
Choosing a Management Strategy: Two Structured Decision-Making Methods for Evaluating the Predictions of Stochastic Simulation Models

KATHERINE RALLS* AND ANTHONY M. STARFIELD†

*National Zoological Park, Smithsonian Institution, Washington, DC 20008, U.S.A.

†Department of Ecology, Evolution, and Behavior, University of Minnesota, St. Paul, MN 55108, U.S.A.

Abstract: *We describe two structured decision-making methods—one using a hierarchy of goals and a second using ranking on the sum of weighted criteria—that may be useful for many practical conservation problems, particularly when advisory groups evaluate the output of simulation models. We illustrate both methods by applying them to the problem of choosing a management strategy to address the “mobbing” problem in endangered Hawaiian monk seals. Both methods require estimates of the probabilities of various outcomes, such as a population size of more than 400 seals after 20 years under a specific management regime. We used a simulation model of a small monk seal population to generate these probabilities. Both methods provide an explicit, well-documented, and reproducible decision process that helps justify the decision. Furthermore, they are easy for those untrained in decision analysis to understand and use, they focus discussion on management objectives, they facilitate an examination of trade-offs in the light of multiple and sometimes conflicting objectives, they are suitable for use in workshops, and, at least in our example, they lead to management recommendations that are not highly sensitive to minor changes in probability estimates or other factors.*

Seleccionando una estrategia de manejo: dos métodos estructurados de toma de decisiones para evaluar las predicciones de modelos de simulación estocásticos

Resumen: *En este trabajo describimos dos métodos estructurados para la toma de decisiones—uno que usa una jerarquía de objetivos y otro que usa una clasificación de acuerdo a la suma de los criterios a los cuales se les ha asignado un peso de acuerdo a su importancia relativa—que podrían ser útiles para muchos problemas prácticos de conservación, particularmente cuando grupos asesores evalúan el resultado de modelos de simulación. Ilustramos ambos métodos aplicándolos al problema de la selección de una estrategia de manejo para atacar el problema de comportamiento sexual agresivo (“mobbing”) en poblaciones de focas Hawaianas. Ambos métodos requieren estimaciones de las probabilidades de obtener los distintos resultados, tales como un tamaño poblacional de más de 400 focas luego de 20 años bajo un régimen de manejo específico. Para generar estas probabilidades usamos un modelo de simulación de una pequeña población de focas. Ambos métodos proveen un proceso de decisión explícito, bien documentado y reproducible que ayuda a justificar la decisión. Más aún, son fáciles de entender y usar por personas no entrenadas en el análisis de decisiones, centran la discusión sobre los objetivos de manejo, facilitan el examen de soluciones intermedias considerando objetivos múltiples y a veces opuestos, son adecuados para ser utilizados en talleres de trabajo y, por lo menos en nuestro ejemplo, conducen a recomendaciones de manejo que no son muy sensibles a cambios pequeños en las estimaciones de las probabilidades u otros factores.*

Paper submitted November 12, 1993; revised manuscript accepted September 13, 1994.

Introduction

Conservation issues are complex, and relevant scientific data are often fragmentary. Various constituencies often have conflicting opinions regarding the best course of action, and decision makers are faced with the difficult task of not only choosing a management strategy but also convincing various parties that the choice is sensible.

"Humans are notoriously poor at making choices when there are significant uncertainties, conflicting objectives and complex interactions," so it is important that we try to make our conservation decision processes as explicit and rational as possible (Maguire 1991). The use of an explicit framework to guide conservation decisions can help us choose strategies that are consistent with our goals, data, and beliefs, that document the decision-making process, and that make it easier to defend our decisions.

Operations research and management science, the academic disciplines dealing with scientific approaches to decision making, offer a number of decision-making frameworks that may prove useful for conservation biologists. So far, however, only one technique—decision trees with an expected value criterion—has been widely applied to problems of endangered species (Thibodeau 1983; Maguire 1986, 1987; Maguire et al. 1987, 1988; Maguire & Lacy 1990; Maguire & Servheen 1992; Soulé 1989). This technique may not be the best for all conservation problems (Starfield & Herr 1991; Ralls et al. 1992).

Conservation biologists have devoted considerable effort to simulation models of small populations, which form the basis of population viability analysis (Soulé 1987; Nunnery & Campbell 1993). These models can be used to show the probable results of various management strategies. Much less effort has been devoted to the still difficult problem of choosing a management strategy based, at least in part, on the modeling results. The choice of a strategy is often made at workshops or other meetings where the participants are likely to have different priorities and agendas. We therefore sought structured decision-making methods that groups could use to evaluate the results of simulation models based on a variety of criteria.

Multiple-criteria decision making methods generally assume that the consequences of a decision can be estimated deterministically, however imprecisely. The consequences of choosing a specific management strategy for an endangered species are inherently stochastic. How to handle both uncertainty and multiple-criteria is an open question in multiple-criteria decision-making research (Stewart 1992). Goodwin and Wright (1991) and Stewart (1992) provide reviews of multiple-criteria decision-making methods.

Given an environment of probabilistic outcomes and group decision making, we prefer techniques that promote explicit discussions of priorities and trade-offs and avoid the use of multi-attribute utility (Starfield & Herr 1991). We believe that a decision-making technique for conservation problems should be easy for the nonexpert to use and understand (Stewart 1992). If all concerned parties understand and participate in the decision-making process, they will be more likely to support the conclusions of the analysis. Finally, we believe that techniques should provide robust results—that is, recommendations should not be highly sensitive to minor changes in probability estimates or other factors (Stewart 1992).

We describe two structured decision-making methods that meet these criteria: establishing and using a hierarchy of goals, and ranking on the sum of weighted criteria. Our purposes are to introduce these methods into the conservation biology literature, show how they can be tailored to fit circumstances, and encourage the use of structured decision-making methods in conservation biology.

Methods

We illustrate both methods by applying them to the selection of a management strategy to address the "mobbing" problem in endangered Hawaiian monk seals (*Monachus schauinslandi*) (Starfield et al. 1995). Some monk seal populations appear to be limited by an aggressive behavior known as mobbing, in which adult males injure and often kill adult females and immature seals of both sexes in mating attempts. The frequency of mobbing deaths appears to increase as the population's adult sex ratio becomes increasingly male-biased, although the exact relationship between these two variables is unknown. We consider eight alternative management strategies, described in Starfield et al. (1995) that encompass the range of suggested approaches to the problem: nonintervention, "wait and see," adding females to the population, and removing males from the population.

Both decision-making methods require estimates of the probabilities of various outcomes, such as the probability of a population size of more than 400 seals after 20 years under a specific management regime. We generated these probabilities by running our monk seal model (Starfield et al. 1995) 1000 times to estimate each probability. Because the exact relationship between the adult sex ratio and the frequency of mobbing-related deaths (the mobbing response) is unknown, we modeled five alternative assumptions about the mobbing response. While the details of the model and these five assumptions are not germane to the present paper,

it is important to note that each assumption produced different estimates of the probabilities of various outcomes.

We describe a hypothetical workshop in which the participants have agreed to use our model and the two decision-making methods to evaluate the management strategies on three major criteria: the probable effect on the demography of the seal population, the number of mobbing deaths likely to be prevented, and the cost of management. Participants disagree, however, on the relative importance of the latter two criteria.

Inspection of the probabilities of various outcomes generated by the model under the eight management strategies and one of the five assumptions regarding the mobbing response does not lead to a clear choice of management strategy (Table 1). In general, strategies that are likely to reduce deaths are also likely to require repeated and expensive intervention, while strategies that minimize intervention are correlated with a reduced probability that the seal population will flourish. The two methods described below are designed to structure the evaluation of the alternative strategies.

A Hierarchy of Goals

Linear programming is a standard technique for maximizing (or minimizing) a single objective subject to a suite of linear constraints. Goal programming is an extension of the linear programming algorithm to multiple objectives (see Porterfield 1974; Spronk 1981). We are not concerned here with linear problems or constraints, but we are concerned with multiple objectives. We therefore borrow from goal programming only the way in which it defines a hierarchy of objectives or goals: we list our most important goal first, our next most impor-

tant goal second, and so on (Table 1). There are two key points to this approach. It is essential to choose priorities thoughtfully, because management strategies that do poorly on high-priority goals will be eliminated and not reconsidered, and goals need not be independent.

Participants in our hypothetical workshop agree on the first three goals in Table 1, which aim to promote the viability of the seal population (a small number of adult females and a male-biased adult sex ratio after 20 years are both good indicators of likely future problems in the population). Some participants, however, believe that preventing mobbing injuries and deaths is the next most important consideration, while others are more concerned about the cost of management intervention, and a third group wants to see the seal population flourish irrespective of cost.

The fact that goals need not be independent is the key to compromise: it allows the desires of one group to be met partway and then readdressed after addressing the needs of another group. We make a partial concession to costs in Goal 4, which minimizes the probability of having to intervene (remove or add seals) more than three times in 20 years. This goal is not ideal for those most concerned with costs: they want to intervene no more than once (or at most twice) in 20 years. We then do our best to promote the growth of the seal population and to minimize the chance of future imbalances in the adult sex ratio in Goals 5 and 6. We then address the concern about seals killed in mobbing events. Those expressing this concern would like fewer than 50 such deaths to occur in 20 years. In Goal 7 we meet them partway by minimizing the probability of more than 100 deaths. We return to the question of cost in Goal 8 and this time try to intervene less than twice in 20 years. Finally, in Goal 9, having satisfied all other goals, we try

Table 1. A hierarchy of goals and the probability of achieving each goal under eight alternative management strategies for addressing the mobbing problem in Hawaiian monk seals.

Goal ^a	Probability							
	1	2	3	Strategy Number ^b		6	7	8
1. Total population <100 (minimize)	0.15	0.00	0.00	0.00	0.00	0.00	0.10	0.01
2. Adult females <60 (minimize)	0.45	0.00	0.00	0.00	0.00	0.00	0.31	0.09
3. Sex ratio ≥ 1.4 (minimize)	0.38	0.00	0.00	0.00	0.00	0.00	0.25	0.01
4. Intervened >3 times (minimize)	—	0.03	0.08	0.05	0.11	0.69	0.98	0.09
5. Total population ≥ 400 (maximize)	0.32	0.98	0.98	0.88	0.90	0.92	0.42	0.52
6. Sex ratio <1.0 (maximize)	0.41	0.97	0.98	0.97	0.98	0.97	0.57	0.92
7. Mobbing deaths ≥ 100 (minimize)	0.67	0.00	0.00	0.00	0.00	0.00	0.55	0.20
8. Intervened <2 times (maximize)	—	0.20	0.06	0.33	0.08	0.02	0.00	0.50
9. Mobbing deaths <40 (maximize)	0.06	0.95	0.99	0.58	0.73	0.76	0.07	0.06

^a All outcomes are either at the end of 20 years (such as total population) or totals for the entire 20-year period (such as mobbing deaths).

^b 1 = no action; 2 = remove 10 males each year in which adult sex ratio is more than 1.2; 3 = remove 10 males each year in which more than three mobbing deaths occur; 4 = remove 10 males each year in which more than six mobbing deaths occur; 5 = remove 10 males if no action was taken in previous year and more than three mobbing deaths occurred in present and previous year; 6 = remove five males each year in which more than three mobbing deaths occur; 7 = add 10 females each year in which sex ratio is more than 1.2; 8 = no action for six years, then remove 10 males each year in which more than six mobbing deaths occur.

to keep mobbing deaths below the threshold of 50 deaths in 20 years.

The flexibility offered by this approach might enable workshop participants to agree on a single list of goals. If not, the workshop could produce more than one list, each list reflecting a different point of view. The analysis would proceed with multiple lists and would show how sensitive the solution is to differences in priorities.

Suppose the participants in our workshop have agreed on the prioritized list of goals shown in Table 1. At this point we depart from classical goal programming and develop a simple filter process for reaching a decision. We begin with all eight strategies as viable options and test them using the probabilities generated by the model against each of the goals in order (Table 1).

Strategies 1 and 7 have the highest probabilities that the final total population will be less than 100. We suppose that the discussion concludes that a 0.15 probability is acceptable, so all eight strategies survive the first goal.

Next we look at the second goal and see that Strategies 1, 7, and 8 have higher probabilities that the final number of adult females will be less than 60. At this stage we might decide to filter out Strategies 1 and 7 but leave Strategy 8 in play. Continuing in this way, discussing the questionable strategies at each stage, dropping some and retaining others, leads to the results in Table 2.

In this case, the filtering process produces one strategy, Strategy 2. But there is little to be lost by opting for Strategy 4 rather than Strategy 2 if Strategy 4 is preferable in some way—less expensive or easier to implement. Thus, we can use Table 2 as a basis for discussing the trade-offs between strategies in terms of issues that are not addressed by the simulation model.

Looking back up the hierarchy of goals in this way helps meet a criticism of this goal-filtering approach, that it is too drastic—once a strategy is dropped from consideration it is irretrievably lost (T. J. Stewart, personal communication). Another way of meeting this criticism is what we might call “provisional filtering.” When we are unsure about whether or not to retain a

strategy, we retain it but mark it with a question mark. A questionable strategy might be dropped further down the hierarchy as it emerges that it does not meet later goals, or it might be reinstated if it outperforms all the other retained strategies with respect to several goals further down the hierarchy that collectively seem more important.

So far we have based the entire analysis on the probabilities in Table 1, which were calculated using only one of the five assumptions regarding the mobbing response. Thus, we have ignored one of the dimensions of our decision space: the uncertainty about the exact relationship between the adult sex ratio and the frequency of mobbing-related deaths. We can repeat the probability of calculations and the provisional filtering process, however, using the other four assumptions about this relationship. The results of this process are shown in Table 3. We conclude that Strategy 3 has a good chance of achieving our goals regardless of the exact mobbing response.

Ranking on the Sum of Weighted Criteria (SMART)

Our second method is a well-known member of a class of techniques known as Simple Multi-Attribute Rating Techniques, or SMART, and is described by Goodwin and Wright (1991) and Von Winterfeldt and Edwards (1986). We call this technique “ranking on the sum of weighted criteria,” or SMART ranking. SMART ranks options on the basis of some measure of performance, in this case the probabilities of various outcomes under the alternative management strategies.

The first step, as in the goal-hierarchy method, is to identify our criteria for evaluating alternative management strategies. (To conform with the operations research literature, we use the term criteria rather than goals because we will not prioritize them). The difference between SMART and the goal-hierarchy method is that our criteria here must be independent (Keeney & Raiffa 1976; Keeney 1992). As before, we have three primary criteria to consider: the viability of the seal population, the cost of management, and the number of

Table 2. Results of filtering out unacceptable strategies using the hierarchy of goals and probabilities in Table 1.

Goals	Acceptable Strategies (X)*							
	1	2	3	4	5	6	7	8
1. Total population not <100	X	X	X	X	X	X	X	X
2. Adult females not <60		X	X	X	X	X		X
3. Sex ratio not ≥1.4		X	X	X	X	X		X
4. Do not intervene >3 times		X	X	X	X			X
5. Total population ≥400		X	X	X	X			
6. Sex ratio <1.0		X	X	X	X			
7. Mobbing deaths not ≥100		X	X	X	X			
8. Intervene <2 times		X		X				
9. Mobbing deaths <40		X						

* See Table 1 for strategy definitions.

Table 3. Retained strategies using the provisional filtering process and five assumptions regarding the way in which mobbing deaths increase as the adult sex ratio becomes increasingly male-biased.

Assumption	Strategies*							
	1	2	3	4	5	6	7	8
1		X	X?					
2		X?	X?					
3		X	X					
4			X	X?	X?			
5		X?	X	X?	X?			

* Strategies without question marks are definitely retained; strategies with question marks are possibly retained. See Table 1 for strategy definitions.

seals killed by mobbing, and they are indeed independent. Under each of these major criteria we can specify secondary criteria such as "the total population size after 20 years should be more than 100 seals." The secondary criteria need not be independent and will look very much like the goals in the previous method. In fact, we can use our experience with that method to reduce their number, because the goals relating to the adult sex ratio had no impact on the previous results.

Our workshop has produced the following list of primary and secondary criteria (secondary criteria are numbered sequentially, irrespective of the primary criteria.):

Population Viability (primary)

- (1) The total population after 20 years should be more than 100 seals.
- (2) The number of adult females should not be less than 60.
- (3) The total population after 20 years should be more than 400.

Management and Costs (primary)

- (4) Do not intervene more than three times in 20 years.
- (5) Intervene less than twice in 20 years.

Number of Mobbing Deaths (primary)

- (6) Prevent more than 100 mobbing deaths during 20 years.
- (7) Prevent more than 40 mobbing deaths during 20 years.

The next step is to measure the performance of the strategies with respect to each criterion. For this we require a score on a scale from the least desirable to the most desirable. The probabilities in Table 1 provide a basis for scoring. We can use the probabilities as they are for the criteria corresponding to goals whose probability of occurrence we wished to maximize in the previous analysis; for the others, we can compute the

required score by subtracting the probability in the table from 1.00. To adjust the scores so that they are all on a scale of 0 to 100, we multiply them by 100 (Table 4). The next step in the computations is to assign weights to the criteria. Some useful ways of helping workshop participants arrive at suitable weights are discussed in Chapter 12 of Goodwin and Wright (1991). To keep our total score for each strategy on a scale of 0 to 100, we must ensure that the weights add up to exactly 1.00. Also, because our secondary criteria are not independent and we have different numbers of secondary criteria within primary criteria, we must first assign weights to the primary criteria and the subdivide each of the primary criterion weights among the secondary criteria in that category.

Suppose our workshop participants agree that population viability is twice as important as either management costs or mobbing deaths and that management costs and mobbing deaths are equally important. We therefore assign a weight of 0.5 to population viability, 0.25 to costs, and 0.25 to mobbing deaths. We then divide the 0.5 among the three secondary population-viability criteria. If we let W_j denote the weight accorded to criterion j , we might set $W_1 = 0.2$, $W_2 = 0.15$, and $W_3 = 0.15$. In a similar fashion, suppose workshop participants set $W_4 = 0.15$, $W_5 = 0.10$, $W_6 = 0.15$, and $W_7 = 0.10$.

We let S_{ij} denote the score of management strategy i with respect to criterion j on the 0-to-100 scale and calculate the weighted total score T_i for each strategy i , using the formula

$$T_i = \text{sum over all } j \text{ of } W_j S_{ij}$$

and we rank the strategies from highest to lowest total score. The results (Table 4) agree well with those of the goal hierarchy analysis: management Strategies 2, 3, and 4 are in the top three positions.

The final step is sensitivity analyses, which are easy to perform if a computer program has been written or a

Table 4. Scores for the SMART ranking analysis based on the probabilities in Table 1.

Strategy	Criterion*							Final Score	Rank
	1	2	3	4	5	6	7		
1	85	55	32	100	100	33	6	61	7
2	100	100	98	97	20	100	95	91	1
3	100	100	98	92	6	100	99	89	2
4	100	100	88	95	33	100	58	87	3
5	100	100	90	89	8	100	73	85	4
6	100	100	92	31	2	100	76	77	5
7	90	69	42	2	0	45	7	42	8
8	99	91	52	91	50	80	6	73	6

* 1 = total population after 20 years >100 seals; 2 = number of adult females ≥ 60 ; 3 = total population after 20 years >400; 4 = no more than three interventions in 20 years; 5 = intervene less than two times in 20 years; 6 = prevent >100 mobbing deaths in 20 years; 7 = prevent >40 mobbing deaths in 20 years.

spreadsheet set up to do the calculations. Two types of sensitivity analysis are relevant. First, we can vary the weights assigned to the various criteria: we could, for example, compare two competing weighting systems. (SMART methods tend to be insensitive to all but very large changes in the weights used [Von Winterfeldt & Edwards 1986]; this can be regarded as an advantage or a disadvantage). Second, we can vary the parameters or the weakest assumptions in our model—in other words, the relationship between mobbing-related deaths and the adult sex ratio.

We performed both types of sensitivity analysis, and in almost every case we found Strategies 2, 3, and 4 (not necessarily in that order) in the top three positions and Strategy 7 in the last position. All three recommended strategies call for immediate action and the removal of 10 males from the population (Starfield et al. 1995). They differ only in the event chosen as a trigger for management action—a particular sex ratio or number of deaths due to mobbing in the previous year. As with the goal-hierarchy method, the results were robust, that is, the recommended management strategies were similar regardless of the weights we assigned to the various objectives or our assumptions about the relationship between mobbing deaths and the adult sex ratio in the population. It is worth noting that, at least in this example, using structured decision-making methods to evaluate the modeling results led to more robust conclusions than did simple inspection of the model's output (Starfield et al. 1995).

Concluding Discussion

The use of structured decision-making methods does not necessarily lead to the right decision, but it does help us reach a decision consistent with stated goals (for example, maximizing final population size and minimizing the number of management interventions), data (such as demographic data for monk seals), and beliefs (for example, mobbing-related deaths tend to increase as the population sex ratio becomes increasingly male-biased).

The two decision-making methods we describe share several important properties: (1) they provide an explicit, well documented, and reproducible decision process that helps to justify the decision; (2) they are easy for those unfamiliar with multicriteria decision-making methods to understand and use; (3) they promote and focus discussion on objectives and priorities; (4) they facilitate a structured examination of multiple objectives and the resulting trade-offs; (5) they are flexible and nonprescriptive and are suitable for use in workshops and meetings; and (6) they tend to produce robust recommendations.

Both methods are capable of dealing explicitly with

uncertainty and perceptions of risk, but in different ways. In the goal-hierarchy method, we make comparative decisions about uncertainty when we decide which strategies to accept and which to reject at each stage of the filtering process. In SMART, we have control over how to derive the score for a criterion from the probability of a particular outcome. We chose in Table 4 to make scores directly proportional to the probability. If we were more adverse to risk, we might choose to make the score proportional to the square of the probability.

The two methods also differ in the way they deal with multiple objectives (goals or criteria). The goal-hierarchy method considers goals sequentially, one at a time in order of importance, while the SMART ranking method considers all the criteria simultaneously. The advantage of considering goals sequentially is that we are required to make conscious decisions about priorities. For example, we chose in Table 1 to minimize the risk of bad outcomes first and then to attempt to maximize good outcomes. This procedure reflects a more cautious approach to conservation decisions than that implied by the use of expected values (Starfield & Herr 1991; Ralls et al. 1992). A potential disadvantage of considering goals sequentially is that we risk the ill-advised rejection of a management strategy that scores poorly on one high-priority goal but well on all other goals. This risk can be minimized by using "provisional filters." The risk of considering all criteria simultaneously is just the opposite: a strategy that scores poorly on a single, heavily weighted, vital criterion might nevertheless attain the highest final score. This risk can be minimized or eliminated by sensitivity analyses.

No decision-making technique will be best for all situations. The choice of method will depend upon the nature and structure of the decision problem and the psychology and group dynamics of the people involved in making the decision. Some groups might want to use more than one method, as in our example, and compare the results.

We believe both methods described here will be useful additions to the conservation biologist's tool kit of analytical techniques. We hope the introduction of these methods to conservation biologists will promote a search for additional useful decision-support tools, and an ongoing discussion about which tools are most appropriate for which situations and how to use them to build consensus.

Acknowledgments

We thank Theodor J. Stewart for advice and references on multicriteria decision-making methods, William G. Gilmartin and Timothy J. Ragen for unpublished data and helpful discussions on monk seals, and J. D. Roth for helping to develop the monk seal model. J. Ballou, D.

DeMaster, R. Frankham, T. Ragen, and D. Siniff made helpful comments on the manuscript. The James Smithson Society, The Friends of the National Zoo, and the National Marine Fisheries Service provided financial support.

Literature Cited

- Goodwin, P., and G. Wright. 1991. *Decision analysis for management judgement*. John Wiley and Sons, Chichester, England.
- Keeney, R. L. 1992. *Value-focused thinking: A path to creative decision-making*. Harvard University Press, Cambridge, Massachusetts.
- Keeney, R. L., and H. Raiffa. 1976. *Decisions with multiple objectives: preferences and value trade-offs*. Wiley, New York.
- Maguire, L. A. 1986. Using decision analysis to manage endangered species populations. *Journal of Environmental Management* 22:345–360.
- Maguire, L. A. 1987. Decision analysis: a tool for tiger conservation and management. Pages 475–486 in R. L. Tilson and U. S. Seal, editors. *Tigers of the world*. Noyes Publications, Park Ridge, New Jersey.
- Maguire, L. A. 1991. Risk analysis. *Conservation Biology* 5:123–125.
- Maguire, L. A., and R. C. Lacy. 1990. Allocating scarce resources for conservation of endangered subspecies: Partitioning zoo space for tigers. *Conservation Biology* 4:157–166.
- Maguire, L. A., and C. Servheen. 1992. Integrating biological and sociological concerns in endangered species management: Augmentation of grizzly bear populations. *Conservation Biology* 6:426–434.
- Maguire, L. A., U. S. Seal, and P. F. Brussard. 1987. Managing critically endangered species: The Sumatran rhinoceros as a case study. Pages 141–158 in M. E. Soulé, editor. *Viable populations for conservation*. Cambridge University Press, Cambridge, England.
- Maguire, L. A., T. W. Clark, R. Crete, J. Cada, C. Groves, M. L. Shaffer, and U. S. Seal. 1988. Black-footed ferret recovery in Montana: a decision analysis. *Wildlife Society Bulletin* 16:111–120.
- Nunney, L., and K. A. Campbell. 1993. Assessing minimum viable population size: Demography meets population genetics. *Trends in Ecology and Evolution* 8:234–239.
- Porterfield, R. L. 1974. A goal programming model to guide and evaluate tree improvement programs. *Forest Science* 22:417–430.
- Ralls, K., R. A. Garrott, D. B. Siniff, and A. M. Starfield. 1992. Research on threatened populations. Pages 197–216 in D. R. McCullough and R. A. Barrett, editors. *Wildlife 2001: populations*. Elsevier Applied Science, London.
- Soulé, M. E., editor. 1987. *Viable populations for conservation*. Cambridge University Press, Cambridge, England.
- Soulé, M. E. 1989. Risk analysis for the Concho water snake. *Endangered Species Update* 6:19–25.
- Spronk, J. 1981. *Interactive multiple goal programming*. Martinus Nijhoff Publishing, The Hague, The Netherlands.
- Starfield, A. M., and A. M. Herr. 1991. A response to Maguire. *Conservation Biology* 5:435.
- Starfield, A. M., J. Roth, and K. Ralls. 1995. "Mobbing" in Hawaiian monk seals: The value of simulation modeling in the absence of apparently crucial data. *Conservation Biology* 9:166–174.
- Stewart, T. J. 1992. A critical survey of the status of multiple criteria decision making theory and practice. *Omega* 20:569–586.
- Thibodeau, F. R. 1983. Endangered species: Deciding which species to save. *Environmental Management* 7:101–107.
- Von Winterfeldt, D., and W. Edwards. 1986. *Decision analysis and behavioral research*. Cambridge University Press, Cambridge, England.

Note Added in Proof

In July and August 1994, NMFS biologists removed 22 male seals from Laysan Island and released them in various locations along the main Hawaiian Islands. Five of these seals were tagged with satellite transmitters so that their movements could be followed. As of December 1994, all five had remained near their release sites.

