, then Monte Carlo simulation methods would be adequate. One would simulate the metapopulation dynamics for each management strategy a reasonable number of times, perhaps 1000 times, to confidently differentiate the impacts of each strategy. However, the utility of SDP lies in its ability to explicitly find the optimal strategy at some initial time, considering the best de-

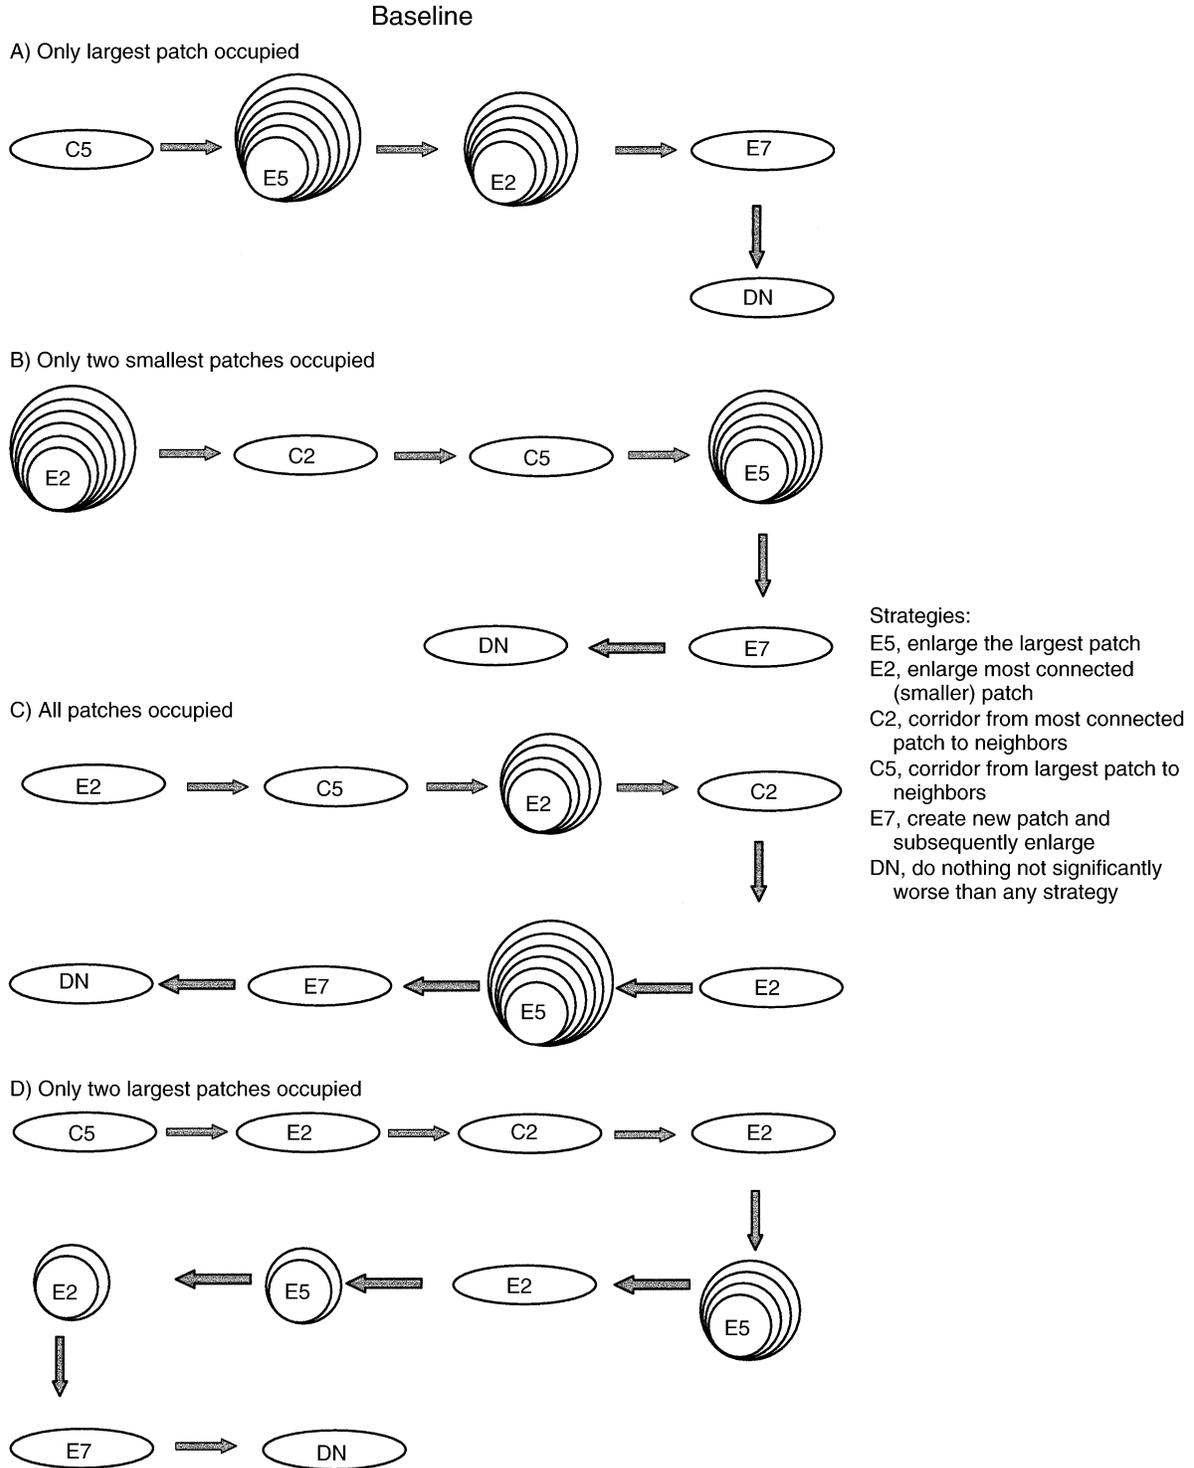


FIG. 5. The trajectories of strategies under the baseline scenario at stationarity for the emu-wren metapopulation under low fire conditions and different patch occupancies (A–D). Each trajectory of strategies represents the long-term optimum for that population state (patch occupancy), assuming the metapopulation remains in that population state throughout. Each circle represents one action. The concentric circles show the execution of the strategy a repeated number of times before the next strategy is implemented in the trajectory. The arrows show the course through time.

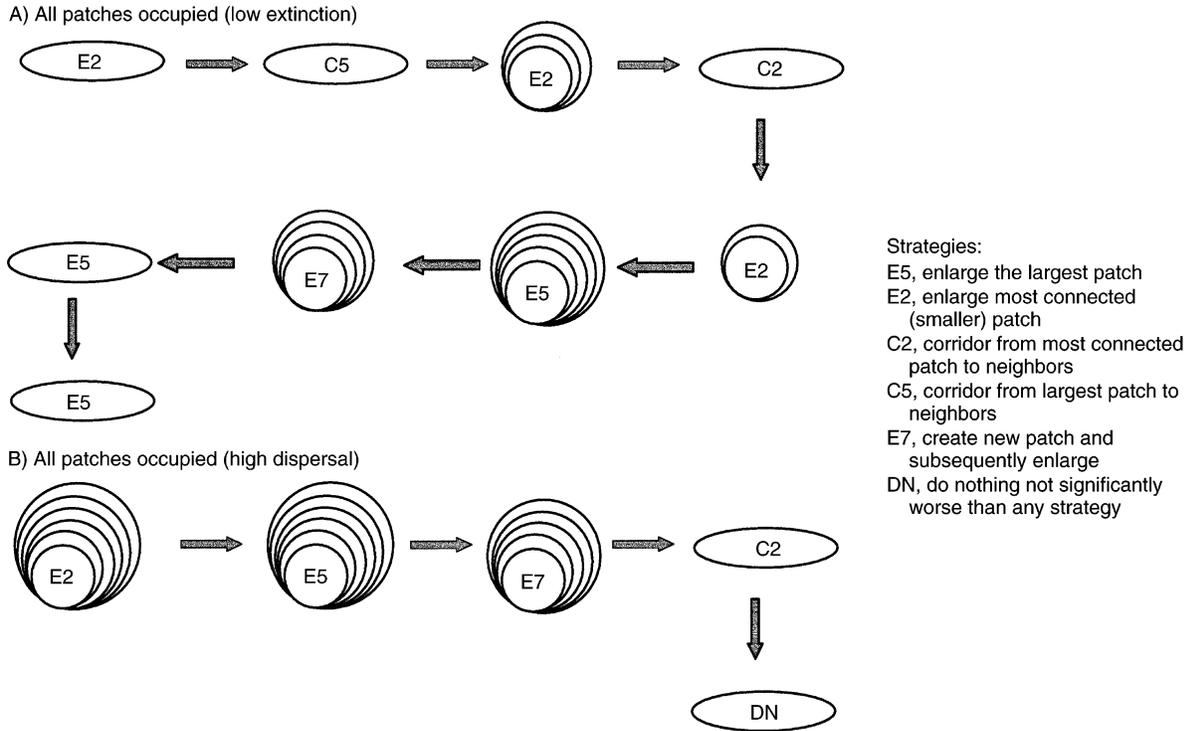


FIG. 6. Trajectories of strategies under the (A) low extinction and (B) high dispersal scenarios at stationarity for the emu-wren metapopulation under low fire conditions. See Fig. 5 for an explanation of the trajectories and circles.

cisions in the future. To find the optimal decision for the emu-wren metapopulation at the initial time using Monte Carlo methods, one would have to simulate all possible strategy combinations over the entire time horizon. To find the optimal decision for all population states, assuming that the metapopulation is in the present landscape state, would require (number of strategy combinations over the entire time horizon) \times (~ 1000 replicates) \times (number of population states) simulations, where t equals the number of time periods. Over 18 time periods, as in our model, this equals an exorbitant number, $(\sim 10^{12}) \times (1000)^{18} \times (128) \approx 1.28$

$\times 10^{34}$ simulations! It is often the case that resources for the conservation of a species are spread over many years, and it is not feasible to allocate area for habitat reconstruction all at once. In such a case, the sequence of actions becomes critical and, as indicated, this can only be efficaciously calculated with stochastic dynamic programming. Moreover, stochastic dynamic programming gives the exact optimum, something that Monte Carlo methods can only approximate.

What lessons for metapopulation management can be learned from this analysis? Firstly, our results highlight the fact that the optimal decision is very state

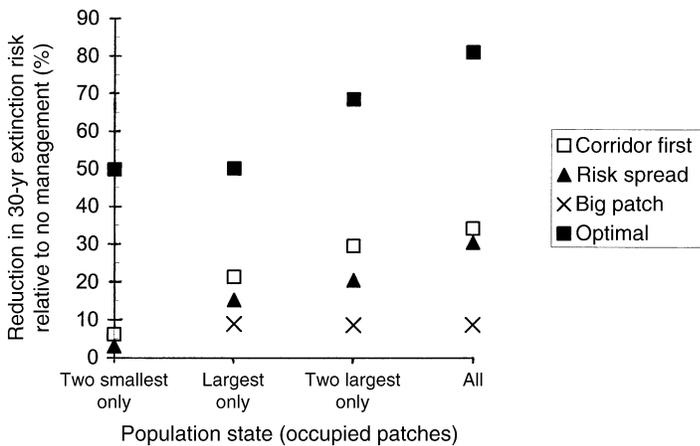
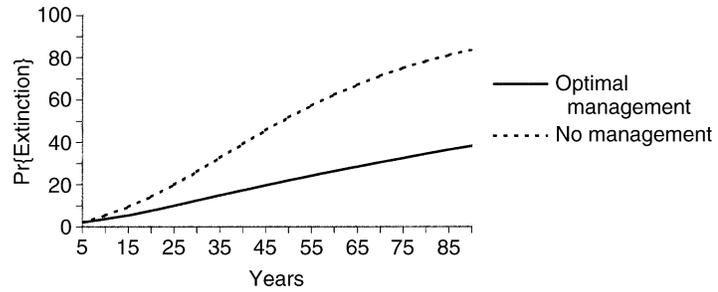


FIG. 7. A comparison of the 30-yr extinction probabilities between the optimal, state-dependent (SDP) set of decisions and fixed, state-independent sets of decisions under low fire conditions.

FIG. 8. A comparison of the optimal set of decisions with no management under the baseline parameter scenario of the model and low fire conditions. The extinction trajectories initially diverge over time, and then the difference remains constant, indicating the diminishing returns of the optimal strategies. The difference would be greater with more ambitious strategies. The extinction probabilities for all non-optimal decisions lie between the two curves.



dependent. Management strategies cannot be executed without cognizance of the pattern of metapopulation occupancy, landscape state, and a good understanding of the extinction and recolonization rates. One cannot hope to manage effectively metapopulations without knowing which patches in the metapopulation are occupied. The complex interplay of strategies and the number of landscape and population state combinations in a spatially realistic metapopulation model make a glib analysis (such as “it is always better to enlarge the biggest patch”) difficult. Secondly, the sequence of actions is critical, as is evident in our comparison of the optimal set of strategies with other sequences of actions. It is not simply what you do, but when you do it that matters. The population state-independent strategies can be significantly worse than the optimal strategy, even when compared to the optimal strategy under a short, 30-year time horizon. This metapopulation management problem is akin to the optimal scheduling problem in Operations Research (Walker 1999), and stochastic dynamic programming is really the only tool that can solve it. With a limited amount of resources available for conservation, spread thinly among many endangered species programs, it is essential that the maximum benefit be obtained. If one has knowledge about the occupancy of the metapopulation and the rates of extinction and recolonization, then the use of the SDP framework will always be superior, as it is the only technique that can give the exact optimal set of strategies. Other techniques or simple qualitative rules can only approximate the optimal solution, and probably crudely at that.

Using an analytical metapopulation model for an idealized landscape, where all patches have equal sizes and rates of extinction and recolonization, Etienne and Heesterbeek (2001) conclude that decreasing the ex-

tinguishing rate is more beneficial than increasing recolonization by the same amount. However, the unit of comparison for managers is: given some fixed cost, which action is most beneficial? Our results indicate that for real landscapes the picture is a complex function of patch geometry, the distribution of patch areas, the extinction and recolonization functional forms (Eqs. 6–8), and most importantly, which patches are occupied. They further recommend, based on varying the patch-specific rates but not including spatial explicitness or variation in patch sizes, that if one is focusing on extinction, one should preferably decrease the lowest local extinction probability. Furthermore, if one is focusing on recolonization, then one should preferably increase the recolonization probability of the patch with the lowest local extinction probability. The former result is discordant with our result in Table 4, (although admittedly this is only one sample metapopulation) and casts doubt on those simple rules for real landscapes. Moreover, with species management it is seldom the case that only one action can be performed; usually a suite of actions can be taken over a given time period. The focus should be on how to optimally schedule those management actions.

In fact, one salient point of our analyses is that it is not easy to glean simple rules of thumb on metapopulation management. Although intuitive general rules merge with regard to the first ($t = 0$) management strategy to execute (Table 3), we have shown that the sequence of actions can be quite complicated, and it is impossible to distill qualitative rules. Conservation has always sought for the holy grail of simple rules, as is evidenced by the SLOSS debate (Diamond and May 1981). This analysis shows that simple, robust rules about the whole sequence of management actions are chimeras, due to the complexity of real landscapes. It

TABLE 4. The optimal management decision for various population states when only one management action can be made (baseline model).

Baseline scenario population state	All strategies	Only enlargement	Only corridors
Only two smallest patches (2, 3) occupied	E2	E2	C2
Only largest patch occupied	C5	E5	C5
All patches occupied	C5	E2	C5
Only two largest patches occupied	C5	E2	C5

Note: Strategy abbreviations are defined in the notes to Table 3.

is necessary to apply a quantitative decision theory framework for every species of concern.

To use this technique, a manager would need to fix the time horizon over which one is interested in minimizing a species' extinction probability. The various strategies need to be delineated, and any constraints on the number of times they can be executed. Here we have limited ourselves to habitat reconstruction, but in practice, many other possibilities could also be incorporated, such as fire management, predator control, and translocation of individuals into patches. In fact, one of us (M. I. Westphal) is using this SDP formulation to look at optimal mowing of pastures for butterfly metapopulations. The manager must then have some estimate of the extinction and recolonization rates, which can be parameterized easily with two snapshots of data (Moilanen 1999). The functional forms of extinction and recolonization can also be altered, but the population dynamics eventually must be described by a discrete-time Markov chain. The model that we developed can be elaborated with more realism, such as including habitat quality, fire spread, time delays, additional species, and habitat succession. Explicit financial costs can even be incorporated into the dynamic programming equation. Here we have standardized the costs of each strategy. After the state space and metapopulation dynamics have been specified, the manager would then survey the occupancy of the metapopulation, and the model would give the exact optimal decision for the current time period and the state of the metapopulation. However, although the technique is quite powerful, it is limited by the size of the state space to rather small networks of metapopulations (Clark and Mangel 2000).

How can our results help to guide policy decisions for the emu-wren? Unfortunately, as our results show, management should not be enacted blindly without regard to exactly which patches are occupied and an understanding of the extinction and recolonization rates. One can choose the most conservative extinction and recolonization scenario presented here, and then after a complete survey of the patches and a specification of the time frame of interest, this model will give the exact optimal strategy. As the parameters become more fine-tuned, this technique can be applied with more confidence for this or other metapopulations of the emu-wren in the Fleurieu Peninsula. It must be emphasized that the optimal sequence of strategies is never known beforehand, because what one should do in the future is a function of which patches will be occupied, a probabilistic process. Our main goal in this paper was to formalize this approach and show how metapopulation management cannot be done without the stochastic dynamic framework. This decision theory technique should be considered an important tool for metapopulation management in general, and provides an "engineering," cost-benefit approach that is needed in conservation (Possingham et al. 2002).

Previous recommendations for landscape planning have not considered the co-dynamics of the population and landscape. One method is to evaluate the relationship between landscape metrics (O'Neill et al. 1988, Turner et al. 1989, Gustafson and Parker 1994, McGarigal and McComb 1995) and species occupancy and design landscapes with a spatial pattern that would have a high probability of species occurrence. Frank and Wissel (1998) make general recommendations about patch placement in the landscape with regard to dispersal range and the correlation length of the metapopulation dynamics. Vos et al. (2001) have proposed using ecologically scaled landscape indices (ESLI) of average patch carrying capacity and connectivity as guides for landscape planning. Hanski and Ovaskainen (2000) have presented a measure of metapopulation capacity, which is the dominant eigenvalue of a matrix of interpatch connectivity measures. The effect of patch alteration or creation can be evaluated by assessing the change in metapopulation capacity. Lambeck (1997) has promoted the focal species approach for landscape reconstruction, in which habitat is created that satisfies the ecological requirements of the most area-limited, resource-limited, dispersal-limited, and process-limited (e.g., by fire) species. Although these methods may be important in providing end goals or final characteristics of landscapes, they provide no tack through time. They decouple landscape dynamics and population dynamics and tacitly assume that the latter occurs at a much faster time scale. Our model is a novel, dynamic approach to landscape reconstruction and is the first attempt to explicitly include these two processes. By analyzing the co-dynamics, we find that the optimal decision logically depends on the current state. Moreover, one cannot talk about optimal metapopulation management only in terms of which actions to undertake, but rather the sequence of those actions. Stochastic dynamic programming is the only method that can effectively solve this problem.

ACKNOWLEDGMENTS

This work is partially supported by a U.S. Environmental Protection Agency STAR Fellowship to M. I. Westphal (91581601-0). We would like to thank the members of the Mt. Lofty Ranges Southern Emu-wren Recovery Program and the Conservation Council of South Australia for their assistance. The Possingham Lab, particularly Andrew Tyre, provided insightful comments on the manuscript.

LITERATURE CITED

- Andrewartha, H. G., and L. C. Birch. 1954. The distribution and abundance of animals. University of Chicago Press, Chicago, Illinois, USA.
- Beissinger, S. R., and M. I. Westphal. 1998. On the use of demographic models of population viability in endangered species management. *Journal of Wildlife Management* **62**: 821-841.
- Brown, J. H., and A. Kodric-Brown. 1977. Turnover rates in insular biogeography: effect of immigration on extinction. *Ecology* **58**:445-449.

- Clark, C. W., and M. Mangel. 2000. Dynamic state variable models in ecology. Oxford University Press, New York, New York, USA.
- Day, J. R., and H. P. Possingham. 1995. A stochastic metapopulation model with variability in patch size and position. *Theoretical Population Biology* **48**:330–360.
- den Boer, P. J. 1968. Spreading the risk and stabilization of animal numbers. *Acta Biotheoretica* **18**:165–194.
- Diamond, J. M., and R. M. May. 1981. Island biogeography and the design of nature reserves. Pages 228–252 in R. M. May, editor. *Theoretical ecology*. Blackwell, Oxford, UK.
- Etienne, R. S., and J. A. P. Heesterbeek. 2001. Rules of thumb for conservation of metapopulations based on a stochastic winking-patch model. *American Naturalist* **158**:389–407.
- Frank, K., and C. Wissel. 1998. Spatial aspects of metapopulation survival—from model results to rules of thumb for landscape management. *Landscape Ecology* **13**:363–379.
- Garnett, S. T., and G. M. Crowley. 2000. The Action Plan for Australian birds 2000. In *Environment Australia*. Department of the Environment and Heritage, Canberra, Australia. [Online, URL: <http://www.ea.gov.au/biodiversity/threatened/action/birds2000/cont.html>.]
- Gustafson, E. J., and G. R. Parker. 1994. Using an index of habitat patch proximity for landscape design. *Landscape and Urban Planning* **29**:117–130.
- Hanski, I. 1998. Metapopulation dynamics. *Nature (London)* **396**:41–49.
- Hanski, I. 1999. *Metapopulation ecology*. Oxford University Press, Oxford, UK.
- Hanski, I., A. Moilanen, T. Pakkala, and M. Kuussaari. 1996. The quantitative incidence function model and persistence of an endangered butterfly metapopulation. *Conservation Biology* **10**:578–590.
- Hanski, I., and O. Ovaskainen. 2000. The metapopulation capacity of a fragmented landscape. *Nature* **404**:755–758.
- Hanson, F. B., and D. Ryan. 1998. Optimal harvesting with both population and price dynamics. *Mathematical Biosciences* **148**:129–146.
- IUCN Species Survival Commission. 1994. IUCN Red List Categories. In *International Union for the Conservation of Nature and Natural Resources*, Gland, Switzerland.
- Johnson, F. A., C. T. Moore, W. L. Kendall, J. A. Dubovsky, D. F. Caithamer, J. R. Kelley, and B. K. Williams. 1997. Uncertainty and the management of mallard harvests. *Journal of Wildlife Management* **61**:202–216.
- Lambeck, R. 1997. Focal species: a multi-species umbrella for nature conservation. *Conservation Biology* **11**:849–856.
- Levins, R. 1969. Some demographic and genetic consequences of environmental heterogeneity for biological control. *Bulletin of the Entomological Society of America* **15**:237–240.
- Levins, R. 1971. Extinction. *Lecture Notes in Mathematics* **2**:75–107.
- Littley, T., and J. Cutten. 1994. Draft Recovery Plan for the Mt. Lofty Ranges Southern Emu-Wren (*Stipiturus malachurus intermedius*). In *Australian Nature Conservation Agency, Endangered Species Program*, Adelaide, Australia.
- Lubow, B. C. 1996. Optimal translocation strategies for enhancing stochastic metapopulation viability. *Ecological Applications* **6**:1268–1280.
- Mangel, M., and C. W. Clark. 1988. *Dynamic modeling in behavioral ecology*. Princeton University Press, Princeton, New Jersey, USA.
- McGarigal, K., and W. C. McComb. 1995. Relationships between landscape structure and breeding birds in the Oregon Coast Range. *Ecological Monographs* **65**:235–260.
- Milner-Gulland, E. J. 1997. A stochastic dynamic programming model for the management of the Saiga antelope. *Ecological Applications* **7**:130–142.
- MLR Southern Emu-Wren Recovery Team. 1998. *Recovery Plan for the Mt. Lofty Ranges Southern Emu-wren (Stipiturus malachurus intermedius): 1999–2003*. In *Conservation Council of South Australia*, Adelaide, Australia.
- Moilanen, A. 1999. Patch occupancy models of metapopulation dynamics: efficient parameter estimation using implicit statistical inference. *Ecology* **80**:1031–1043.
- Moilanen, A., and I. Hanski. 1998. Metapopulation dynamics: effects of habitat quality and landscape structure. *Ecology* **79**:2503–2515.
- Moilanen, A., A. T. Smith, and I. Hanski. 1998. Long-term dynamics in a metapopulation of the American pika. *American Naturalist* **152**:530–542.
- Nicholson, A. J. 1933. The balance of animal populations. *Journal of Animal Ecology* **2**:132–178.
- O'Neill, R. V., J. R. Krummel, R. H. Gardner, G. Sugihara, B. Jackson, D. L. DeAngelis, B. T. Milne, M. G. Turner, B. Zygmont, and S. W. Christensen. 1988. Indices of spatial pattern. *Landscape Ecology* **2**:63–69.
- Pickett, M. 2000. The Mount Lofty Ranges Southern Emu-wren (*Stipiturus malachurus intermedius*) Recovery Program: banding and monitoring 1994–1999. In *Mt. Lofty Ranges Southern Emu-wren Recovery Program*, Conservation Council of South Australia, Adelaide, Australia.
- Possingham, H. P. 1996. Decision theory and biodiversity management: how to manage a metapopulation. Pages 391–398 in R. B. Floyd, A. W. Sheppard, and P. J. De Barro, editors. *Frontiers of population ecology*. CSIRO (Commonwealth Scientific and Industrial Research Organization) Publishing, Melbourne, Australia.
- Possingham, H. P., D. B. Lindemayer, and G. N. Tuck. 2002. Decision theory for population viability analysis. Pages 470–489 in S. R. Beissinger and D. R. McCullough, editors. *Population viability analysis*. University of Chicago Press, Chicago, Illinois, USA.
- Ralls, K., and A. M. Starfield. 1995. Choosing a management strategy: two structured decision-making methods for evaluating the predictions of stochastic simulation models. *Conservation Biology* **9**:175–181.
- Richards, S. A., H. P. Possingham, and J. Tizard. 1999. Optimal fire management for maintaining community diversity. *Ecological Applications* **9**:880–892.
- Schodde, R. 1982. *The fairy wrens: a monograph of the Maluridae*. Lansdowne Editions, Melbourne, Australia.
- Schodde, R., and I. J. Mason. 1999. *The directory of Australian birds*. Collingwood, Victoria, Australia.
- Shea, K., and H. P. Possingham. 2000. Optimal release strategies for biological control agents: an application of stochastic dynamic programming to population management. *Journal of Applied Ecology* **37**:77–86.
- Spencer, P. D. 1997. Optimal harvesting of fish populations with nonlinear rates of predation and autocorrelated environmental variability. *Canadian Journal of Fisheries and Aquatic Sciences* **54**:59–74.
- Turner, M. G., R. V. O'Neill, R. H. Gardner, and B. T. Milne. 1989. Effects of changing spatial scale on the analysis of landscape pattern. *Landscape Ecology* **3**:153–162.
- Vos, C. C., J. Verboom, P. F. M. Opdam, and C. J. F. ter Braak. 2001. Toward ecologically scaled landscape indices. *American Naturalist* **157**:24–41.
- Walker, R. C. 1999. *Introduction to mathematical programming*. Prentice Hall, Engelwood Cliffs, New Jersey, USA.